



Department of Social and Organizational Psychology

A Flexible View of Spontaneous Trait Inferences

Tânia Montenegro Ramos

A dissertation presented in partial fulfillment of the Requirement for the degree of

Doctor in Social and Organizational Psychology

Speciality in Social Psychology

Supervisor:
Doctor Leonel Garcia-Marques,
FPCEUL

Co-Supervisor:
Doctor David Hamilton,
UCSB

August, 2009

The present research was conducted under the Doctoral Project – SFRH/ BD/ 16379/ 2004 – funded by the Science Foundation and Technology of the Portuguese Ministry of Science and Higher Education.

American Psychological Association (PsycINFO Classification Categories and Codes):

3000 Social Psychology

3040 Social Perception and Cognition

2340 Cognitive Processes

2343 Learning and Memory

“We must attempt to figure out the significance of what we have seen, making assumptions about what it all means. This problem is called *inference*.”

Schank (AI Magazine, 1987)

Acknowledgments

It is a great pleasure to have the opportunity to express my gratitude to those who helped me in this journey.

I would like to thank my supervisor Leonel Garcia-Marques. If you think that it is impossible to have a brilliant mind and also an enormous heart, then Leonel is the proof that sometimes it happens. I want to thank him for his guidance, his audacious ideas, and his friendship. Leonel's influence on me goes much beyond the present work. He makes me understand the deepness and complexity of reality, the nonlinear facet of facts. Every time I have a meeting with him I leave with a hunger for knowledge. It has been an enormous honor to be her student.

I would also like to thank my supervisor David Hamilton. I couldn't believe when I knew that David Hamilton, one the founders of Social Cognition, would be my supervisor. He is dedicated, caring, and extremely wise in his comments. I want to thank Dave for the warm reception in Santa Barbara, and for integrating me so well in his research group. I am very happy for the interesting discussions we have, for which I have learned a lot. I also want to thank him for the careful revision of this thesis, especially during a time when he was overloaded with work.

I thank my colleagues and friends from SOCAS: Ana, André, David Kellen, David Rodrigues, Ludmila, Margarida, Marília, Mário, Marta, Pedro, Ricardo, Rita Jerónimo, Rita Silva, Rui, Sara, Sérgio, Sofia, Tammy, Tomás. Thank you all for creating such a comfortable environment of research, and also for your support, friendship, and comments regarding my work. It is a pleasure to work with so wonderful people. Thanks to Marília and Ricardo for the memorable times we shared in Santa Barbara. Special thanks to Marília for her help with formatting issues.

I also want to thank my colleagues from Santa Barbara (Kaat, Joel, Jackie, Nate, Sarah, Debbie, Diana, Devin, Pam, Angie, Zoe, Wes) for their assistance and support in everything that I needed, and for the engaging discussions in Dave's lab meetings. A

special thanks to Kaat Van Acker for her collaboration in Experiment 4, and for her enthusiastic friendship.

I am very grateful to Mécia de Sena for making me feel home, and for enriching the period I passed in Santa Barbara.

Because emotional stability is so fundamental in all this process, I want to thank my closer friends and family. Thanks Mother for your unconditional love, for your support in all my professional choices, and for always making me feel safe. One of the greatest rewards of my work is certainly your admiration. Thank you Dad for loving me, and believing in me. I miss you. Thank you to Tiago, Pedro, Paulo, Telmo, Ana, Amália, and Paula my brothers and “sisters”, and also my best friends.

To Henrique, my partner in life. You know how important you were in all this process. Thank you for your love, comprehension, and encouragement in my worst stressing moments. Thank you for making me smile when I felt tired, for your peaceful silences when I needed to be only concentrated in writing, and for your ability to analyze things in such a balanced way. Much of my strength came from you. I guess now we can finally take some days off...

Lisbon, August 2009
Tânia Montenegro Ramos

ABSTRACT

For a long time, social psychologists focused on understanding how perceivers interpret other people's behavior. One central question has been the identification of the conditions under which perceivers infer personality traits of others. Recent studies (e.g., Winter & Uleman, 1984) suggested that the inference of a trait about an actor from his behavior is a spontaneous process, with characteristics of automaticity. In the present thesis, evidence is presented in favor of a more flexible view of the spontaneous trait inference (STI) process. First, we tested whether STIs are weaker when traits are not so easily inferable from behaviors (Experiment 1). Second, in two sets of experiments we examined whether STIs are guided by coherence requirements. In the first set, we explored whether both STI and spontaneous situational inferences are influenced by the social category of the actor (Experiments 2, 3, and 4). In the second set, we analyzed whether STIs are influenced by the previous presentation of behavioral information about the same actor (Experiments 5 and 6). Finally, the nature of the STI process was further explored by examining whether previous STIs are deliberately used in subsequent tasks (Experiment 7) and by analyzing how previous STIs influence the processing of congruent and incongruent information (Experiment 8). Results, in general, support a flexible view of the STI process. The implications of our work for the debate about the automaticity of the STI process, and for the analysis of the cognitive mechanisms underlying STIs are discussed.

Key Words: spontaneous trait inferences; flexibility; social perception; text comprehension

RESUMO

Desde há muito que os psicólogos sociais se preocupam em compreender os processos envolvidos na interpretação comportamental. Uma questão central é tentar identificar as condições em que os percipientes inferem traços de personalidade acerca dos outros actores sociais. Estudos recentes (e.g., Winter & Uleman, 1984) sugerem que inferir um traço acerca de um actor, a partir do seu comportamento, é um processo espontâneo, com características de automaticidade. Na presente proposta, são apresentados resultados que favorecem uma visão mais flexível do processo de inferências espontâneas de traço (IET). Primeiro, testou-se em que medida as IET são menos prováveis quando os comportamentos não são tão implicativos de traço (Experiência 1). Segundo, em dois conjuntos de experiências, examinou-se em que medida as IET são modeladas por requisitos de coerência. No primeiro conjunto, explorou-se a influência da categoria social do alvo na ocorrência de IET e de inferências espontâneas situacionais (Experiências 2, 3, e 4). No segundo conjunto, analisou-se o efeito da apresentação de informação comportamental prévia acerca do actor na magnitude das IET (Experiências 5 e 6). Por último, explorou-se a natureza das IET. Na Experiência 7, verificou-se em que medida IET prévias são usadas deliberadamente em tarefas subsequentes. Na Experiência 8 analisou-se a influência das IET no processamento subsequente de informação consistente ou inconsistente. Os resultados, em geral, favorecem uma visão flexível das IET. Serão devida as implicações dos resultados para o debate acerca da automaticidade das IET, assim como para a análise dos processos subjacentes às IET.

Palavras-Chave: inferências espontâneas de traço; flexibilidade; percepção social; compreensão de texto

TABLE OF CONTENTS

GENERAL INTRODUCTION	1
<hr/>	
PART I – THEORETICAL BACKGROUND	15
<hr/>	
CHAPTER I – Attribution Theory	17
<hr/>	
1.1. The Problem of other Minds	20
1.2. The Early Contribution of Heider	21
1.3. The Correspondence Inference Theory	22
1.4. Kelley’s Attribution approach	24
1.4.1. Covariation	25
1.4.2. Configuration	26
1.5. The Fundamental Attribution Error and the Correspondence Bias	28
1.6. Why do people prefer dispositional attributions?	32
1.7. Quattrone Anchoring and Adjustment Model	33
1.8. Trope identification-inference Model	33
1.9. Gilbert Three-Stage Model	34
1.10. Questioning the existence of a Fundamental Attribution Error	35
1.10.1. Salience of the dispositional goal	36
1.10.2. Perceived intentionality	39
1.11. The “When” of Attributional Processes	41
1.12. Differentiating Dispositional Inferences and Causal Attribution	44
1.13. Summary of the Chapter	46
<hr/>	
CHAPTER II – Spontaneous Trait Inferences	49
<hr/>	
2.1. Initial Evidence for Spontaneous Trait Inferences	52
2.2. Spontaneous Trait Inferences Paradigms	55
2.2.1. Cue-Recall Paradigm	56
2.2.2. Recognition Probe Paradigm	60
2.2.3. Savings in relearning Paradigm	62
2.2.4. False Recognition Paradigm	64

2.2.5.	Other Paradigms	66
2.3.	The automaticity of STIs	68
2.3.1.	Intentionality	68
2.3.2.	Awareness	72
2.3.3.	Efficiency	76
2.3.4.	Controllability	78
2.3.5.	Conditional Automaticity View	81
2.4.	Spontaneous Trait Transference Effect	87
2.4.1.	Theoretical Accounts of STTs	88
2.4.2.	Limits on STTs	90
2.5.	Other Spontaneous Inferences	93
2.5.1.	Gist Inferences	94
2.5.2.	Situational Inferences	95
2.5.3.	Goal Inferences	97
2.5.4.	Other Spontaneous Inferences	100
2.5.5.	The Multiple Social Inferences View	102
2.6.	Open Questions	103
2.6.1.	Inferences at Encoding Stage	104
2.6.1.1.	The Contamination Problem	104
2.6.1.2.	Encoding vs. Retrieval	105
2.6.2.	Inference about the actor or about the behavior	109
2.6.3.	Association or inference?	111
2.7.	Summary of the Chapter	116
 CHAPTER III – Text Comprehension		 119
3.1.	The Explosion of Inferences	123
3.2.	Scripts approach	125
3.2.1.	Empirical Evidence for the Role of Scripts	127
3.3.	Situation Models	131
3.4.	Constructionist Approach	133
3.5.	Minimalist Approach	138
3.5.1.	Minimalist Approach Development	142
3.5.2.	The Gradual View of Inferences	143

3.6.	General Implications for STI Research	146
3.7.	Consequences of a minimalist approach to a flexible view of STIs	147
3.8.	Summary of the Chapter	149
CHAPTER IV – Present Proposal: A Flexible View of the STI Process		151
<hr/>		
4.1.	State of the Art	153
4.2.	The Present Proposal: A Flexible View of the STI process	154
4.2.1.	Evidence from STI research	154
4.2.1.1.	STIs are affected by stereotypes: A strong argument in favor of the STI flexibility	155
4.2.2.	Evidence from Text Comprehension Research	158
4.3.	First general principle: STIs are more likely when easily available	158
4.4.	Second general principle: STIs work in line with local coherence requirements	159
4.5.	Third general principle: Spontaneous trait inferences vary in a continuum of strength	161
PART II – EMPIRICAL STUDIES		163
<hr/>		
SECTION I – Spontaneous Trait Inferences are More Likely When Easily Available		165
<hr/>		
	Experiment 1: Probe Recognition with Ambiguous Behaviors	167
	Method	171
	Results	174
	Discussion	177
SECTION II – Spontaneous Trait Inferences Work in Line with Local Coherence Requirements		181
<hr/>		
	Experiment 2: Probe recognition with Behaviors and Situational Continuations	183
	Method	188
	Results	190
	Discussion	192

Experiment 3: Reverting the Order of Behaviors and Situations	197
Method	200
Results	201
Discussion	204
Experiment 4: Probe recognition with Trait/Situational Ambiguous Behaviors	209
Method	211
Results	213
Discussion	214
Discussions of Experiments 2-4	216
Experiment 5: Cued Recall Paradigm with Pairs of Behaviors	219
Method	223
Results	224
Discussion	225
Experiment 6: Probe Recognition Paradigm with Pairs of Behaviors	227
Method	230
Results	232
Discussion	233
Discussions of Experiments 5 and 6	234
SECTION III – STIs Vary in a Continuum of Strength	237
<hr/>	
Experiment 7: Trait-cued Recall	239
Method	244
Results	245
Discussion	248
Experiment 8: False recognition and reading times	253
Method	260
Results	263

Discussion	267
Discussions of Experiments 7 and 8	271
PART III – GENERAL DISCUSSION AND CONCLUSIONS	275
<hr/>	
1. Main Findings	277
2. Potential Limitations	279
2.1. Paradigms Weaknesses	279
3. Implications for STI Debates	283
3.1. Automaticity of STIs	284
3.2. Other Spontaneous Social Inferences	284
3.3. Debates concerning Actor vs. Behavior and Association vs. Inference	285
4. A Closer Look into the Three Principles Underlying STIs	287
5. Future Venues	290
5.1. Effects of Trait Accessibility on STIs	290
5.2. Contributions of STI research for Text Comprehension	291
6. Conclusions	292
REFERENCES	293
<hr/>	
APPENDICES	337
<hr/>	
APPENDIX A: Stimulus Material and Instructions of Experiment 1	339
APPENDIX B: Stimulus Material of Experiment 2	345
APPENDIX C: Stimulus Material of Experiment 3	349
APPENDIX D: Stimulus Material of Experiment 4	353
APPENDIX E: Stimulus Material and Instructions of Experiment 5	357
APPENDIX F: Stimulus Material of Experiment 6	369
APPENDIX G: Stimulus Material and Instructions of Experiment 7	375
APPENDIX H: Stimulus Material and Instructions of Experiment 8	381

LIST OF FIGURES AND TABLES

EXPERIMENT 1

Table 1: Means Errors as a Function of Type of Trial	175
--	-----

EXPERIMENT 2

Figure 1: Proportion of errors as a function of probe type and category dimension	191
Figure 2: Response Times for Correct Responses as function of probe type and category dimension	192

EXPERIMENT 3

Figure 3: Response Times for Correct Responses as function of probe type and category dimension	202
Figure 4: Response Times for Correct Responses as function of probe type and category dimension	203

EXPERIMENT 4

Figure 5: Percentage of errors as function of probe type and category dimension	214
---	-----

EXPERIMENT 5

Figure 6: Percentage of errors as function of probe type and category dimension	225
---	-----

EXPERIMENT 6

Table 2: Means of errors as a function of type of pair	233
--	-----

EXPERIMENT 7

Table 3: Free recall and ARC scores as a function of processing goal and cues condition	246
---	-----

GENERAL INTRODUCTION

A couple of days before starting to write the first lines of my dissertation I went to the Lisbon Gulbenkian Park with a friend. We seated ourselves on the grass and started to talk about the last events in our lives. My friend complained that her boss always asks her to work extra hours. Now she could finally take some days off to do what her children prefer the most – go to her house in the south, near the beach. While we were talking I noticed a woman seated in a yoga position with her eyes closed. Nearby a couple was lying down on the grass. He was reading the newspaper, while she was reading a Dostoyevsky book. Close by, I observed a young man juggling balls and a girl dressed all in white walking with her dog. After a while chatting, my friend and I went to a bar to eat something. When we were paying my friend insisted to pay the bill. After a while, the park guard, using what seemed to me an unnecessary harsh voice, warned us that the park was closing and that we should leave immediately.

Just in a couple of hours, what do you think I have learned about the different people I received information from? What have I learned about the young men, the park guard, the woman seated in the yoga position, the girl with the dog, the couple reading, my friend, my friend's children, and her boss? How have I processed and interpreted the different pieces of behavioral information about each one of these individuals? Or, in other words, what have these pieces of information told me about who these people are?

As this example illustrates, the amount of information about others we are exposed to, even in very brief periods of time is impressive. Every day, we communicate and interact with dozens of people, and each one of them performs innumerable actions that reflect their intentions, internal mental states, and personalities. Most of these behaviors are not neutral in meaning. Quite the opposite, each one of them encloses a potential of meaning that can be penetrated by anyone who happens to be observing them.

In addition, as inveterate social animals that we are, nothing captures our attention more than other people. Even if we were on the most wonderful island in the world, contemplating the most amazing sunset, if a person shows up in the scenery, our attention will automatically abandon the sunset and focus on the person. Probably, this capacity to catch our attention reflects how important it is to anticipate other's intentions (e.g., Fiske, Cuddy, & Glick, 2007)

Interactions with other people are numerous, and other people may catch our attention more than anything else, but more important than all, understanding others is crucial to our sense of adaptation. Comprehending the stable causes of behavior allows us to have a sense of control over reality. It allows us to predict how people will behave in the future, and to adjust our own behavior according to those predictions. Also, it gives us the power of influencing others' actions, by trying to interfere with the causes of their behavior.

Ultimately, all our social life, including our ability to maintain a happy wedding, educate our children, establish friendship relations, and also the capacity to get work, money, and power are dependent on our ability to understand others in an accurate manner. It's not surprising that cognitive therapy is based on a model that gives a predominant role to the replacement of maladaptive attributions (e.g., Beck, 1995; Brewin, 1986; Massad & Hulse, 2006; Stratton, 2003; Bradbury & Fincham, 1990; Peterson et al., 1982; Peterson & Seligman, 1984).

Given the large amount of social information that we have to deal with every day, and since the understanding of others is the basis of adapted interpersonal relations, it would make sense that our mind had been shaped by evolution to be sensitive to interpret other people's minds, based on the external signs they send.

Understanding these mechanisms has been one of the main topics of research in the social psychology. According to the most influential authors in the area (Asch, 1952; Heider, 1958), the clarification of the mechanisms underlying the processing and interpretation of other people's behavior would be the key to understanding social life, including our perceptions, judgments, and behaviors towards others.

One of the important questions studied in social psychology is specifically concerned with the understanding of the processes by which perceivers go from the observation of behaviors to inferences about the actor's stable personality traits. Initially, researchers' approach was guided by the assumption that, upon the observation of a behavior, perceivers engage in an attributional analysis in order to examine the conditions and situational constraints surrounding the behavior occurrence. This analysis would provide a basis for perceivers to determine the validity of inferring a trait disposition about the actor. Thus, trait dispositional inferences were viewed as a

relatively late processing stage, conditional upon the analysis of the causal locus of behavior (Heider, 1958; Jones & McGillis, 1976; Kelley, 1967).

However, subsequent formulations suggested that the search of the causal locus of the behavior is not a precondition for dispositional trait inferences to occur. According with this view, dispositional trait inferences can take place without consideration of causal relations (e.g., Hamilton, 1998; Krull, 2001; Smith & Miller, 1983; Winter & Uleman, 1984), and attributional analysis (i.e., exploring the causes of behavior) can also take place independently of making trait dispositional inferences (Krull, 2001; Hamilton, 1998; Hastie, 1984; Jerónimo, 2007). Thus, attributional analysis and dispositional inferences were better conceptualized as distinct mental processes. While attributional processes typically involve systematic and deliberative thinking, dispositional inferences more often occur in a less deliberative way, free from causal considerations.

Despite the invaluable contributions of the initial attribution researchers (Heider, 1958; Jones & Davis, 1965; Jones & McGillis, 1976; Kelley, 1967), and regardless of the fact that posterior attributional models tried to incorporate a less rational picture of the social perceiver (Gilbert, Pelham, & Krull, 1988; Quattrone, 1982; Trope, 1986), the study of the trait dispositional inference process could be carried out independently of the examination of deliberative (even if faulty) attributional processes. In addition, the typical utilization of overt instructions was probably not suitable to fully capture the nature of the dispositional inference process.

With the development of implicit methods of research (e.g., Jacoby & Dallas, 1981; Jacoby, 1983a; 1983b; Warrington & Weiskrantz, 1968, 1970, 1974) and with the concomitant explosion of studies interested in discovering the powers of the unconscious (for relevant examples see Bargh & Chartrand, 1999; Dijksterhuis & Van Knippenberg, 1998), soon researchers in this domain tried to understand whether perceivers could infer trait dispositions of actors from their behaviors, without intention or awareness.

In line with this reasoning, a new program of research, initiated with Winter and Uleman (1984), provided evidence that inferring personality traits of actors from their behaviors is a spontaneous process, since it occurs even when people have no explicit intention to interpret other's behavior. The spontaneity notion had such a strong impact

that studies conducted within this domain are referred to as the spontaneous trait inferences literature (for reviews see Uleman, 1999; Uleman, Newman, & Moskowitz, 1996; Uleman, Saribay, & Gonzalez, 2008).

Subsequent research has been highly focused on testing whether or not the process fulfills all the automaticity criteria (Lupfer, Clark, & Hutcherson, 1990; Moskowitz & Roman, 1992; Newman & Uleman, 1990; Todorov & Uleman, 2003). Certainly, this route of research is appealing, since it can provide controversial findings. Proving the automaticity of the trait inference process would suggest that the initial stages of social perception are relatively outside of our control.

The goal of the present dissertation is to argue for a more flexible view of the spontaneous trait inference (STI) process. We base our work on a review of different literatures (attribution, spontaneous trait inference, and text comprehension literatures) that together give us a clear picture of how the process of STIs has been studied, and provide crucial inputs in order for us to sustain a flexible view of the process.

The dissertation is organized in the following manner. We first present the theoretical background of the studies (PART 1); then we describe the empirical experiments that were conducted in order to test the flexibility hypothesis about the STI process (PART 2); and finally, we finish with a general discussion, in which we consider the implications of our findings, we point out potential limitations in our studies, and we explore new ideas for future research (PART 3).

In the first chapter of the theoretical review (CHAPTER 1) we focus on the attribution literature. The main goal of this chapter is to demonstrate how the dispositional inference process was initially studied, as a stage within the attribution process. We start the chapter by describing the classic contributions of Heider (1958), Jones and Davis (1965) and Kelley (1967). The initial attributional models represented the perceiver as a rational agent, who would search for the causal locus of the behavior before making a dispositional inference about the actor. We continue the chapter presenting the most well known attributional bias, the fundamental attribution error (Ross, 1977), which is the tendency to overestimate the role of dispositional over situational constraints in behavior explanation. Here, we highlight the distinction between the fundamental attribution error (Ross, 1977) and the correspondence bias (Jones & Harris, 1967). After that, we pass to the description of subsequent attributional

models, developed with the intent of portraying a less rational and potentially biased social perceiver: the anchoring and adjustment two stage model (Quattrone, 1982); the identification-inference model (Trope, 1986); and the three-stage model (Gilbert, Krull, & Pelham, 1988; Gilbert, et al., 1988). We then summarize a set of studies (e.g., Krull, 1993; Krull & Erickson, 1995; Quattrone, 1982) that brought into question the generality of the fundamental attribution error, suggesting that social perceivers can be either dispositional or situational biased, depending on the circumstances. After that, we specify the conditions that tend to instigate attributional thinking (e.g., Hastie, 1984; Kanazawa, 1992; Lau & Russell, 1980). We close the chapter by presenting both theoretical and empirical arguments for the view of dispositional inferences and causal attributions as distinct cognitive processes.

In the second chapter (CHAPTER 2) we review and discuss the spontaneous trait inference literature. This literature represents a fundamental change in the way dispositional trait inferences were viewed. Within this framework, dispositional trait inferences are conceived as occurring every time a behavior is comprehended, independently of the perceiver intention or awareness. This chapter describes the evolution of the spontaneous trait inference research and outlines, specifically, how the literature has been focused on testing the automaticity of the process. We begin the chapter by presenting the first empirical findings, provided by Winter and Uleman (1984), which indicated that the trait inference process could be a spontaneous process (i.e., non-intentional and unconscious). We then describe the different main paradigms that have been used to explore the spontaneity of the trait inference process: the cued recall (e.g., Winter & Uleman, 1984); the recognition probe (e.g., McKoon & Ratcliff, 1986; Uleman, Hon, Roman & Moskowitz, 1996); the savings in relearning (e.g., Carlston & Skowronski, 1994; Carlston et al., 1995); and the false recognition (Todorov & Uleman, 2002; 2003) paradigms. After that, we explore whether the process of spontaneous trait inferences fits the different criteria for automaticity. Specifically, we present studies that test whether the process is non-intentional (e.g., Winter & Uleman, 1984); unconscious (Winter & Uleman, 1984; Todorov & Uleman, 2002); efficient (Lupfer et al., 1990; Uleman, Newman, & Winter, 1992; Winter, Uleman, & Cunniff, 1985; Todorov & Uleman, 2003), and controllable (Uleman et al., 1996; Uleman & Blader, 2001; see Uleman, Blader, & Todorov, 2005). We then analyze the ambiguities

that have surrounded the definition of automaticity, and we describe how most cognitive processes cannot be described as completely automatic (e.g., Bargh, 1989, 1994). We continue the chapter by considering the spontaneous trait transference effect (Skowronski, Carlston, Mae, & Crawford, 1998; Carlston & Skowronski, 2005), which represents the surprising fact that perceivers spontaneously infer a trait of an actor, even when the actor is merely describing a behavior of a third person. This effect has important implications to the discussion about the mechanisms underlying STIs (e.g., Crawford, Skowronski, Stiff, & Scherer, 2007). We then briefly summarize the other types of spontaneous inferences reported in the literature (e.g., Ham & Vonk, 2003; Gernsbacher, Goldsmith, & Robertson, 1992; Hassin, Aarts, & Ferguson, 2005), given that trait inferences are not the only type of inferences that occur spontaneously. We close the chapter by citing the fundamental debates that are still open in the spontaneous trait inference literature. These debates concern (a) whether the paradigms used in the literature can prove that the trait inference occurs during encoding (e.g., Corbett & Doshier, 1978; McKoon & Ratcliff, 1981, 1986, 1992; Wyer & Srull, 1989), (b) if the trait inference that is said to occur describes the actor or is a summary for the behavior (Bassili, 1989a, 1989b; Claeys, 1990; Higgins & Bargh, 1987; Newman & Uleman, 1993; Park, 1989; Uleman et al., 1993; Whitney, Davis, & Waring, 1994), and (c) whether the process underlying the effects reported in the literature is an inference or an association (Brown & Bassili, 2002; Carlston et al., 2005; Carlston et al., 2007; Carlston et al., 2008; Todorov & Uleman, 2004; Uleman, 1999).

In order to clarify the previous controversies, as well as the debate about the automaticity of the trait inference process, it is important to examine the vast literature that investigates the types of inferences that occur during text comprehension. In the third chapter (CHAPTER 3), we focus in the inferences that occur during text comprehension. The main goal of this chapter is to specify how the inferential process generation has been modeled and to present data that explore the types of inferences that are said to occur online during reading. Initially, we present a problem of simulating artificial systems of comprehension: the “explosion of inferences” problem (Rieger, 1975; Schank, 1975). After that, we describe the role of *Scripts* in making available knowledge-based expectations and in avoiding the explosion of inferences problem (e.g., Schank & Abelson, 1977; Sharkey, 1986). Then, we consider the

hypothesis that readers construct *Situation Models* that guide inference generation (Johnson-Laird, 1983; van Dijk & Kintsch, 1983; Zwaan & Radvanski, 1998). After that, we describe two different perspectives about inference generation: the constructionist theory (Graesser, Singer, & Trabasso, 1994) and the minimalist approach (Mckoon & Ratcliff, 1992). According to the constructionist theory, inferences are made in order to achieve a coherent and meaningful representation of what is being processed. In opposition, the minimalist approach states that the occurrence of inferences is restricted to inferences that are easily available in memory and to those that are needed for text local coherence. Based on the minimalist framework, we propose three principles underlying STIs occurrence.

In a last chapter (CHAPTER 5), we describe our proposal in detail. We start by outlining how the existing literature is consistent with a flexible view of the STI process. A crucial argument in favor of this view is the recent finding that STIs can be inhibited when an incongruent stereotype is associated with the actor of the behavior (Wigboldus, Dijksterhuis, & van Knippenberg, 2003). After that, we specify the three principles proposed to underlie STI occurrence: (1) STIs are more likely when easily available, (2) STIs work in line with local coherence requirements; and (3) STIs vary in a continuum of strength. Finally, we explain how the three proposed principles are addressed in our studies.

In the second part of the dissertation (Part II), the flexibility of the spontaneous trait inference process is empirically tested. Three different lines of studies were developed in order to test each one of the proposed principles. First, the principle according to which STIs are more likely when easily available was examined in Experiment 1 (SECTION I). We used the recognition probe paradigm (McKoon & Ratcliff, 1986; Uleman et al., 1996) with ambiguous behaviors (i.e., behaviors that imply simultaneously two personality traits). The typical use of behaviors that clearly imply a trait may have contributed to an overestimation of the automaticity of the process. We predicted that STIs would be less likely to occur when the ambiguous behavior is presented with a neutral category, than when the same behavior is presented with a social category that favors one of the trait interpretations of the behavior.

Second, we examined whether STIs occurrence is modeled by local coherence requirements. According with this principle, STI are facilitated when they increase the

local coherence of the text and are inhibited when they provoke text local incoherencies. This principle was explored in two different sets of experiments (SECTION II). On one hand, we tested whether the occurrence of both spontaneous trait inferences and spontaneous situational inferences (SSIs) is influenced by the social category that is ascribed to the actor of the behavior (Experiments 2-4). According with the local coherence principle we expected that when the category-label is inconsistent with the behavior, spontaneous trait inferences from the behavior would be less likely to occur, compared with a condition in which the same behavior is presented with a consistent category-label. In contrast, we expected that spontaneous situational inferences would be *more* likely to occur when the category-label is inconsistent than when it is consistent with the behavior (Experiment 2). In Experiment 3, we tested whether the presentation order of the dispositional and situational components of the sentences would influence the occurrence of both types of spontaneous inferences. Other goal of Experiment 3 was to try to rule out any explanation of the previous results (Experiment 2) based on mere category activation effects. Finally, in Experiment 4 we used sentences that were equally likely of eliciting a spontaneous situational inference and a spontaneous situational inference, in order to test whether the magnitude of both spontaneous inferences is dependent on the behaviors that are being processed.

The local coherence principle was tested in another set of experiments (Experiments 5 and 6). If STIs are cognitively malleable, and if they work in line with coherence requirements, then their occurrence should be lower every time the trait inference is disruptive to the process of comprehension. Accordingly, spontaneous trait inference magnitude should not only be influenced by well learned stereotypes (e.g., Wigboldus et al., 2003) but also by previous behaviors of the same actor. In Experiment 5 we tested this idea by using a new paradigm that combines features of the traditional impression formation paradigm with features of the cued recall paradigm (Winter & Uleman, 1984). In this paradigm, participants were presented with pairs of behaviors and later were asked to recall the second behavior of each pair. During recall, either the first behavior of each pair, or the traits implied by the second behaviors, were provided as cues. Our hypothesis was that when a behavior is preceded by a congruent behavior during encoding, STIs from that behavior will be stronger, in comparison with a condition in which the same behavior is preceded by an incongruent behavior. Since

STIs are expected to be stronger after the presentation of a congruent behavior, we predicted that trait-cues would facilitate recall in congruent-pair conditions. In contrast, STIs should be weaker when the behavior is preceded by an incongruent behavior. However, we predict that in this case the second behavior is more likely to be compared with the previous received information. As a result, first-behaviors would become more associated with first behaviors. Because of that, first-behaviors cues should be more efficient for recall in incongruent-pair conditions. In Experiment 5, we tried to prove the on-line nature of the previous findings by applying the probe recognition paradigm (e.g., McKoon & Ratcliff, 1986) with pairs of behaviors.

Third, the two final experiments were conducted in order to test the principle that STIs vary in a continuum of strength (SECTION III). In line with this principle, stronger inferences are characterized by the possibility of being deliberately accessed in posterior moments, and by acting as expectations about the future behavior of the actor. In Experiments 7 and 8 intentional and spontaneous inferences were compared in relation to these two characteristics. Experiment 7 explored whether STIs may be deliberately accessed in posterior tasks. Participants were presented with 24 sentences illustrative of 4 different personality traits, under memory or impression formation instructions. They were then asked to recall the behaviors, either with the traits implied by the behaviors as cues, or without cues. Both the number of behaviors and the level of trait-clustering were recorded. Since under impression formation conditions participants tend to have access to previous inferred traits (Ferreira, Garcia-Marques, Ramos, Hamilton, & Jerónimo, in press; Hamilton, Katz, & Leirer, 1980), we predicted that trait-cues would make no difference under impression formation conditions. In contrast, because memory participants fail to monitor trait inferences that occur during encoding, they probably are not able to spontaneously use the inferred traits as retrieval cues, *unless* the cues are explicitly provided. Thus, we predict that trait cues will tend to lead to better recall and greater trait clustering than no cues condition, but only in the memory condition. Experiment 8 tested whether a spontaneous trait inference is mentally represented as a dispositional characteristic of the actor. If that is the case, this mental representation should work as an expectation about the actor and therefore it should guide subsequent processing of information. We applied the false recognition probe paradigm typically used in the literature (Todorov & Uleman, 2002) and added a

reading time measure. We wanted to explore whether the trait inferences that are made when participants read trait-implying sentences influence the time taken to read subsequent congruent and incongruent behaviors. If the trait inferred is perceived as a characteristic of the actor, subsequent inconsistent behaviors should be more difficult to process, and this should result in longer reading times for inconsistent behaviors (Experiment 8).

Results from this set of experimental studies favor a more flexible view of the spontaneous trait inference process. Such a view makes sense from an adaptation perspective, since the mind would not be forced to constantly make trait inferences about the personalities of others, even under circumstances where those inferences turn out to be inadequate or irrelevant. In the General Discussion (PART 3), we summarize our main findings and discuss potential limitations of our studies. We also examine the implications of our findings to the main debates in the STI field, and suggest future venues for research.

Our Approach

When we write a dissertation about “inference processes”, we are faced with the fact that people can give different meanings and connotations to the concept of “inference”. Different academics have approached the inference concept in different ways. Logicians and philosophers study the inference process by exploring the rules of valid inferences (e.g., Braine, 1978; Quine, 1970), statisticians develop formal models of inferences, as for example the Bayes Theorem (Bernardo & Smith, 1994), and researchers on artificial intelligence are concerned with developing automatic inference systems (Russell & Norvig, 2003).

We differentiate ourselves from these approaches because we are concerned with the actual rules that people use when making inferences, and not with prescriptive models of reasoning. Influences from other domains are certainly useful, but only in the sense that they could help us pursuing this goal. Another difference is that we are specifically interested in studying the inference processes that people use when interacting with *other people*, and the consequences that those inferences have for social perception.

Within this psychological framework, we follow a social cognitive approach. The socio-cognitive perspective is a tradition of research within social psychology that emerged in the 1970s (Fiske & Taylor, 1991; Hamilton, 2005; Higgins, 2000; Kunda, 1999). Social cognition makes use of metaphors and techniques from cognitive psychology (see Eysenck & Keane, 2005), and its main goal is to provide precise descriptions of the mental representations and processes underlying social judgment and behavior. Although the development of social cognition was influenced by cognitive psychology, the findings from social cognition provide unique insights into the nature of fundamental cognitive processes that would remain uncovered if not explored under social meaningful conditions (e.g., Wyer & Srull, 1994).

The central idea of this approach is that in order to explain a psychological phenomenon we need to specify the mental operations that are in its origin. We embrace this mentalist view (Sperry, 1993). The social perceiver is an active processor of information that constructs perception (Bruner, 1957; Halberstadt & Niedenthal, 2001), elaborates the information that is stored in memory (Craik & Tulving, 1975), and retrieves the stored information in more heuristic or systematic ways (e.g., Chaiken & Trope, 1999; Garcia-Marques & Hamilton, 1996). Thus, there is no isomorphic relation between the external reality and the way that reality is constructed in our mind. The mind has its own rules in order to deal with the external environment, and it is according to these rules that behavior is driven (Kelley, 1967). Thus, if we want to understand the social perceiver we have to understand the way our mental architecture encodes, processes, stores, and retrieves information.

In the present work, we apply a socio-cognitive perspective in order to analyze the cognitive processes underlying the processing and interpretation of social behavior.

PART I

THEORETICAL BACKGROUND

CHAPTER I

ATTRIBUTION THEORY

The initial study of dispositional trait inferences was integrated within the study of attributional processes. Because dispositional inferences were at first studied as part of attributional analysis, a consideration of this literature is crucial for introducing and critically evaluating our own program of research. In the present chapter, our aim is not to exhaustively describe this literature but to providing a summary of the main theoretical perspectives and empirical findings. This will give us an idea of how trait inferences were initially approached by social psychologists researchers.

We start the chapter by introducing the focus of attributional research: explaining how people explain others' behaviors. Then, we describe the view of the classic attributional authors (Heider, 1958; Jones & Davis, 1965; Jones & McGillis, 1974; Kelley, 1967)), according to which dispositional trait inferences occur only after a careful analysis of the causes of the behavior. We then look into the Fundamental Attribution Error (Ross, 1977) and to the Correspondence bias (Jones & Harris, 1967; Gilbert & Malone, 1995). Both these terms have been generally used to refer to the tendency to overestimate the personal determinants of behavior. Despite the distinction proposed by Hamilton (1998), according to which the fundamental attributional error is a bias that occurs as part of attributional analysis, and the correspondence bias is an error that reflects the tendency to make trait inferences from behaviors, both terms have been used interchangeably. As a consequence, the same evidence is usually presented to support the occurrence of both errors. Later in the text, we describe the attribution models of Quattrone (1982), Trope (1986), and Gilbert et al. (1988). These models already conceive the possibility that trait inferences occur before attributional processes, as well as their automatic occurrence (especially the model of Gilbert, see Gilbert et al., 1988). However, studies continue to make use of overt instructions, which turned difficult a clear examination of the proposed mechanisms. We continue the chapter by briefly considering the conditions that tend to instigate attributional causal processes. As we will see, factors that increase causal reasoning (for example, unexpected events) do not necessarily increase trait inference dispositions. This is line with the view that both processes are distinct. We finish the chapter by delineating the arguments in favor of the differentiation between attributional causal processes and dispositional inferences. The emergence of the spontaneous trait inference literature is easily comprehended in light of these arguments.

1.1. The problem of other Minds

The question of how we are able to understand the meaning of other people's behavior, without having access to their internal mental states, is a classic philosophical problem (Wittgenstein, 1953). The psychological distance that exists between any two people is well captured by the subjective feeling that "no one can really understand me" that probably everyone already has had.

The more extremist philosophical approach to this problem is represented by the solipsism "the problem of other minds", which expresses, in a general way, the difficulty of proving the existence and nature of other people's minds, based on physical evidence (Ames, 2005, Farah, 2008; Minar, 1998).

However, we *do* explain and try to understand the reasons for other people's behaviors, and we do it all the time. Despite the impossibility of getting inside other people's heads, we use our own mind to find explanations for why they act as they do. As Proust (1930) states, we could ultimately say that "Man is the creature...that knows his fellows only in himself". Inside our mind we search the meaning of other's behaviors.

Imagine that a given person X observes a person Y performing a behavior. There is no guarantee that what Y wanted to show with the behavior will be accurately captured by X. Person X can, however, based on general knowledge stored in memory as well as on specific previous knowledge that he may have about Y, try to explain the causes and reasons of the behavior. The way X interprets the behavior of Y will determine how Y will be perceived and evaluated, and will also influence the way X will behave toward Y in the future. For example, if someone fails in an academic exam it will make a big difference whether we explain the failure by laziness, low intellectual abilities, or due to the extreme difficulty of the exam. The causes that we attribute to the behavior are crucial in the way we perceive the person.

Clarifying the mechanisms that we use when trying to explain other people's behavior has been the focus of attribution literature (for reviews see Harvey & Weary, 1984; Kelley & Michela; 1980; Malle, 2004; Hewstone, 1989). Heider (1944, 1958) was one of the first authors that explored the mechanisms underlying attributional reasoning.

1.2. The Early Contribution of Heider

Fritz Heider (Heider, 1920, cited by Malle, 2004; Heider, 1944, 1958) is considered the father of attribution research. His book *The Psychology of Interpersonal Relations* had inspired innumerable theoretical and empirical works, and some authors still claim that “the theoretical richness of Heider’s book has not yet been fully explored” (Gollan & Witte, 2008, p.189).

According to Heider, because we are successful in using the attribution process in our lives in order “to build up and support the constancy of our picture of the world” (Heider, 1958, p.92), common sense psychology could be very useful in the construction of a scientific theory of social behavior¹. Thus, based on the analysis of how lay people perceive and explain different behaviors in everyday relations, Heider proposed a causal model underlying human interactions.

Heider started by making a distinction between personal causality and impersonal causality. When an actor intentionally makes something, he talks about *personal causality*. In contrast, when an actor causes something, without intention, he refers it as *impersonal causality*.

Heider focused mainly in how people tend to perceive and explain actions that are intentionally performed in order to reach a goal (i.e., personal causality). According to Heider, two components characterize an intentional action: the motivation (“*try*” component) and the capacity to perform the action (“*can*” component). Both factors are necessary for action, but neither of them is sufficient. Specifically, we can have the required abilities to achieve a goal, but if we have no motivation, the goal will never be reached. In the same way, no matter how hard we try, we will not be able to perform an action if we don’t have the required abilities.

In addition, Heider specified the constituents of “can” and “try”. According to Heider, can “refers to a relation between the person and the environment” (Heider, 1958, p. 87). That is, the power to do something is not only dependent on personal characteristics, but also on the properties of the environment, which can either facilitate or prevent the action. If the required personal abilities are stronger than the opposing

¹ This view was in opposition to the dominant perspective, according to which theory should contradict common-sense (Epstein, 1997; McGuire, 1997). From another viewpoint, however, common-sense approach is not surprising since, as Kelley asserts, “it is precisely common-sense with which attribution theory is concerned” (Kelley, 1973, p.108).

environmental forces, then it is said that the person “can” cause the event. On the motivational side, Heider defined the constituents of “try” as the *intention* to carry the action and the *exertion* that is applied. These concepts refer to whether a person tries to perform the action, and how hard she tries.

Heider’s formulations were so rich and diverse that we can find in his book cues to almost all ideas that were explored in subsequent studies within the attribution and trait inference domains. This may be one of the reasons why his book had the highest number of citations in 2006 and 2007 since its publication (see Reizenzein & Rudolph, 2008). Among his numerous insights, he calls attention to the role of stereotypes in the type of inferences that are made about the person or about the situation (“If a child, for example, successfully bakes a cake or reads a book, we conclude that the recipe or book was easy,” p. 90), and for the influence that previous knowledge about the actor has on causal thinking (“if we have a very low opinion of a person’s ability then any success will be attributed to luck”, p.91). For that reason, Heider is cited many times, also in the present work.

Jones and Davis (1965) and Kelley (1967) were some of the authors that explicitly extended Heider’s ideas. They developed theories more prone to empirical tests, fulfilling the scientific aim of substituting untested generalities by “piecemeal, detailed, and verifiable results” (Russell, 1914, p.14)

1.3. The Correspondence Inference Theory

Jones and Davis (1965; see also Jones & McGillis, 1976) proposed a theory to explain the conditions under which a perceiver makes a *correspondent* dispositional inference upon the observation of a behavior. A correspondent inference occurs when a perceiver attributes the disposition that describes the behavior to the actor. The disposition implied by the behavior *corresponds* to the one that is attributed to the person.

According to the correspondent theory there are certain conditions that favor the attribution of a correspondent inference. First, the more the actor is perceived to have free choice, the more likely it is that a correspondent inference occurs. That is, if we perceive the act as intentional, then we assume that the behavior reflects something

about the personality of the actor. Jones saw this criterion as a “relatively self-evident proposal” (Jones, 1979, p. 219).

Once the behavior is perceived as intentional, a second factor that influences the attribution of a correspondent trait is whether the behavior is unexpected (i.e., whether the action contradicts social norms, and has effects that are non-common to other actions)². The more a behavior disconfirms our expectancies about the target (target-based expectancies) and about the social group of the target (category-based expectancies) the more we assume that the behavior reveals something unique about the actor, and the more likely is that we make a correspondent attribution (see Weisz & Jones, 1993, for differences between target and category based expectancies).

Subsequent research had, however, provided contradictory findings in relation to both of these assumptions. In relation to the criterion of intentionality, despite attributions being stronger in conditions of perceived free choice, even when participants explicitly know that the actor had no choice to perform the behavior, they still attribute an attitude to the actor according to the behavior observed (Jones & Harris, 1967; Jones, Worchel, Goethals, & Grumet, 1971; Snyder & Jones, 1974).

Concerning the role of expectancies, some studies provided data against the hypothesis that expectancy-inconsistent behaviors lead to stronger dispositional inferences. First of all, it has been shown that perceivers have a set of different strategies to disregard inconsistent information. For example, ambiguous inconsistent information can be reinterpreted in a manner consistent with stereotypes (Kunda, 1990). In addition, even when behaviors are interpreted as truly inconsistent, despite leading to longer processing times (Bargh & Thein, 1985; Stern, Marrs, Millar, & Cole, 1984), and instigating stronger attributional processing (Clary & Tesser, 1983; Hamilton, 1988; Hastie, 1984; Pyszczynski & Greenberg, 1981), they are less likely to be coded in terms of dispositional traits (Jerónimo, 2007). This pattern reflects the interesting fact that more attributional processing is not equivalent to stronger dispositional attributions. Consistently, existing data suggest that inconsistent behaviors are more likely to be attributed to situational factors (Crocker, Hannah, & Weber, 1983).

² In initial formulations, Jones and Davis talk about these two principles: social undesirability and non-common effects. However Jones and McGillis (1976) summarize these two aspects in the general assumption that behaviors that are incongruent with expectancies are more informative about the target.

A different line of research, however, has shown that behaviors that are consistent with social prescriptions are seen as non-informative about the idiosyncratic personality of the individual, and lead to weaker dispositional inferences (Ajzen, 1971; Jones, Davis, & Gergen, 1961; Lay, Burrton, & Jackson, 1973). For example, if John is smiling, talking and drinking in a party we have no basis to infer that John is an extroverted person, since those behaviors are normative in a party social situation. This pattern may be taken as an indirect indication that, in relation to social norms, inconsistent behaviors actually result in stronger dispositional inferences about the target.

Together, this set of studies clearly indicates that expectancies have different origins (Kassin, 1979; Miller & Turnbull, 1986). *Actor-based expectations* make perceivers expect consistency of the actor's behavior across situations, whereas *situation-based expectations* lead perceivers to expect different people to behave in the same way in specific social settings. Studies suggest that behaviors that are incongruent with expectancies about the target lead to weaker dispositional inferences than do congruent behaviors. In contrast, behaviors that are incongruent with social norms probably result in stronger dispositional inferences than behaviors that are in line with such prescriptions. A question open to empirical scrutiny concerns how these different expectancies interact to influence the probability of inferring dispositional characteristics.

.1.4. Kelley's Attribution approach

Harold Kelley had an unquestionable impact in the field of social perception (Garcia-Marques & Garcia-Marques, 2003), being probably the author that developed in more detail Heider's ideas. Kelley (1967, 1971, 1972; see also Kelley, 1973) distinguished between two different circumstances in which a person is faced with the problem of finding a causal explanation. One is when the person has multiple pieces of information about the event, and the other is when the person has information from a single observation of the event. According to Kelley, in order to search for a causal explanation, different strategies will be applied in these two situations. In the former case *covariation* principles are applied, while in the latter case *configuration* information based on causal schemata is used.

1.4.1. Covariation

When people have multiple observations about the event, causal judgments are assumed to be based on analyses of covariances, similar to the analysis of variance method (ANOVA). Those analyses provide information about which effects covary with which causes. Despite the inadequacy of the ANOVA analogy (see Garcia-Marques, 1988; Jaspars, Hewstone, & Fincham, 1983), the level of elaboration of the model is remarkable.

Kelley proposes that people use three types of covariation information. *Distinctiveness*, whether the effect occurs only in the presence of a specific entity (high distinctiveness) or whether the effect occurs in the presence of other entities as well (low distinctiveness). For example, does John only laugh when he talks with Susan or does he laugh with anyone he talks to? *Consensus*, whether other people respond in the same way to the same stimulus (high consensus), or whether other people do not respond in the same way to the same stimulus (low consensus). That is, does everyone who talks to Susan laugh, or only John laugh when talks to Susan? And *Consistency*, whether the effect occurs in the presence of the specific cause, in different moments in time (high consistency) or whether the effect doesn't occur in different moments in time (low consistency); that is, does John laugh in every circumstance when he talks to Susan, or only in some circumstances?

The analysis of the eight different patterns of covariation that result from combining high and low instances of the three information sources determines the type of causal explanation provided. Specifically, the pattern of the three types of information determines whether the causal factor is the person, the entity, or the circumstance. For example, in conditions of high distinctiveness, high consensus, and high consistency the behavior is assumed to be caused by the entity (Kelley, 1967). That is, if John only laughs when talks to Susan, if everyone who talks to Susan laughs, and if he laughs in every occasion he talks to Susan, then something about Susan is probably responsible for the behavior. In contrast, a pattern of low distinctiveness, low consensus, and high consistency suggests a personal attribution (see McArthur, 1972). Specifically, if John laughs with everyone he talks to, if he laughs every time he talks with Susan, and if he is the only person who laughs when talking to Susan, then it is something about John that causes the behavior.

However, according to Kelley (1967), people do not take in consideration all the different patterns of covariation, under all circumstances. In conditions in which people lack information, time, or motivation to perform complex computations, they base their causal judgments on *Causal Schemata* based-knowledge (Kelley, 1972).

1.4.2. Configuration

People have a repertoire of causal schemata about how things usually happen in the world. Causal schemata refer to general knowledge about the relation between different causes and effects. According to Kelley (1972), schemas derive from explicit and implicit learning about how events are causally related in the world, and also from the experience that perceivers have on controlling causes and effects.

These causal structures are applied not only when perceivers lack processing resources, but also when causal channels are so clear that covariation analysis is viewed as unnecessary. As Kelley (1972) states, causal schemata enclose simple beliefs about the world as, for instance, the belief that a man who dresses like a woman is homosexual. These beliefs, even if wrong, provide a quick and simple causal explanation for the behavior. Thus, if observed events fit with the schema causal structure, a reasonable causal inference can be made.

In addition, Kelley (1972, 1973) enunciated two causal principles based on causal schemata that people usually apply when think about causes and effects: the discounting principle and the augmentation principle. Discounting reflects the fact that the role of a potential cause is reduced when other cause is likely to be true. For example, if we wake up and observe that the streets are wet, it may be because it has rained or because the streets have been washed. If for some reason we came to know that streets were washed, then it is less likely we assume that it rained. This happens regardless of the independence of the causes (for studies about the discounting principle, see Ahn & Bailenson, 1996; Goedert & Spellman, 2005; Hansen & Hall, 1985; Morris, Smith, & Turner, 1998; Van Overwalle & Van Rooy, 2001; Van Overwalle & Timmermans, 2005). The augmentation principle states that a facilitative cause is perceived as stronger when the effect occurs in the presence of a perceived inhibitory cause. For example, if someone do well on a test despite of the fact that the

test is hard, then the person must be truly smart (for studies on this principle see Hansen & Hall, 1985; Van Overwalle & Van Rooy, 2001).

One important question concerning Kelley's approach is when people use covariation or causal schemata information to make causal judgments, and how the two types of information interact. According to Kelley, previous beliefs or schemas may distort the detection and analysis of covariance information (Orvis, Cunningham, & Kelley, 1975; see also Chapman & Chapman, 1969; Hamilton & Rose, 1980). However, exactly how and when these two types of reasoning interact is not completely clear. Notice that this problem is perpendicular to many social and cognitive models that distinguish between an effortful and controlled mode of processing, on one hand, and a quick processing mode with associative nature, on the other hand (Chaiken & Trope, 1999; Smith & DeCoster, 2000).

Despite advancing with the notion of causal schemata, Kelley was more often mentioned in the literature due to his ANOVA covariation model. This model was criticized by reflecting an excessively rational view of the human mind. Describing the perceiver as capable of computing complex covariation probabilities was not in line with empirical findings (e.g., Hamilton & Gifford, 1976) and with the bounded rationality ascribed to the human mind (Simon, 1957); a mind that, due to computational limitations, makes use of simpler rules of reasoning, or heuristics, and biases in order to deal with an overloading external environment (Tversky & Kahneman, 1974; Gilovich, Griffin, & Kahneman, 2002).

The rational view of the social perceiver was apparent both in the covariation model of Kelley (1973) and in the correspondence inference theory (Jones & Davis, 1965; Jones & McGill, 1974). Both frameworks describe dispositional inferences as a late stage within attributional process. The number of complex analyses that perceivers have to engage in before making a dispositional inference about the actor is well illustrated in the words of Jones and McGillis (1976):

The would-be attributor appraises the effects of the observed act and of plausible alternative acts, considers the effects in terms of his prior expectancies of people in general and of the actor in particular, validates the knowledge-of-effects

assumption, and makes inferences about intentions, and ultimately more stable dispositions. (p.417)

However, subsequent research contradicted a rational picture of the social perceiver. Instead, perceivers' biases and flaws started to become apparent in a number of studies. Attributional researchers' attention turned, then, to the analysis of a recurrent attributional bias: the Fundamental Attribution Error (Ross, 1977) or Correspondence Bias (Jones, 1979).

1.5. The Fundamental Attribution Error and the Correspondence bias

The idea that perceivers explain behavior by reference either to personal internal factors or to situational external factors is a core distinction within attribution research (see Gilbert & Malone, 1995). This dichotomy was already apparent in the work of Heider (1958), and even in works from previous social researchers such as Lewin (1931) and Ichheiser (1949).

Knowing whether the behavior is predominantly guided by the situation or by internal forces is linked to the debate about the existence of free will – the central question whether agents have the power to control their decisions and actions, or whether their actions are mainly controlled by situational constraints. This debate is also crucial to the question of whether the agent can be regarded as responsible for his actions, which has strong morality consequences (Reeder, Kumar, Hesson-McInnis, & Trafimow, 2002; Woolfolk, Doris, & Darley, 2006). As Gilbert and Malone (1995) mentioned, this issue is particularly striking when we look to certain historical events, from which the most distinctive is, undoubtedly, the holocaust tragedy. Here, more than ever, the world asked *how was that possible?* How could we explain the behavior of each person that in some way or another contributed to the occurrence of such brutality?

According to attribution research, social agents tend to perceive behaviors as depending on personal or situational forces. As Hamilton (2005) points out, when perceivers explain behavior by a dispositional characteristic, they have a basis to predict that the future behavior of the actor will be consistent with the inferred disposition. On the contrary, if perceivers explain the behavior by the situation, then little is learned

about the actor. Drawing a dispositional or a situational attribution makes a big difference in what is learned from the event.

Empirical studies on this issue have demonstrated that perceivers show a dispositional bias. Perceivers tend to overestimate the role of personal factors when explaining other's behaviors and to underestimate the influence of situational constraints (Ross, 1977; Gilbert & Malone, 1995). Researchers refer to this bias as the fundamental attribution error (Ross, 1977) or correspondence bias (Jones & Harris, 1967; Gilbert & Malone, 1995).

However, these two terms, although often used as synonyms, are not interchangeable. Both Hamilton (1998) and Krull (2001) noticed that the fundamental attribution error refers to the overestimation of dispositional over situational causes, whereas the correspondence bias refers to the bias to think that behaviors reflect dispositional properties. The main difference lies in the fact that the fundamental attribution error involves the consideration of both dispositional and situational factors, whereas the correspondence bias is defined by a direct link between dispositions and behaviors. Hamilton (1998) specifies the argument by saying that the fact that we tend to infer that a specific behavior reflects a disposition (e.g., the behavior is a friendly behavior, and so the person must be friendly) doesn't necessarily imply the consideration of the causes of the behavior (i.e., Why did the behavior occur?). This same argument is also mentioned by Krull (2001) when he claims that only the fundamental attribution error involves a true causal judgment. Thus, it is theoretical soundness to conceive that the fundamental attribution error involves causal attribution thinking, while the correspondence bias involve dispositional trait inferences, without consideration of the causes of the behavior. Of course, despite referring to different processes, the two errors are probably not mutually exclusive. They can probably co-occur. Specifically, because correspondent inferences are assumed to have a more automatic nature (e.g., Smith & Miller, 1983), they can be the basis for causal attributional errors, when causal reasoning is involved (see Krull, 2001)

Regardless of the distinction proposed by Hamilton (1998), both biases are usually perceived as being intrinsically associated, and the same empirical studies were frequently used to support the occurrence of both phenomena. The fact is that, in these studies, it is impossible to disentangle the intervention of both types of errors.

Therefore, when we describe this research, we will refer to the bias reflected in participants' responses as a general "dispositional bias". The reason is that it is not clear whether a "fundamental attribution error" or a correspondence bias" is involved. We'll come back later to the crucial discussion about the distinction between the causal attributions and correspondence inferences. For now, we'll review some of the evidence that supports the existence of a general dispositional bias.

In a classic study, Jones and Harris (1967) asked participants to read an essay that was either pro-Castro or anti-Castro. Participants were informed that essays were written by other students, in conditions of *free-choice* (students were free to express their true attitude about Fidel Castro), or in conditions of *no-choice* (students were asked to write the essay defending a given position, as part of an assignment for a course). Participants were then asked to make a judgment about the true attitude of the writer of the essay.

Not surprisingly, when there was free choice participants rated writer's attitudes in line with the position advocated in the essay. The unexpected finding was that, even when participants knew that students had no choice about which position to support, they still attributed the attitude according to the direction of the essay. If the person was anti-Castro in the essay, they assumed that the person was anti-Castro, and even if the person was pro-Castro in the essay (an unexpected position within Americans) participants assumed that they were truly pro-Castro. This pattern of results contradicts predictions from correspondence theory concerning intentionality (Jones & Davis, 1965) and, more important, they indicate that participants overlook the importance of situational constraints in the production of the behavior. This effect was replicated in subsequent studies (Forgas, 1998; Gilbert & Jones, 1986; Jones, 1979, 1990; Jones & Berglas, 1976; Jones et al., 1971; Snyder & Jones, 1974).

In another well-known study, Ross, Amabile, and Steinmetz (1977) randomly assigned participants to the role of contestants or questioners in a quiz game. The task of questioners was to compose questions, and then pose those questions to the contestants, while contestants should attempt to answer the questions. Because the questions were made up by the questioners they knew all the answers, while contestants showed some difficulty in responding. However, both contestants (Experiment 1) and external observers (Experiment 2) rated questioners higher on general knowledge than

contestants. Again, participants seem to fail to consider the powerful role of the situation; in this case by neglecting the situational constraints that gave questioners an advantage position in the game.

It is interesting to note the impossibility of differentiating between causal attributional and dispositional inferences processes in both these studies (Jones & Harris, 1967; Ross et al., 1977). The dispositional bias reported in both studies can be explained either by a “correspondence bias” or by a “fundamental attributional error”, if we follow the distinction made by Hamilton (1998). Specifically, we don’t know whether participants overestimated the degree to which behaviors reflected dispositional characteristics, without engaging in causal analysis, or whether participants did consider both situational and personal causes, but gave more weight to the personal factors. Briefly, we it is not clear whether causal reasoning was or not involved in participants’ responses.

Another line of research has argued that this dispositional bias is stronger when observers explain actor’s behavior than when actors explain their own behavior. Jones and Nisbett (1971) proposed that while observers tend to explain actors behaviors mainly in dispositional terms, actors tend to explain their own behaviors by situational constraints. This bias has been called the *actor-observer bias* and some support for its occurrence has been provided (Jones & Harris, 1967; Jones, Rock, Shaver, Goethals, & Ward, 1968; Lay, Ziegler, Hershfield, & Miller, 1974; McArthur, 1972; Nisbett, Caputo, Legant, & Marecek, 1973). Later research has qualified this argument by taking into consideration the type of action. While actors tend to attribute failures to external factors, they attribute success to internal factors. This *Self Serving Bias* (Miller & Ross, 1975) would have the function of preserving our self-esteem.

Despite the popularity of these ideas, subsequent research have provided contradictory findings, putting into question the generality of both the actor-observer bias (Knobe & Malle, 2002; Malle, 2006; Malle & Knobe, 1997; Robbins, & Spranca, & Mendelsohn, 1996; Storms, 1973; Taylor & Fiske, 1975; Taylor, Koivumaki, 1976; Watson, 1982) and the self serving bias (Bradley, 1978; Ender & Bohart, 1974; Miller & Ross, 1975; Ruble, 1873; Ross & Sicoly, 1979). In general, the different patterns of data don’t seem to be explained by a simple person-situation dichotomy (Malle, 2008).

This program of research is fascinating, not only because it indicates that self-attributions are often biased but also because it speaks directly to the question of whether actors have introspective access to the causes that determine their own behavior (Nisbett & Wilson, 1977; Wilson & Nisbett, 1978) or whether they use a priori causal theories (Wilson & Nisbett, 1977) and their past behavior (Bem, 1965, 1967; Kiesler, Nisbett, & Zanna, 1969) in order to conclude about the reasons of their actions. However, in the present thesis we are specifically interested in understanding how perceivers process and interpret the behaviors of *other* targets.

For that reason, we return to the most well known interpersonal bias – the dispositional bias (Jones & Harris, 1967; Ross et al., 1977). Why do people seem to prefer dispositional over situational factors to make sense of others actions? We will describe some contemporary attribution models (Gilbert et al., 1988; Gilbert et al., 1988; Quattrone, 1982; Trope, 1986) that provided useful insights into this question.

1.6. Why do people prefer dispositional attributions?

The first theoretical contribution to the clarification of the tendency to overestimate the importance of personal factors was given by Heider (1958). Heider noticed that the behavior “*tends to engulf the total field*” of the observer (Heider, 1958, p.54). The fact that observers’ attention is focused toward the behavior may result in neglecting the influence of the surrounding context.

However, for different reasons, this view is not totally satisfactory. First, subjects show a dispositional bias even when the perceptual salience of the behavior cannot explain the effect (Jones & Harris, 1967). Second, as some authors properly questioned (Gilbert & Malone, 1995), it is not clear why the perceptual salience of the behavior should favor dispositional factors, since it is the behavior that is salient, not dispositions. In fact, the claim that stronger salience of the behavior should result in stronger dispositional responses, despite initially supported (e.g., Storms, 1973), was not confirmed in later studies (Arkin & Duval; 1975; Ellis & Holmes, 1982; McArthur & Post; 1977; Taylor & Fiske, 1975). Third, Heider’s account lacked a clarification of the specific underlying mechanisms and processes, which make some authors question whether it can be considered an explanation (Gilbert, Pelham, & Krull, 2003; Jones &

Harris, 1967). Thus, other accounts were needed in order to provide a better explanation for the tendency to overestimate dispositional factors.

1.7. Quattrone anchoring and adjustment Model

Quattrone (1982) proposed a two stage model, based on anchoring and adjustment principles (Tversky & Kahneman, 1974). The model assumes that people start by making an inference in line with the initial context and then the initial inference is adjusted by considering other possible constraints on the behavior. Since these adjustments are usually insufficient, people tend to show an overattributional bias, in favor of the initial inference.

Because previous experiments, and also natural settings, tend to promote initial dispositional inferences, an insufficient correction of the inference by situational factors is usually obtained, which results in a dispositional bias. However, Quattrone (1982) noticed that the effect would be reversed if the initial context favors situational causation. This prediction was confirmed in a study conducted by Quattrone (1982) that applied the attitude-attribution paradigm (Jones & Harris, 1967). In this study, it was shown that when participants are asked to judge the importance of situational forces, rather than to infer the attitude of the participant, a failure to make sufficient adjustments to the underlying dispositions is observed (Quattrone, 1982).

1.8. Trope identification-inference Model

The model proposed by Trope (1986) distinguishes between two different processes: the *identification* process, and the *dispositional inference* process.

The identification process corresponds to a categorization of the event. In this phase the perceiver (a) encodes the actor's cues in terms of behavioral categories (for example, the behavior is identified as a generous behavior), and (b) represents situational cues in terms of situational categories (for example, the actor works in a charitable institution). The identification depends not only on the present external cues (bottom-up processing), but also on knowledge based-expectations about the actor, and about the context (top-down processing) (e.g., McClelland, 1979; Morton, 1969; Rumelhart & McClelland, 1982). Also, it is assumed that situational and actor's cues influence the initial identification of each other. That is, the cues provided by the actor

influence how the situation is identified and the cues provided by the situation influence the initial identification of the actor's behavior. This influence is stronger the more the cues are ambiguous. Empirical data gave support to this notion. It was shown that the emotiveness evoked by the situation influences in a consistent manner the emotion that is identified in a face, being this effect particularly evident with ambiguous faces (Trope, 1986; Trope, Cohen, & Maoz, 1988).

The product of the identification phase serves as input to the dispositional inference process. At this stage, perceivers use a subtractive rule, according to which the probability of attributing a disposition to the actor is higher when the actor is performing a correspondent behavior, and the behavior is inconsistent with situational constraints. Results consistent with such predictions were also obtained, showing, for example, that the same angry reaction is judged as less dispositional when the situation is frightening (i.e., consistent) (Trope et al., 1988).

While the identification process was described as automatic and unconscious, the dispositional inference process was portrayed as controlled. However, Trope (1986) observes that the dispositional inference phase may also become automatic.

An important point concerning this theoretical account is that it favors the view of identification and dispositional inference as distinct cognitive processes. More important, according with this view these two processes may not match. The behavior can be categorized as an angry behavior, but depending on the situation, an "angry" disposition can be, or not, attributed to the target.

1.9. Gilbert three-stage Model

Gilbert's model apparently combines aspects of both previous models, and is probably the model more suitable to the incorporation of the fundamental attribution error. According to the model (Gilbert et al., 1988; Gilbert et al., 1988) the social inference process is composed by three distinct sequential processes. The first stage is *categorization* of the behavior, which is assumed to be a pre-attributional process, in which the perceiver describes the type of behavior observed. This is identical to the identification process proposed by Trope (1986). Second, a *characterization* or dispositional inference is made about the actor. And finally, the inference is *corrected* by taking into consideration the influence of possible situational constrains.

The two first operations are said to occur outside awareness and to be efficient, while correction is viewed as a deliberative and resource demanding process. It is suggested that characterization should be viewed as a perception-like process, whereas correction is a higher order reasoning process

From this analysis, it follows that resource consuming activities (which are regular in everyday life) would not interfere with the categorization and characterization stages, but would make perceivers unable to consider the influence of the situation. Gilbert and colleagues (Gilbert et al., 1988; Gilbert et al., 1988) provided support for these assumptions by showing that participants under cognitive load (i.e., performing concurrently another task) are more dispositional in their judgments than are participants in no depletion conditions. For example, Gilbert, et al. (1988) used a paradigm adapted from Snyder and Frankel (1976) where participants are shown a video of a woman manifestly anxious. Participants were either told that the woman was discussing anxious-related topics or relaxing topics. Participants in the low load condition made an appropriate analysis of the event, rating the target more dispositional anxious when discussing relaxing topics than when discussing anxious topics. In opposition, participants in the high load condition rated the woman as equally dispositionally anxious in both conditions, apparently failing to take the situation into account.

Subsequent studies reported, however, that this pattern can be reversed (Krull, 1993; Krull & Erickson, 1995). These and other studies provided a more exact, but also more complex, picture of the conditions under which the so called fundamental attribution error will occur. We will consider some of these studies next.

1.10. Questioning the existence of a Fundamental Attribution Error

Despite initial evidence (Jones & Harris; 1965; Ross et al., 1977) favoring the tendency to make internal dispositions, even under conditions in which situational constraints are sufficient to explain actor's behavior, later research questioned the psychological reality of such phenomenon. One reason to question these results was related with the particular salience of the dispositional goals.

1.10.1. Salience of the dispositional goal

Some authors (Krull, 1993) argued that a dispositional bias is found because of the fact that in previous studies participants had a dispositional inference goal. If the goal was a situational one, then no dispositional bias should occur. Quattrone (1982) was the first to provide evidence supporting this notion. Using the attitude attribution paradigm (Jones & Harris, 1967), Quattrone demonstrated that participants explicitly informed about the previous attitude of the author of a given essay, under essay free-choice conditions, continue to attribute the attitude expressed to features of the situation.

In addition, Krull and Erickson (1995), reversing the paradigm of Gilbert et al. (1988), informed participants that the woman exhibiting anxious signs in the video was either dispositionally anxious or calm. Participant's task was to diagnose the anxious tenor of the discussed topics. Results showed that busy participants, contrary to non-busy, rated the nature of the topics without taking the dispositional information into account (see also Krull, 1993). Thus, when participants have a situational goal, apparently they first characterize the situation, and then fail to correct for dispositional factors.

These results are in agreement with the notion that the sequential inference process is not fixed (Krull, 1993), contrary to what was assumed by Gilbert (Gilbert, et al., 1988; Gilbert et al., 1988). Instead, it seems that people infer a disposition (personal *or* situational) according with experimental requirements, and in agreement with what attracts more their attention. Then, they adjust the initial inference either for situational or for dispositional factors. This account is in line with Quattrone's anchoring and adjustment model (Quattrone, 1982).

Also, these studies suggest that it is not necessarily the case that participants are inherently more dispositional than situational. A more likely explanation is that, in general, it is harder to think about factors that are not salient in the context. In previous studies it was more difficult to take situational factors in consideration because the requirements of the experiment were dispositional. Subsequent studies supported this notion, demonstrating that the dispositional bias diminishes or disappear when the aspects of the situation are made salient (Ajzen, Dalto, & Blyth, 1979; Fein, Hilton, &

Miller, 1990; Fleming & Darley, 1989; Wright & Wells, 1988), or when the event is processed in a more systematic ways (Forgas, 1998; Tetlock; 1985; Webster, 1993).

Ajzen et al. (1979) proposed that the dispositional bias reported in the attitude-attribution paradigm (Jones & Harris, 1967) is explained by the fact that ambiguous background information about the actor is re-interpreted as consistent with the dispositional nature of the essay. They replicated the study from Jones and Harris (1967) using different materials. Results showed that, when no background information is provided, there are no differences in attitude attribution in the constrained condition. In this study, it is not clear whether results can be explained by the fact that ambiguous profile information could no longer be biased according to the attitude reported in the essay, or because removing background information makes situational constraints more salient, or due to both aspects. In any case, results suggest that the fundamental attribution error is not as pervasive as previously thought.

In a different replication of the attitude-attribution paradigm (Jones & Harris, 1967), Fleming and Darley (1989, Experiment 1) showed that when the random nature of the assignment is made more salient, and cues suggesting behavior intentionality are totally removed, there is no evidence for a dispositional bias under constraining conditions. In agreement, Trope and Gaunt (2000), manipulating the salience of the situation in different ways, showed that participants correctly take the situation into account even in a high-load condition, as long as situational factors are salient.

Wright and Wells (1988) called attention to the arbitrariness of the situation, by alerting participants that the information could not be enough to draw an attitude attribution about the writer. Under these conditions, the dispositional bias decreased drastically. In a consistent way, when experimenters raise a suspicion of ulterior motives (e.g., suggesting that the essay was written to ingratiate an important person), even in conditions of free choice, participants fail to attribute an attitude consistent with the essay (Fein, 1996; Fein et al., 1990).

A different line of studies shows that the inducement of a more systematic processing also diminishes the dispositional bias magnitude. For example, in a study conducted by Forgas (1998) it was demonstrated that participants in good mood are more likely to make dispositional judgments based on coerced written essays, while participants in bad mood are less likely to make dispositional judgements. They explain

the results by the fact that bad mood tends to lead to a more systematic processing of information (Schwarz, 2001; Schwarz & Bless, 1991) and thus increase the likelihood of considering the role of situational factors.

A study by Webster (1993) showed that the need for closure, defined as the need to avoid ambiguity and to quickly reach an answer for problems (Webster & Kruglanski, 1994), also affects the magnitude of the dispositional bias. It was demonstrated that when the task required an attribution to the person, high need for closure participants showed a greater dispositional bias, and low need for closure were less dispositionally biased. When the task required a situational attribution (Experiment 3) high need for closure participants revealed higher situational bias. We can presume that results can be, at least in part, explained by the fact that the desire to reach closure lead individuals to avoid cognitive complexity (Webster & Kruglanski, 1994), preventing them from processing the details of the event.

In a similar vein, Tetlock (1985) showed that when participants are made accountable, no dispositional bias occurs. This can also be related with the fact that under accountable conditions, information is processed more carefully, since participants have to justify their decisions to others.

Finally, Leyens, Yzerbyt, and Corneille (1996) provided evidence supporting the fact that the specific topics of previous essays (e.g., abortion, drugs) in general, favor dispositional explanations. They showed that when other topics are used, different types of explanations are preferred, and no dispositional bias occurs (see also Corneille, Leyens, Yzerbyt, & Walther, 1999).

These studies are in line with the argument that, in previous experiments reporting a fundamental attribution error, it is in general more difficult to take the situation into account. When experiments are constructed in such a way that it becomes easier to consider situational factors than dispositional factors, the effect is reversed (Krull, 1993; Krull and Erickson, 1995; Quattrone, 1982). In addition, when the original paradigm is used, but is designed in such a way that facilitates the consideration of the situation (either by making the situation more salient, or by promoting a more systematic processing of information) there is no evidence for a dispositional bias occurrence. As Gawronski (2004) appropriately noticed, the fact that participants take into account the situation, when it is made salient, or when they process more carefully

the event, is inconsistent with the view that participants consider situational factors to be irrelevant for judgments.

A different way to look into this pattern of data is suggested by Erickson and Krull (1999). They support the view that dispositional inferences and causal attributions should be regarded as distinct cognitive processes, and they use this distinction to explain previous results. Specifically, they suggested that in the attitude-attribution paradigm a dispositional inference (that can be either personal or situational, depending on how the experiment is set up) is readily made. However, because additional constraints on the behavior are often in the form of causal information, its consideration requires causal attributional thinking. Once causal attribution and dispositional inferences are supposedly guided by different principles, it might be difficult for participants to integrate attributional and inferential information.

This reasoning is supported by a study from Johnson, Jemmott, and Pettigrew (1984), where it was confirmed that participants can exhibit a dispositional bias, even when they accurately take into account the constraining character of the situation. Using the quizmaster paradigm (Ross et al., 1977), they found that participants rated the questioner higher in general knowledge, even when they recognized that the way the experiment is constructed make questioners look more intelligent. It seems, then, that dispositional inferences and causal attributions operate in parallel ways, and the information gained from one process is not adequately adjusted within the other.

1.10.2. Perceived intentionality

Other research questioned the fundamental attribution error on different grounds. It was suggested that, given the features of the attitude attribution paradigm, making a dispositional inference is not totally unreasonable. For instance, Fleming and Darley (1989) hypothesized that the pattern observed by Jones and Harris (1967) was due to the fact that participants attribute some degree of intentionality to the actors. In line with this argument, Miller and Rorer (1982) argued that participants assumed that writers would not have written such strong arguments, if those arguments were not in agreement with their true attitudes.

Supporting this perspective, Miller and colleagues (Miller, Ashton, & Mishal, 1990, Experiment 1) reported that, when participants are asked to write essays in

agreement with their true attitudes, they rate those essays as stronger, and as more likely to evoke correspondent inferences than essays inconsistent with their true attitudes. Notice that these participants are anticipating the dispositional bias. Accordingly, it was found that, when easy essays are provided to observers, under constraining conditions, the correspondence bias is reversed (Miller & Rorer, 1982; see also Gawronski, 2003).

In sum, it might be the case that observers are not committing any *fundamental error*, but are using implicit theories according to which good essays are written by true believers, while bad essays are written by actors that don't believe in the position advocated. If that is the case, it is debatable whether making a dispositional judgment in this case is an error. However, Snyder and Jones's (1974) results add ambiguity to this debate. They asked participants to write essays and subsequently those essays were given to new participants. These later participants were asked to rate the true attitude of writers. Under these conditions, participants still report that the writer's true attitude is in agreement with the essay. So, results don't provide a clear answer about the role of essay difficulty in the dispositional bias occurrence.

A different study approaching the intentionality issue was conducted by Fleming and Darley (1989, Experiment 2). They demonstrated that if the actor exhibits a facial expression conveying disappointment (i.e., a sign that explicitly removes any perception of intentionality) after receiving the direction of the random assignment, there is no evidence of dispositional bias.

Given the studies that showed that the magnitude of the dispositional varies with a great number of variables, some authors claimed (Harvey & McGlynn, 1982; Harvey, Town, & Yarkin, 1981) that there is no basis to talk about a fundamental attribution error (but see Reeder, 1982). If, depending on the experimental circumstances, people may overestimate dispositional factors, or overestimate situational factors, it can be legitimately asked whether the fundamental attribution error is a *fundamental bias* in person perception. Maybe a more prudent way of describing these facts is illustrated by Harvey and McGlynn (1982):

There are conditions (maybe many such conditions) that produce a great tendency to attribute to dispositions. There are conditions that produce a great tendency to attribute to situations. There are conditions that produce a great tendency to

attribute to both dispositions and situations in some interactive fashion...There are conditions that produce a tendency not to attribute to anything...(p. 346)

Considered in this way, the existing research raises some intriguing questions. On one hand, we may question whether the dispositional bias observed in previous studies was just a reflex of researchers suffering themselves from a bias, being much more concerned with trait determinants than with situational determinants of behavior. On the other hand, we may ask if the fact that some studies prove that people take situational constraints into account necessarily means that outside the laboratory people are not generally dispositional. In fact, we usually don't have much situational information about people around us, which can result in a dispositional bias.

Finally, from a theoretical point of view, neither the disregard of situational causal constraints, nor the existence of a correspondence bias, are necessary conditions for fulfilling the basic human needs of control and stability. Contrary to what is sometimes implicitly assumed in the literature, we can easily think that perceivers would have a greater control over reality if they consider both personal and situational constraints. Certainly, more empirical work is needed to clarify these questions. Mainly, the data presented legitimates the search for more dynamic attribution models, able to account for the perceiver's perception of the interplay between dispositions and situations (Gawronski, 2004).

We have reviewed the main research and theoretical explanations about how people tend to perceive other people's behavior. However, it's also important to know *when* people engage in attributional analysis in the first place. Do perceivers carefully search for the causes of all events they observe, or are there some events that trigger attributional causal analysis more than others?

1.11. The “When” of Attributional Processes

As noted by Kelley (1967), attribution research portrays social perceivers as if they were highly motivated causal searchers. However, the fact that participants may engage in attributional analysis when explicitly instructed to do so doesn't mean that, in the absence of such instructions, they would spontaneously perform the same kind of analysis.

In line with this reasoning, Enzle and Shopflocher (1978) demonstrated that explicit instructions instigate attributional processes that do not occur in the absence of overt attributional requirements. It was shown that participants instructed to make an attributional analysis rated an actor that performed a helpful behavior under unconstrained conditions more positively than when the behavior was performed under experimental inducing instructions. However, when participants weren't required to engage in attributional processes, no differences in ratings were observed.

Subsequent research explored which conditions more likely to spontaneously instigate attributional thinking. This research showed that the occurrence and magnitude of attributional analysis is stronger in face of unexpected behaviors (Clary & Tesser, 1983; Hastie, 1984; Kanazawa, 1992; Lau & Russell, 1980; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981); in conditions of subjective loss of control (Pittman & Pittman, 1980; Swann, Stephenson, & Pittman, 1981); when others are personally relevant (Berscheid, Graziano, Monson, & Dermer, 1976; Harvey, Yarkin, Lightner, & Town, 1980; Yarkin-Levin, 1983); and under failure circumstances (Diener & Dweck, 1978; Wong & Weiner, 1981).

A number of studies confirmed that unexpected events instigate stronger causal processing. Lau and Russell (1980), for instance, analyzed the attribution analysis in sport newspapers and verified that attributional causal explanations are more likely to occur in face of unexpected results. Pyszczynski and Greenberg (1981) showed that participants prefer to analyze helping items about an actor when the actor had previously exhibited an unexpected helping behavior. In addition, Clary and Tesser (1983) reported that participants that are instructed to retell previous stories spontaneously introduce a larger number of causal explanations for unexpected than for expected events. Finally, in Hastie's (1984) study participants were presented with behavioral-describing sentences under impression formation conditions, and were asked to write a continuation for each sentence. Results showed that participants were more likely to complete expectancy-incongruent sentences with explanations about why the behavior occurred, in comparison with expectancy-congruent behaviors. Thus, unexpected events tend to activate causal attributional analyses.

In addition, it was extensively argued in the attribution literature that one of the biggest motivations underlying attribution causal explanations is the need for control.

Few studies had empirically scrutinized this assumption. However, some experiments provide data that confirms the influence of the subjective feeling of control on attribution. Pittman and Pittman (1980), for example, presented data suggesting that participants have a great desire to learn about others after experiences of helplessness. In another study, Swann et al. (1981) deprived some participants control by giving non-contingent feedback about their performance. It was shown that control deprived subjects sought more information about others, compared with non deprived control participants.

Berscheid and her colleagues (Berscheid et al., 1976) reported the effect of personal relevance of actors on attribution. They showed that when men and women anticipate a date with a stranger they are more likely to engage in causal reasoning, as well as on processing activities inherent to causal analysis. Harvey et al. (1980) further showed that when participants are induced to have a more emphatic position (Experiment 1), or to anticipate interaction with one of the persons observed in a video (Experiment 3; see also Yarkin-Levin, 1983), they make more attributions about the observed actors, compared with individuals that are not given such instructions.

Lastly, some studies provided evidence in favor of the idea that failures trigger a search for causal explanations. Diener and Dweck (1978) reported that children that have more difficulty dealing with failure (i.e., helpless oriented children) tend to make more causal attributions in order to explain failure, while mastery-oriented children make few attributions and report more thinking about ways of overcoming failure. More direct evidence was provided by Wong and Weiner (1981), who asked participants which questions they would ask themselves after observing certain events. They reported that participants make more “why” questions under failure than under success conditions.

In sum, the likelihood of engaging in an attributional thinking in order to explain “why things happened as they did” depends on the event in question. While some actions apparently don’t need explanation, saving our efforts, and being easily integrated in our base of knowledge, other actions are intriguing and tend to instigate causal reasoning. An important point is that some of the conditions that tend to instigate stronger causal processes are also known for leading to *less* dispositional trait inferences. For example, unexpected behaviors result in deeper causal reasoning (Clary

& Tesser, 1983; Hastie, 1984; Kanazawa, 1992; Lau & Russell, 1980; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981), but are less likely to lead to correspondent inferences (Jerónimo, 2007). As Hamilton (1998, p.106) asserts “attributional thinking is most likely to occur precisely under those conditions when correspondent inferences are least likely to be made, specifically, when the behavior violates the operative expectancy”. As we will see, this is a fundamental argument in favor of the distinction between causal attributional and dispositional inference processes.

1.12. Differentiating Dispositional Inferences and Causal Attribution

Early researchers (Heider, 1958; Jones & Davis, 1965; Jones & McGillis, 1974; Kelley, 1967) saw causal attributional analysis as an antecedent of dispositional inferences. Later models (Gilbert et al., 1988; Gilbert et al., 1988; Quattrone, 1982; Trope, 1986) placed the dispositional inference in a relatively early position within the inference process. However, these models were not sensitive to the distinction between dispositional inferences and causal attributions, and not explicitly incorporated this distinction into their proposals. We can presume that the initial phases in the different models tend to involve dispositional inferences, while the adjustment (Quattrone, 1982), subtracting rule (Trope, 1986), or correction (Gilbert et al., 1988) are more likely to involve causal thinking. But the differentiation between dispositional inference and causal attribution is not explicitly mentioned by the models.

There are, however, strong theoretical (Hamilton, 1988, 1998; Krull, 2001) and empirical reasons (Bassili, 1989; Erickson & Krull, 1999; Johnson et al., 1984; Smith & Miller, 1983; see also Hilton, Smith & Kim, 1995; Reeder & Spores, 1983) that support the view that dispositional inferences and casual attributions are distinct mental processes. Hamilton (1988; 1998) firstly argued for this distinction. While dispositional trait inferences involve a direct inference of a personality trait from a behavior, attributional processes involve the consideration of the causes of the behavior. Hamilton (1998) noted that a way of capturing the distinction is by looking to differences of processing of congruent and incongruent behaviors. Dispositional inferences are easily drawn from behaviors that confirm our expectations about the actor. For example, if we have a friendly expectation about John, it is likely that we will easily draw a dispositional trait inference when we observe John performing a friendly behavior. The

expectation facilitates the encoding of the behavior in terms of the corresponding trait (see Jerónimo, 2007). In contrast, when the observed behavior is inconsistent with existing expectations, different processes came into play. In this case, a causal attributional analysis is likely to occur that allows the perceiver to make sense of the puzzling observed behavior. Thus, consistent-behaviors and inconsistent-behaviors are likely to instigate different processes. While consistent behaviors tend to trigger dispositional trait inference processes, inconsistent behaviors tend to activate causal attributional reasoning.

In agreement with Hamilton (1988; 1998), Erickson and Krull (1999) proposed a way to define both processes. According to these authors, dispositional inferences occur when the trait that characterizes the behavior (e.g., anxious) is used to characterize the person (e.g., he is an anxious person), or the situation (e.g., the situation is anxiety provoking); and causal attribution is the process of searching for causal understanding (e.g., Why he is acting so anxiously?).

Empirically, Smith and Miller (1983) showed that there are no significant time differences between categorization and the time participants take to make a dispositional inference. On the contrary, causal attributions take significantly more time than categorization processes. They concluded that dispositional inferences are closely linked to the identification of the behavior and can occur spontaneously.

Reeder and Spores (1983) provide empirical results showing that causal attributional judgments are more affected by the different demands of the situation, than judgments of the actor that not involve causal attribution processes. Specifically, it was shown that judgments of morality based on immorality behavior are unaffected by situational pressures that facilitate the immoral behavior. By contrast, causal attributions are affected by the situational demands.

Research by Johnson et al., (1984) is also consistent with this distinction. In this study it was shown that even when participants recognize the situational constraints of the event, they still exhibit a dispositional bias. They concluded that correspondence inferences may be highly independent from causal judgments.

Bassili (1989) further reported that while impression formation participants exhibit a high level of trait activation, participants instructed to allocate causality

showed no effects of trait activation. These results clearly suggest that causal reasoning do not necessarily involve thinking in terms of trait.

Erickson and Krull (1999) also tested the validity of this distinction. They thought that because the dispositional inference process is seen as being more strongly linked to behavior interpretation than are causal attributions, the way a behavior is interpreted would have a greater impact on dispositional inferences than on causal attributions. The predictions were supported by their results, demonstrating that dispositional inferences were in fact closer to the interpretation of the behavior than are causal attributions.

Overall, results support the view of dispositional inferences and causal attributions as distinct processes. The automaticity of both these processes is still a debatable question, open to empirical test. Attributional processes seem to usually occur in a systematic and elaborated way (Kelley, 1967), but it was also suggested that attributions can occur in more automatic ways, supported by causal-knowledge structures and schemas (e.g., Abelson & Lalljee, 1988; Abraham, 1988; Hilton & Knibbs, 1988; Kelley, 1974). Dispositional inferences (personal or situational), on the other hand, usually occur in less deliberative ways, but we can also question if there are circumstances where the process can be totally disrupted; or for instance, ask whether under some circumstances the dispositional inference will only occur after a careful causal attributional analysis.

In any case, dispositional inferences and causal attributional processes are distinct processes that may occur independently of each other. Thus, despite trait inferences being initially studied within an attributional framework, and regardless of the fact that attribution research had provided important information about how people go from behaviors to trait inferences, the study of dispositional trait inferences could be carried out independently of an attributional background.

1.13. Summary of the Chapter

The initial study of dispositional trait inferences was mixed up with the study of attributional processes. Classic attributional researchers (Heider, 1958; Jones & Davis, 1965; Kelley, 1967, 1973) portrayed the social actor as quite rational. According to these authors, the social perceiver analyzes the determinants of observed behaviors

before making a dispositional inference. However, the social perceiver doesn't seem to rely on rational attributional analysis, as proposed by classic attributional researchers (Heider, 1958; Jones & Davis, 1965; Kelley, 1967, 1973), but to incur in different suboptimal strategies.

It was extensively claimed that perceivers tend to exhibit a dispositional bias, overestimating the power of personal factors in determining behavior. This dispositional bias has been referred as the fundamental attributional error (Ross et al., 1977) or correspondence bias (Jones & Harris, 1967). Despite some authors having point out that the fundamental attributional error involve attributional thinking while the correspondence bias involve inferential thinking (see Hamilton, 1998; Krull, 2001), the studies that explored the dispositional occurrence were not sensitive to this distinction.

Different models were developed to describe the way perceivers perform dispositional inferences (Gilber et al., 1988; Quattrone, 1982; Trope, 1986). All of these models were concerned with accounting for the interplay between situational and personal factors, and usually defined a set of serial stages involved in making a dispositional inference. Quattrone (1982) suggested that an initial inference is made according to the context, and then an insufficient adjustment occurs. Both Trope's (1986) and Gilbert's (Gilbert et al., 1988) models make a distinction between an identification stage and a dispositional inference stage. The identification process is regarded in both models as an automatic, perception-like process, but the dispositional process, while described by Trope (1986) as a controlled process, is seen by Gilbert (Gilbert et al., 1988) as being automatic. Empirical studies that were done within this framework are not able to clarify the automaticity question, since overt instructions were usually applied. The types of processes that are used when explicitly asked to judge or explain others' behavior can be different from those that people used without such requirements.

Attributional research further showed that attributional analysis tends to be triggered under certain special circumstances, as when perceivers observe unexpected events (Clary & Tesser, 1983; Hastie, 1984; Kanazawa, 1992; Lau & Russell, 1980; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981). It is also known that unexpected behaviors are less likely to be encoded in terms of traits (Jerónimo, 2007).

Thus, these findings favor the distinction between causal attribution and dispositional inferences.

We also highlight the arguments that support the distinction between causal attribution and dispositional inferences (Bassili, 1989; Erickson & Krull, 1999; Hamilton, 1988, 1998; Johnson et al., 1984; Krull, 2001; Reeder & Spores, 1983; Smith & Miller, 1983). This means that the study of the dispositional inference process does not need to be necessarily linked to the study of casual attribution. This was the main impetus for the development of the spontaneous trait inference literature.

CHAPTER II

SPONTANEOUS TRAIT INFERENCES

In the previous chapter, we looked into the literature on Attribution. The initial approach to the study of trait inferences was conducted within an attributional framework. The attribution literature made important contributions to the clarification of dispositional inferences processes. However, causal attributional and dispositional inferences are different cognitive processes that may occur independently of each other (Bassili, 1989a; Erickson & Krull, 1999; Krull, 2001; Hamilton, 1988; 1998; Smith & Miller, 1983). That is, despite sometimes trait dispositional inferences could result from attributional processes, the nature of dispositional inferences processes could probably be better captured if detached from the study of attributional processes. This explains the emergence and rapid growth, of the spontaneous trait inference field of research.

The spontaneous trait inference (STI) literature approached the study of dispositional inferences independently of considering deliberative attributional processes. Although other dispositional inferences about people may be considered (e.g., desires, goals, beliefs, attitudes, abilities, interests; see Heider, 1958; Malle, 2008), the focus was predominantly on the conditions in which perceivers infer personality *traits*. This tendency was grounded on the generally held assumption (Asch, 1946), supported by empirical data (e.g., Fiske & Cox, 1979), that personality traits are central constructs in person perception, playing a predominant role in the organization of “perception, storage, and retrieval of information about people” (Hastie & Kumar, 1979, p.26).

Because previous studies from attribution research mainly relied on explicit instructions, asking participants to causally explain behavior, or to form impressions about a given actor, it was not possible to determine whether perceivers would engage in the same processes in the absence of such instructions. It was important to clarify whether processes that were observed in the laboratory would also occur under more natural settings (Enzel & Schopflocher, 1978). STI researchers were driven by this question, the main goal being to explore whether trait inferences occur *spontaneously* (for reviews see Uleman et al., 1996; Uleman et al., 2008).

By definition (Uleman et al., 1996) a spontaneous trait inference is said to occur when a personality trait of an actor is inferred from his behavior without an explicit intention to form an impression, or to infer a personality trait, about the actor. The possibility that trait inferences could have a spontaneous nature was already apparent in

the way Solomon Asch characterized the impression formation process: “We look at a person and immediately a certain impression of his character forms itself in us. (...) We know that such impressions form with remarkable rapidity and great ease” (Asch, 1946, p.258). With the emergence of STI literature, Asch’s intuition was empirically and systematically examined.

The STI approach implies a new conceptualization of the dispositional inference process that is totally different from the one portrayed by attributional research. While dispositional inferences were previously seen as a relatively late stage of an attributional process instigated by the explicit motivation to understand others (Jones & Davis, 1965; Kelley, 1967, 1973), now it was suggested that the mere attention to and comprehension of a behavior would be enough to prompt a trait inference about an actor, even when the participant has no intention of inferring the trait and no awareness of having done so.

This new perspective changes the characterization of the person perception process, in at least two fundamental ways. First, seeing trait inferences as inherent in the comprehension of the behavior, and not as a distinct mental stage (Winter & Uleman, 1984), would imply that inferring personality traits about others is a highly recurrent process. Second, if trait inferences occur without our explicit intentions, it would mean that our impressions of others are founded in part on processes that go unnoticed by us. STIs can, thus, influence the way the actor is perceived, as well as the processing of future information about him, without perceiver’s conscious awareness (Moskowitz & Roman, 1992). Because of its implicit nature, these processes should be difficult to control or avoid. Initial evidence supporting the spontaneity of trait inferences was provided by Winter and Uleman (1984).

2.1. Initial Evidence for Spontaneous Trait Inferences

According to Winter and Uleman (1984), STIs occur as part of the normal process of behavior comprehension. As a result, upon the observation of a behavior, both the behavioral information itself and the inferred personality trait are encoded in memory.

In order to test the spontaneity of trait inferences, Winter and Uleman (1984) relied on the encoding specificity principle (Tulving & Thomson, 1973). This principle states that “what is stored is determined by what is perceived and how it is encoded, and

what is stored determines what retrieval cues are effective in providing access to what is stored” (Tulving & Thomson, 1973, p.353). A central idea of this principle is that a retrieval cue is effective when the relation between the cue and the retrieval item has become associated during the encoding of the item (Thomson & Tulving, 1970).

Support for the encoding specificity principle was reported in several studies (Thomson & Tulving, 1970; Tulving & Osler, 1968). For example, Thomson and Tulving (1970) presented participants to be remembered words (e.g., man) in the presence of weak semantic associates (e.g., hand). Recall of the words was higher when the same weak semantic associates were provided as retrieval cues, compared with both a no-cue condition and with a condition in which strong semantic associates not presented at encoding were provided as cues (e.g., woman). Actually, the two last conditions didn’t differ in terms of recall efficacy. The authors viewed this data as evidence that a cue, even if previously associated with the word, will not be effective in retrieval unless it was part of the overall encoding representation of the item.

Applying the encoding specificity principle to STI research, Winter and Uleman (1984) predicted that, if traits are spontaneously inferred during behavior encoding, they will be effective cues for the retrieval of behaviors. Thus, trait-cue effectiveness at recall would serve as an indicator that a trait inference had occurred during behavior encoding.

In Winter and Uleman’s (1984) experiments, participants were presented with trait implying-behaviors (many of them taken from Smith & Miller, 1983) under memory instructions. Each behavior was paired with an occupation that was shown in pre-testing to not be related with the behavior (for example, the occupation “reporter” was paired with the behavior “steps on his girlfriend’s feet as they foxtrot”). After an anagram distractor task, participants were asked to recall the behaviors under different cue conditions. In the first experiment (Winter & Uleman, 1984, Experiment 1) the efficacy of trait cues (e.g., the trait “clumsy” as cue for the behavior “the reporter steps on his girlfriend feet as they foxtrot”) was compared with the efficacy of semantic-cues associated with the actor (e.g., “newspaper” when the actor was “reporter”), and also with a no-cue condition. The main argument was that, if trait cues are at least equally effective as strong semantic cues, it could only be due to episodic links between

sentences and traits that was created (inferred) during encoding, since a priori associations between traits and behaviors were shown in pre-tests to be weak.

Results were analyzed using a score consisting of the number of sentence parts (i.e., actor, verb, object, and preposition) recalled. Based on this score, it was observed that the percentage of recall was higher in the trait-cue than in the no-cue condition, but there was no difference between trait-cues and semantic-cues. However, semantic and trait cues were differently effective for recall of different parts of the sentences. Semantic cues were more effective in recall of the actor, while trait-cues were more effective in recall of the other parts of the sentence (verb, object, and preposition). In addition, when recall of the entire sentence is considered (the most frequently recalled combination), recall was equally high in the semantic-cue and trait-cue conditions, and higher in both these conditions than in the no-cue condition.

This first experiment gave some support to Winter and Uleman's (1984) predictions. However, as the authors observed, it could be argued that the semantic-cue condition used in the first experiment did not represent an adequate control condition. Because semantic cues were pre-tested to be associated only with the actor, and not with other sentence elements, the efficacy of semantic cues could have been underestimated. In order to control for this problem, a second experiment was run in which trait-cue efficacy was compared with the efficacy of semantic-cues associated with the verb phrases. This pre-test revealed that six of the eighteen verb phrases used in the first experiment generated dispositional traits as stronger associates, which raises the possibility that the efficacy of dispositional cues in the first experiment could have, at least in part, an associative nature. These six sentences were removed in the second experiment.

Results showed, first, that in terms of the overall recall index, trait-cues were more effective than semantic-cues and no-cues. The last two conditions didn't differ. Second, for recall of the entire sentence, trait-cued recall was more effective than semantic-cued recall, and marginally higher than no-cued recall. Third, trait-cues were more effective than both semantic and no-cued recall for all sentence parts, except for recall of the actor (that didn't differ across conditions).

The overall efficacy of trait-cued recall was taken as evidence for dispositional trait inferences during encoding. However, since recall of the actors themselves didn't

differ across conditions, results cannot clarify whether the inferred trait is a description of the actor or merely a description of the behavior. If inferred traits were descriptions of the actors, traits should be mostly effective in retrieving the actor, which was not the case.

Regardless of these aspects, Winter and Uleman (1984) interpreted the fact that trait-cued recall (i.e., recall with trait cues) was equally (Experiment 1) or more effective (Experiment 2) than semantic-cued recall (i.e., recall with semantic cues), and higher than no-cued recall (i.e., recall without cues), as evidence that trait inferences are unintentionally made during behavior encoding. Moreover, because participants reported almost no conscious awareness of making trait inferences, it was suggested that trait inferences are not only unintentional but also highly unconscious. Since lack of intentionality and awareness are traditionally seen as two important criteria of an automatic process (LaBerge & Samuels, 1974; Posner & Snyder, 1975; Shiffrin & Schneider, 1977), Winter and Uleman (1984) raised the possibility that the process of trait inferences could be automatic.

Winter and Uleman (1984) introduced a new perspective about the dispositional inference process, according to which personality inferences are a frequent and non-optional process. This conceptualization was challenging, also in experimental terms, requiring the development of paradigms that were able to explore the implicit nature of the process.

2.2. Spontaneous Trait Inference Paradigms

Several paradigms have been developed in order to test the occurrence of spontaneous trait inferences. In all of them, participants are presented with sentences describing behaviors performed by hypothetical actors. Behaviors are pre-tested to clearly and unambiguously evoke a specific trait inference. For instance, the sentence “solves the mystery half-way through the book” describes a behavior that clearly implicates the trait *clever* (for other examples see Winter & Uleman, 1984, p. 241).

The central requirement of spontaneous trait inference paradigms is that they provide participants with a processing goal that directs their attention toward behavioral information, without explicitly instructing them to form an impression, or to infer a trait, about the actor. In addition, the activated processing goal should not indirectly

involve an explicit trait inference (see Uleman, 1999). In most spontaneous trait inference paradigms, memory instructions are applied because it is assumed that memorization demands attention toward behavioral information, without necessarily inducing explicit trait inferences. However, independently of theoretical pre-assumptions, proving that participants do not engage in explicit trait inferences, even in the absence of impression formation instructions is a problem inherent to this research. For example, if participants use explicit trait inference processes when asked to memorize the material, it cannot be concluded that trait inferences occurred spontaneously. For that reason, it is theoretically pertinent to clarify what kinds of processes are prompted by different processing goals (see Uleman & Moskowitz, 1994).

Four main paradigms have been used in the spontaneous trait inferences domain: the cued-recall paradigm (Winter & Uleman, 1984); the recognition probe paradigm (Newman, 1991; Uleman, Hon, Roman, & Moskowitz, 1996); the savings in relearning paradigm (Carlston & Skowronski, 1994); and the false recognition paradigm (Todorov & Uleman, 2002).

2.2.1. Cued-Recall Paradigm

The cue-recall paradigm was initially applied by Winter and Uleman (1984), and it was already described. Because it was the first procedure specifically applied to test STIs, the cued-recall paradigm has been one of paradigms most recurrently used in the literature (e.g., Claeys, 1990; Uleman, Winborne, Winter, & Shechter, 1986; Uleman, et al., 1993; Winter et al., 1985), both to provide support and to raise questions about the initial findings of Winter and Uleman (1984).

Winter et al. (1985), for example, adapted the paradigm by introducing a digit recall task. Participants were instructed to remember a series of digits, with trait-implicating sentences presented as distractors between digits presentation. Under these conditions there are no apparent reasons for participants to intentionally memorize the sentences or engage in any type of elaborative trait processing. Thus, this method represents an interesting way of testing the unintentional nature of the trait inference process.

In this study, the efficacy of trait-cued recall was compared with three different conditions: semantic-cued recall (i.e., semantic associates of the actors); gist-cued recall

(i.e., the title or theme of the sentence); and no-cued recall. Results replicated Winter and Uleman's (1984) findings, showing that trait-cues are more effective than both no-cues and semantic-cues. However, trait-cued and gist-cued recall didn't differ, with both being more efficient than other recall conditions. The same digit version of the cued-recall paradigm was used by Lupfer et al. (1990) and by Uleman et al. (1992), with similar results.

In this study, the efficacy of trait-cued recall (i.e., recall of the sentences with the implied traits provided as cues) was compared with three different conditions: semantic-cued recall (i.e., recall with semantic associates of the actors provided as cues); gist-cued recall (i.e., recall with the title or theme of each sentence provided as cues); and no-cued recall (i.e., recall without cues). These terms (trait-cued recall, semantic-cued recall, gist-cued recall, and no-cued recall) will be often used through our work to refer to the different types of cue recall conditions.

Uleman and Moskowitz (1994) further explored the effect of processing goals on the magnitude of STIs, using the original cued-recall paradigm (Winter & Uleman, 1984). Trait-cue efficacy was compared with gist, semantic, and no cue conditions. In the first experiment (Uleman & Moskowitz, 1994, Experiment 1) the usual memory condition was compared with a memory condition in which participants were instructed to ignore the meaning of the sentences. Results showed that trait-cued recall (as well as gist-cued and no-cued recalls) was higher with usual memory instructions than with ignoring-memory instructions, suggesting that a meaningful processing increases STIs. However, the overall evidence for the occurrence of STIs in this study was weak, for several different reasons. First, trait-cued predicate recall was more effective than no-cued and semantic-cued recall, but only with ignoring-memory instructions (not with the usual memory instructions). Second, actor trait-cued recall was worse than actor semantic-cued recall, in both instructions conditions. Thus, evidence for trait inferences was generally weak in this study.

In a second experiment (Uleman & Moskowitz, 1994, Experiment 2), participants were given three different processing goals: analyzing graphemes (i.e., locate specific letters in the sentences); analyzing phonemes (i.e., identify a particular sound); or analyzing a feature of each word (i.e., decide the gender of each word). Results showed that these three processing goals produced lower levels of recall (for all

cue conditions), than did the usual memory instructions (from the first experiment). This suggests that a memory goal involves a greater elaboration, which facilitates both trait and gist spontaneous inferences.

It was also found that trait-cued recall was higher under semantic processing conditions than under phonetic and graphemic conditions, respectively. In contrast, gist-cued recall showed an inverse pattern, being weaker with semantic than with phonetic and graphemic conditions, respectively. Based on these results, the authors argued that trait and gist inferences are guided by different mechanisms.

Finally, the authors used the usual criterion for testing the occurrence of spontaneous inferences (i.e., whether trait and gist cued recall were higher than other cue conditions). With a graphemic goal, only predicate gist-cued recall was higher than other cue conditions. Thus, there was evidence for gist, but not for trait, spontaneous inferences. With a phonemic goal, both trait and gist cued recalls were higher than other cue conditions, indicating both trait and gist spontaneous inferences. Finally, with a semantic goal, there was only evidence for trait inferences, with predicate trait-cued recall being more effective than both no-cued recall and semantic-cued recall.

In a third experiment (Uleman & Moskowitz, 1994, Experiment 3) the goal was to explore whether trait inferences would occur as a consequence of other non-trait social judgments. In this study, an explicit trait inference condition was compared with two non-trait judgments conditions: (a) a behavioral similarity judgment (i.e., how likely it would be that participants perform the same behavior as the actor), and (b) a personal similarity judgment (i.e., how similar participants were to the actor of the behavior). Results showed that trait-cued recall and gist-cued recall were not significantly different under the three processing conditions. In addition, under behavioral similarity instructions, there was evidence for both trait and gist spontaneous inferences, with higher predicate and actor recall in both gist-cue and trait-cue conditions, than in control conditions. Under personal similarity instructions, predicate trait-cued and predicate gist-cued recall was higher than recall in the control conditions, but actor recall was higher only with gist-cues. Finally, with an explicit impression formation goal, evidence for both trait and gist inferences were reported, both in predicate and actor recall.

From these three experiments, Uleman and Moskowitz (1994) concluded that the type of processing goal influences spontaneous inferences processes. Specifically, processing goals influence the magnitude of spontaneous inferences, the type of spontaneous inferences that are made (trait or gist), and the likelihood of linking inferred traits to actors.

Bassili and Smith (1986, Experiment 1) also applied the cued-recall paradigm, but added an explicit impression formation condition. Trait-cue recall was significantly more effective with explicit instructions. The authors concluded that the magnitude of trait inferences occurring spontaneously is limited and cannot be comparable with the magnitude of explicit trait inferences. Since STIs are influenced by processing goal, Bassili and Smith argued that they should not be characterized as an automatic process. Using the same paradigm, Claeys (1990) replicated the initial findings of Winter and Uleman (1984), demonstrating that trait-cued recall is more effective than semantic-cued recall (using associates of the last words of the sentences). Further, like Bassili and Smith (1986), they showed that trait-cued recall was higher with an explicit categorization instruction, than with memory instructions. Again, this supports the idea that trait inferences are less prevalent in the absence of explicit instructions.

In summary, in most of the studies using the cued-recall paradigm, trait-cued recall is at least as effective as semantic-cued recall, and more effective than no-cued recall. However, the question about whether trait inferences are linked to the representation of the actor cannot be clarified. There is weak evidence of higher actor recall under trait-cued than under control conditions (e.g., Winter & Uleman, 1984; Winter et al., 1985). This has been one of the major difficulties of the cued-recall paradigm. A second problem is that participants under memory instructions can engage in impression formation processes in order to facilitate recall (Hamilton, 1981). A third limitation is that results obtained from the trait-cued paradigm can be alternatively explained by retrieval processes (e.g., Wyer & Srull, 1989). Finally, D'Agostino and Beegle (1986) argued that a problem with the trait cued recall paradigm is that cues are typically manipulated within participants, with trait-cued recall preceding no-cued recall. Because of that, the recall superiority typically observed for trait-cued recall may be explained by output interference effects (e.g., Nickerson, 1984; Roediger, 1974). Specifically, the previous trait-cued recall may interfere with the subsequent recall

without cues. According with their formulations, D'Agostino and Beegle (1996) found that the typical trait cue advantage is eliminated when cues are manipulated between participants.

2.2.2. Recognition Probe Paradigm

In the recognition probe paradigm (e.g., Ham & Vonk, 2003; Newman, 1991, 1993; Uleman et al., 1996; Van Overwalle et al., 1999; Wigboldus et al, 2003; Wigboldus, Sherman, Franzese, & Van Knippenberg, 2004), participants are presented with trait-implying sentences (e.g., “wins the science quiz”), and each sentence is followed by a probe word. In the experimental trials, probe words are the traits implied by the sentences (e.g., “smart”). Participants’ task is to indicate whether the probe word was part of the previous sentence, as fast and as accurately as possible. If the trait is spontaneously inferred from the behavior, it will harder to indicate that it was not part of the sentence. Thus, longer response times and more errors are expected when trait-implying sentences are followed by implied trait probes than when the same sentences are followed by neutral probe words. Despite predictions being made both in terms of error rates and response times, patterns of results are not consistent. Some times there are differences in response times and not in errors (McKoon & Ratcliff, 1986; Wigboldus et al., 2003), while in other studies differences are obtained in errors but not in responses times (see Uleman et al., 1996, Experiment 1). In part, these mixed patterns may be explained by a speed-accuracy trade off (Reed, 1973) (i.e., people may choose to respond faster with less accuracy or slower with more accuracy). However, the speed-accuracy trade-off seems not to be sufficient to explain all the patterns observed (see Uleman et al., 1996). Thus, the reasons for the observed inconsistent patterns across studies in the probe recognition paradigm are not completely understood.

The recognition probe paradigm was imported from text comprehension literature, which usually applied it to the study of online inferences (e.g., Corbett & Chang, 1983; McKoon & Ratcliff, 1980, 1986). For example, McKoon and Ratcliff (1980, Experiment 1) used the paradigm to investigate the occurrence of predictive inferences. Participants were presented with a predictive paragraph, and immediately after the paragraph, a probe word was presented. In one condition, the paragraphs described a situation that predicted an event captured by the probe word, while in the

other condition (i.e., control condition) paragraphs were rearrangements of the same words but in a way that did not result in the probe word prediction. It was shown that participants take more time to correctly say that the probe word was not in the paragraph in predictive than in control conditions.

Uleman et al. (1996) used a version similar to the one used by McKoon and Ratcliff (1980), but substituted trait implying paragraphs for the predicting paragraphs. In the first experiment, trait-implying paragraphs resulted in more errors for implied trait probe words than control paragraphs. In the second experiment, feedback on speed and accuracy was introduced in order to increase participant's motivation. This time, the predicted pattern emerged in terms of response times, with participants taking more time to correctly reject the trait after trait implying paragraphs than after control paragraphs.

Contrary to what might happen in the cue-recall paradigm, success in the recognition probe task works in opposition with the trait inferential process. That is, the more participants engage in trait inferences, the more difficult will be for them to reject the trait. This means that, if participants want to be strategic in this task, they should *avoid* making trait inferences. Thus, evidence of trait inferences with this paradigm is compelling in proving the unintentional nature of the process.

However, some concerns have also been raised regarding this paradigm. Specifically, it was argued that results arising from this paradigm can be explained by context-checking processes occurring at the time the probe word is presented (e.g., Forster, 1981; McKoon & Ratcliff, 1986; but see McKoon & Ratcliff, 1989b). According with this explanation, when the probe word is presented, participants compare the meaning of the probe with the meaning of the previous sentence. Since a high compatibility between exist between the two, rejecting the probe is more difficult. However, the fact that fast responses are required makes less likely the possibility that participants engage in reconstructive strategies when the probe word appears. In agreement with this idea, McKoon and Ratcliff (1986) defended that a deadline response procedure can be used in order to clearly classify participants' responses as automatic.

2.2.3. Savings in Relearning Paradigm

The savings paradigm is a classic research method developed by Ebbinghaus (1885/1964) in order to study the nature of memory and learning processes. Participants learn information, and after different periods of time they are instructed to relearn the same information. A *savings* effect is typically observed in relearning trials, correspondent to a decrease in the time needed to learn the information. This learning advantage is observed even years after the first learning episode, and is independent of conscious recollection of information (Burt, 1941; Titchener, 1923).

Carlston and Skowronski (1994) adapted the paradigm and applied it to the STI domain. The typical procedure used in the STI domain is divided in three main phases. In the first phase (exposure phase), participants are presented with series of photos, each one paired with a behavioral sentence. In the second phase (learning phase), participants are presented with pairs of photos and traits. Some of them are *relearning pairs*, composed by pairing previously presented photos with the traits implied by the behaviors previously paired with those photos. Others are *control pairs*, which consist of new photos and traits. In the third phase (the recall phase), participants are given photos presented during the learning phase, and are asked to recall the corresponding traits. It is assumed that, if participants had spontaneously inferred traits from behaviors during the first exposure phase, relearning pairs should be easier to learn than control pairs.

In the first experiment, Carlston and Skowronski (1994, Experiment 1) applied different sets of instructions: trait generation (“think in a trait that describes the behavior”); impression formation (“form an impression about the actor”), or familiarization instructions (“read and get familiar with the material”). Results showed that recall was significantly higher in relearning trials than in control trials (reflecting a savings effect), and the pattern didn’t differ across instruction conditions. A savings effect was also demonstrated using memory instructions, although this effect was slightly smaller than with impression and familiarization instructions (Carlston & Skowronski, 1994, Experiment 4).

Carlston, Skowronski, and Sparks (1995) adapted the paradigm in five additional experiments, with the savings effect being replicated in all of them. First, it was shown that the magnitude of the savings effect is similar if participants are forced to infer a

trait from the behavior during the exposure phase (Carlston et al., 1995, Experiment 1). However, the savings effect is only significant when the trait inferred is the same as the one that is presented as cue (Carlston et al., 1995, Experiments 1 and 2). The savings effect is not observed if the trait inferred doesn't match the one relearned, even when the trait is evaluatively similar to the one that is relearned (Carlston et al., 1995, Experiment 2).

Second, the savings effect is observed both when descriptions of behaviors are presented as self descriptions, as in the traditional paradigm (e.g., "I don't like to spend money"), as well as when behaviors are presented in the third person form (e.g., "He does not like to spend money") (Carlston et al., 1995, Experiment 3).

Third, a savings effect for learning trait-photos pairs is observed even when the photos are said to be of individuals that are merely *describing* other's behaviors (Carlston et al., 1995, Experiment 4). This surprising effect became known as the Spontaneous Trait Transference (STT) effect (Carlston & Skowronski, 2005; Crawford, Skowronski, & Stiff, 2007; Crawford et al., 2008; Mae, Carlston, & Skowronski, 1999; Skowronski et al., 1998). Under these conditions, participants should not apply the trait implied by the behavior to describe the person portrayed in the photo because they know that the person is merely describing someone else's behavior.

Proponents of the savings in relearning paradigm have argued that one of advantages of this method consists in proving that trait inferences are descriptions of the actors, and not merely behavior descriptions. Thus, it was argued that the savings in relearning paradigm was more adequate to explore the spontaneity of the trait inference process. Ironically, experiments using the savings and relearning paradigm ended up revealing a great difficulty of spontaneous trait inference research (Carlston et al., 1995). Specifically, the fact that a savings effect occurs even when the picture portrays an actor describing a third person's behavior raises the possibility that spontaneous trait inferences result from *non-inferential* associative processes. For that reason, subsequent studies using the savings in relearning paradigm have been mainly focused in clarifying the conditions under which trait transference effects occur (Mae et al., 1999; Skowronski et al., 1998), and in testing whether STT and STI are different cognitive processes (Carlston & Skowronski, 2005; Crawford, Skowronski, & Stiff, 2007;

Crawford et al, 2007; Crawford et al., 2008). This question will be further addressed later.

2.2.4. False Recognition Paradigm

Todorov and Uleman (2002) were the first to apply the false recognition paradigm to the study of spontaneous trait inferences. In this paradigm, participants are presented with pairs of photos and behaviors under memory instructions. In a second phase, the previous photos are presented, each one paired with a trait. In some of these trials, each photo is paired with the trait implied by the behavior previously paired with that photo (*match* trials). In other control trials, the photos are randomly paired with traits implied by behaviors previously paired with other faces, or photos are paired with new traits (*mismatch* or *new* trials). Participants have to respond, as fast and accurately as possible, whether the presented trait was included in the sentence previously paired with the presented photo. If spontaneous trait inferences had occurred during initial encoding, and became linked to the actors, participants should be more likely to falsely recognize traits in match than in mismatch trials.

The first study using the false recognition paradigm revealed a number of different aspects (Todorov & Uleman, 2002). First, it was confirmed that participants are more likely to falsely recognize traits that are paired with the actor of the corresponding behavior than traits that are randomly paired with different faces (Todorov & Uleman, 2002, Experiments 1).

Second, when a control condition with new traits paired with previous photos is incorporated, the percentage of false recognitions is higher in the mismatch condition (in which previously implied traits are randomly paired with the photos) than in this new control condition. This result suggests that a trait inference increases the number of false recognitions, regardless of being associated with the actor (Todorov & Uleman, 2002, Experiment 2).

Third, when photos are paired with antonyms of the traits implied by the behaviors, participants are less likely to falsely recognize the antonym traits, than new traits (Todorov & Uleman, 2002, Experiments 3). Todorov and Uleman interpreted this result as evidence that the inferred trait is incorporated into the mental representation of the actor and therefore can affect subsequent trait judgments.

Fourth, in a further experiment, memory and impression formation instructions were compared (Todorov & Uleman, 2002, Experiments 4). Another difference in the procedure was that an exposure time of 10 seconds was incorporated, in addition to the typical 5 seconds exposure time. Under memory instructions, false recognition of traits was higher in the 5s condition than with 10s, while under impression formation instructions the exposure time didn't make a difference. These data was taken as an indication that memory and impression participants are using the extra time engaging in different mental operations. While memory participants supposedly try to retain the exactly presented sentence, impression participants use the additional time to engage in extra elaborative processes.

Todorov and Uleman applied the false recognition paradigm in a number of subsequent studies. In three of these experiments, the efficacy of the STI process was tested (Todorov & Uleman, 2003, Experiments 1, 2, and 3), either by decreasing the time of stimulus presentation to 2 seconds (Experiment 1), by inducing a grammatical processing goal (Experiment 2), or by imposing a recall digit concurrent task (Experiment 3). If spontaneous trait inferences occur independently of available attentional resources, the typical false recognition effect (i.e., higher number of trait false recognitions when the trait is paired with the actor than when traits are randomly paired with another actor) should be independent of these manipulations. Results confirmed predictions, showing a trait false recognition effect under rapid stimulus presentation, with shallow processing (despite the effect being weaker in this case), and under cognitive load conditions. Importantly, the trait false recognition effect occurred regardless of the overall level of recognition being weaker in any one of these three conditions. This evidence was taken as suggesting that the process of making trait inferences about actors is a highly automatic one.

Authors using the false recognition paradigm (Todorov & Uleman, 2002, 2004) have argued that one of its advantages is that it pits the demands of the task against the trait inference process. The logic of the task is similar to the one underlying the recognition probe paradigm (e.g., Uleman et al., 1996). However, while in the recognition probe paradigm the trait is presented immediately after each sentence, in the false recognition paradigm traits are presented only after the entire set of sentences. For this reason, the opposed nature of the recognition task is less clear-cut in proving the

unintentional nature of the process in this case. At encoding, participants have no knowledge that the best way to perform the task is remembering the exact words that form the sentence, and may probably engage in elaborative strategies in order to facilitate recall. In addition, if participants explicitly make trait inferences during encoding, they may use that information to facilitate the monitoring recognition decision (Johnson, 1997; Johnson, Hashtroudi, & Lindsay, 1993). That is, they may remember that they explicitly inferred the trait and may use that knowledge to correctly reject the trait. Thus, it is not clear that trait inference and task requirements work against each other in this case, or if trait inferences may even facilitate performance.

Another concern about the false recognition paradigm is that the presentation of pictures may automatically trigger impression formation processes (see Bargh, 1990, Bargh & Barndollar, 1996; Chartrand & Bargh, 1996), a difficulty that is also inherent to the savings in relearning paradigm.

In contrast, one of the great advantages of the false recognition paradigm is that it is highly effective in proving the existence of an associative link between the trait and the actor. However, as with the savings in relearning paradigm (Carlston & Skowronski, 1994), it is not clear whether the trait-actor link has an associative or inferential nature. This is a crucial issue, and it will be discussed later in greater detail.

2.2.5. Other Paradigms

The previous described paradigms have been the ones mostly frequently applied in the spontaneous trait inference domain. However, explicit memory measures were used in most of them. Participants are explicitly asked to remember a prior known episode. An exception is the savings in relearning paradigm (Carlston & Skowronski, 2004), but even in this case the relation between the initial exposure phase and the task is easily apparent. Implicit memory measures indirectly examine the effects of a prior episode, without participants being aware that their memory is being tested (for reviews on implicit memory see Richardson-Klavehn & Bjork, 1988; Roediger, 1990; Schacter, 1987; Shimamura, 1986). For that reason, in studies in which implicit measures are employed no reference is made to the relation between learning and test phases, and the existing relation is masked.

A number of studies have made use of implicit memory measures to explore the spontaneity of the trait inference process. For example, Lea (1995) used a lexical decision task in order to measure trait activation after participants read paragraphs that could include, or not, a trait-implying sentence. In some of the paragraphs, the trait-implying sentence was incorporated within a context that facilitated the trait inference, while in others the context undermined the trait inference. As predicted, participants took less time to decide if the implied trait was a word after context supporting paragraph, than after inhibiting paragraphs. Since lexical decision time is used as a measure of concept accessibility (but see Balota & Chumbley, 1984), results were interpreted as evidence that spontaneous trait inferences are more likely to occur when a behavior is included in a congruent context.

Whitney and Williams-Whitney (1990) presented participants with a word-stem completion task (e.g., CL _____ for CLUMSY) immediately after participants read trait implying, or control paragraphs. Participants were instructed to complete the word-stem with the first word that comes to their mind. Results showed that word-stems were more likely to be completed with the implied traits after trait implying paragraphs than after control paragraphs.

In a subsequent study, however, Whitney, Waring, and Zingmark (1992, Experiment 1) used a stem completion task and didn't find evidence for STIs. Participants were asked to read trait implying paragraphs with different processing goals: verbatim recall, factual recall, elaborative recall, and impression formation. After each paragraph participants performed a stem completion task. In this study, evidence for spontaneous trait inference was reported only when participants engaged in elaborative processing goals (i.e., elaboration and impression formation instructions).

The reported studies made use of implicit memory measures. However, they are not immune to contamination problems (Jacoby, 1991). Specifically, if participants became aware of the existing relation between the two tasks they could intentionally use conscious strategies in order to improve performance in the supposed implicit task. Since performance in both lexical decision and stem-completion tasks is facilitated by explicit trait inferences, there is the risk that participants may consciously engage in trait inferences.

In sum, several paradigms have been applied to study the occurrence of spontaneous trait inferences. Different limitations have been pointed out to all paradigms, and all of them seem to allow alternative interpretations. In any case, the large number of existing studies certainly provided us important data about the trait inference process, and helped us to better understand whether the process can be defined as automatic, as initially suggested by Winter and Uleman (1984). We move now to a closer analysis of existing data about the automaticity nature of the STI process.

2.3. The automaticity of STIs

According to some authors, the spontaneous trait inference process can be characterized as an automatic process. For example, Winter and his collaborators concluded that “inferences about traits can be initiated in a largely automatic way” (Winter et al., 1985, p.914), and Todorov and Uleman (2003, p. 561) suggested that “much of the process of binding spontaneous trait inferences to actors representations can be characterized as automatic”.

The automatic perspective is controversial, since it implies that participant’s may have less control over their impressions than they might think. Given the important implications of this question, it is worthy to consider the automaticity issue in detail. In order to do that, we will analyze whether the trait inference process satisfies the four most commonly mentioned criteria for automaticity (e.g., Bargh, 1994): lack of intentionality, lack of awareness, efficiency, and uncontrollability. At the end, we will conclude that the existing experimental evidence does not support a characterization of the trait inference process as fully automatic, and we will discuss how the spontaneous trait inference process is better understood as a *conditional* process (Bargh, 1989).

2.3.1. Intentionality

The intentionality criterion has to do with the control that we have over the start of a process (Bargh, 1994; Zbrodoff & Logan, 1986). We can either intentionally engage in a process, or the process can occur without our intention. If a process takes place in the absence of conscious intention, it is said to be *spontaneous* (Uleman, 1999).

Following Uleman's definition (Uleman, 1999), the trait inference process is *intentional* when it is activated by an explicit impression formation goal, while it is *spontaneous* when it is triggered by the mere attention to behavioral information. Thus, in order to prove the lack of intentionality (or spontaneity) of the trait inference process, research must demonstrate that trait inferences occur even when participants do not have the goal of forming an impression, or of making a trait inference. Usually, participants are given processing goals that demand attention to the behavioral information, without necessarily involving explicit trait inferences.

Research made using the trait-cued recall paradigm with memory instructions has shown that traits are effective cues to behavior recall (e.g., Winter & Uleman, 1984). Since participants are instructed only to memorize the information, and excluding the possibility of retrieval accounts, this pattern is taken as evidence for the unintentional nature of trait inferences.

However, by showing that the efficacy of trait-cues is much lower under memory instructions than under explicit trait inferences conditions, both Bassili and Smith (1986, Experiment 1), and Claeys (1990) questioned the spontaneity of the trait inference process. According to Bassili and Smith (1986), if a process is automatic, it should not be influenced by an instruction to explicitly engage on it. On the contrary, if the magnitude of a process is highly different under intentional and spontaneous conditions, as was observed for trait inferences (Bassili & Smith, 1986; Claeys, 1990), the automaticity of the process is questionable.

Differences in the magnitude of trait inferences under memory and impression conditions were not always confirmed when other paradigms were applied. For example, using the false recognition paradigm, Todorov and Uleman (2002, Experiment 4) found that the false trait recognition effect was not significantly different under memory and impression formation conditions. Additionally, using the savings in relearning paradigm, Carlston and Skowronski (1994, Experiment 4) compared memory, impression formation, and familiarization instructions, and found no effect of instructions in the observed savings effect. Also, no differences were detected between familiarization, explicit trait generation, and impression formation instructions (Carlston & Skowronski, 1994, Experiment 1).

How can we explain these different patterns of results, when different paradigms are used? A main difference between the false recognition and the savings in relearning paradigms, on one side, and the cued-recall paradigm, on the other side, is the presentation of faces. The strong individuating character of the pictures may possibly trigger explicit impression formation processes. Thus, it is not clear whether the lack of differences observed by Todorov and Uleman (2002) and Carlston and Skowronski (1994) proves that the magnitude of trait inferences is equally strong under spontaneous and intentional conditions, or whether results reflect a stronger influence of explicit processes under memory conditions when pictures are incorporated.

The fact that participants may be engaging in impression formation processes when instructed to memorize behaviors, or when simply asked to get familiar with photo-behavior pairs, constitutes one of the major problems of proving the unintentionality of the trait inference process. A number of studies directly addressed this question.

Some of them used the cued-recall paradigm, but without instructing participants to memorize behaviors. In these studies, behaviors were presented as mere distractors between trials of a digit recall task (Lupfer et al., 1990; Uleman et al., 1992; Winter et al., 1985). Results showed that trait-cues were more effective than control conditions, indicating the occurrence of spontaneous trait inferences. However, the level of trait-cued recall was lower than what is usually observed with the traditional procedure, which shows that the magnitude of spontaneous trait inferences is not totally independent of processing goals.

There are different possible interpretations for why, specifically, evidence for spontaneous trait inference is weaker when behaviors are presented as distractors. First, it could be that participants pay less attention to behaviors when they are presented as distractors. Since there are no meaningful reasons to process the distracting behaviors, some of them may be ignored. However, the procedure required that participants read each sentence aloud twice, which doesn't provide support for this possibility. Participants are not merely ignoring the sentences. A second possible interpretation is the previous notion that memory participants' may be engaging in explicit trait inferences. That is, conscious trait inference processes may be influencing performance (Jacoby, 1991). Finally, a third interpretation is that, even if participants *are not*

engaging in explicit trait inferences, the higher level of semantic processing involved in memorization facilitate the occurrence of spontaneous trait inferences.

The last interpretation is in agreement with existing studies that show that the magnitude of spontaneous trait inferences is usually weaker under shallow processing condition, than under memory conditions (Uleman & Moskowitz, 1994). Moreover, there is no evidence for spontaneous trait inferences with a graphemic processing goal (in which participants are asked to search for a specific letter within the sentences).

Again, data arising from other paradigms fail to replicate the same pattern of results. Using the false recognition paradigm, Todorov and Uleman (2003, Experiment 2) reported a false trait recognition effect even under shallow processing conditions (i.e., counting nouns in the sentences), although the effect was weaker than under memory instructions. Todorov and Uleman (2003) took this result as providing “the strongest evidence that STIs in the false recognition paradigm are unintentional” (p.560). However, it should be noticed that instructions, in this experiment, induced participants to pay attention to behaviors. In the procedure section, it is explained that: “in the shallow processing condition (...) participants were also told to read the information carefully because they would be asked additional questions at the end of the experiment” (Todorov & Uleman, 2003, p.554). Thus, it is questionable whether this can be considered a mere shallow processing requirement.

Probably the most compelling paradigm in order to prove the unintentional nature of the trait inference process is the recognition probe paradigm (excluding alternative retrieval explanations for STI occurrence). In this paradigm, the strategy most adequate for participants would be to memorize the presented sentences, without engaging in any type of inferential activity. Since evidence for trait inference is usually obtained (Uleman et al., 1996) this suggests that trait inferences can occur spontaneously.

In sum, despite the fact that the label “spontaneous trait inference” has been used to designate this field of research, evidence for the lack of intentionality of the trait inference process is much less clear than we may think. Evidence varies greatly with the paradigm that is applied, and in most paradigms the possibility that explicit impression formation processes are intervening cannot be rejected. The more

straightforward paradigm to prove that trait inferences can occur spontaneously seems to be the recognition probe paradigm (e.g., Uleman et al., 1996).

However, even if trait inferences can occur without intention, its occurrence seems to be much weaker than when they occur intentionally (Bassili & Smith, 1986; Claeys, 1990), and is dependent on some level of semantic processing (Uleman & Moskowitz, 1994). This is line with conclusions from Bargh (1989), who refers that a meaningful understating of the sentences seems to be a pre-condition for spontaneous trait inference occurrence.

2.3.2. Awareness

The question about whether perceivers are cognitively aware of making trait inferences has been approached in several studies. Winter and Uleman (1984) observed that most of their participants reported no awareness of making causal or personal judgments during behavior encoding. Moreover, no significant correlation was observed between reported trait inferences and trait-cue efficacy, indicating that participants' awareness is not a pre-condition for the occurrence of trait inferences.

However, as Winter and Uleman (1984) pointed out, the 10 minutes interval between behavior presentation and the awareness report (presented at the end of the study) may have undermined participants' ability to remember the mental processes engaged during behavior encoding. In order to try to provide a more accurate measure of awareness, Winter et al. (1985) introduced an awareness questionnaire³ immediately after the last behavior, and asked participants about the thoughts they had during reading it. Again, the reported level of awareness about causal or personality thoughts was minimal. This time, a marginally significant correlation was obtained between awareness and trait-cued recall (taking off the recall data of the last sentence). However, the pattern of results interacted with participant's gender. Causal thoughts awareness was superior for women than men, and the correlation between awareness and trait-cued

³ This is usually considered an awareness measure. However, since verbal reports examine the type of mental operations used by participants during the task, they can also inform us about the intentionality of the process. However, the fact that an awareness questionnaire can be informative about the intentionality of a cognitive process doesn't mean that awareness and intentionality are intrinsically related. It is possible that we have conscious awareness of a process that was not intentionally activated, and we can intentionally engage in a process without being aware of its operation.

recall was only significant for women. The reasons for this difference are not clear, but it opens the possibility of existing individual differences in the probability of STIs.

Uleman et al. (1986), using the cued-recall paradigm, applied both an awareness general questionnaire at the end of the study (that included both open ended questions and rating scales) and an open ended question after the last sentence being read. Each awareness measure was applied to different sets of behaviors. In the rating scales, about half of the participants indicated having made personality inferences, both during encoding and during retrieval. In contrast, in the open ended questions participants rarely mentioned trait thoughts.

In this study, a number of correlations between awareness and trait-cued recall were found to be significant. First, a correlation between trait thoughts occurring during the last sentence presentation and the percentage of trait-cued recall was observed, but only for high authoritarianism sentences, and not for low authoritarianism sentences (in this experiment authors were interested in studying the role of authoritarianism as an individual difference variable in the magnitude of STIs). Second, low authoritarianism participants showed a correlation between trait thoughts reported on the general open ended question and trait-cued recall of low authoritarianism sentences. Finally, high authoritarianism participants exhibited a correlation between personality inferences on rating scales and trait-cued recall of high authoritarianism sentences.

Lupfer et al. (1990), Moskowitz (1993a), and Uleman and Moskowitz (1994) all opted to apply an awareness questionnaire after the presentation of the last sentence. In Lupfer et al. (1990) almost all participants reported no awareness of making personality inferences, and no correlation was observed between awareness and trait-cued recall efficacy. In the Moskowitz (1993) study, participants reported using trait inference strategies as much as semantic and imagistic strategies, but no correlation was observed between trait awareness and percentage of trait-cued recall. Finally, in the Uleman and Moskowitz (1994) study, only about 14% of the participants reported trait thoughts. In this case, however, correlations reflected participants' correlation accurate awareness. Trait thoughts reported both on rating scales and on open-ended questions correlated with trait-cue recall of the actors. Trait-cued recall of the predicate was correlated with trait thoughts.

Studies that used awareness questionnaires are not totally clear in relation to this issue. In any case, self-report measures must be evaluated with caution. It is well established that participants' introspective access to over their cognitive processes is limited (Nisbett & Wilson, 1977). Thus, in order to interpret the meaning of a self-report measure, we should accurately understand the mechanisms in which responses are being based (Ericsson & Simon, 1980). In the studies reported here, it is possible that instead of trying to retrieve from memory, participants base their responses on intuitive causal theories about the processes that they might use in such circumstances (Nisbett & Wilson, 1977). Also, self reports may be restricted by perceived demanding aspects of the experimental situation. In this case, participants may not want to say that they used strategies not explicitly mentioned by the experimenter. Finally, an additional problem in studies that introduce an awareness questionnaire after the last sentence (between behavior presentation and recall) is that noise may be introduced into the procedure. Participants may become aware of the goal of the study, and be more likely to use explicit retrieval strategies.

Due to the demanding aspects of self-reports, Moskowitz and Roman (1992) argued that unobtrusive awareness measures are more adequate. In Moskowitz and Roman's (1992) study, participants were instructed to read a set of behavior, under either memory or impression instructions. It was demonstrated that while memory participants show an assimilation effect in the evaluation of a subsequent ambiguous target (i.e., *Donald*, see Higgins, Rholes, & Jones, 1977), impression formation participants exhibit a contrast effect. As previous studies (Lombardi, Higgins, & Bargh, 1987) reported contrast effects when participants remembered the prime, and assimilation effects when participants were unable to remember the prime, Moskowitz and Roman explained their data by asserting that STIs occur without awareness under memory instructions, and under participant's awareness in impression formation conditions. It should be noted, however, that awareness is only one of the possible factors used to explain assimilation and contrast effects (see Banaji, Hardin, & Rothman, 1993; Ford & Kruglanski, 1995; Herr, 1986; Sedikides, 1990; Schwarz & Bless, 1992; Thompson et al., 1994).

Stapel, et al. (1996; see also Stapel, Koomen, & Van der Pligt, 1997) presented an alternative interpretation to Moskowitz and Roman's (1992) results. According to

them, when a general trait category is activated (e.g., friendly), assimilation effects are expected to be observed in subsequent judgments, as those are typically found in category accessibility studies (e.g., Higgins et al., 1977; Bargh & Pietromonaco, 1982; Srull & Wyer, 1979). On the other hand, the activation of specific actor-trait links (e.g., John is friendly) tends to be used as a comparison standard in relation to which subsequent actors are contrasted. Stapel and colleagues (Stapel, et al., 1996) argued that, in a memory condition, behaviors tend to activate general trait descriptions, while under impression formation conditions a specific actor-trait link tends to be established. This would explain the assimilation and contrast pattern obtained by Moskowitz and Roman (1992), without postulating awareness as a determinant factor.

Stapel and colleagues (Stapel, et al., 1996) provided support to their account by showing that when personal information that favours the establishment of a trait-actor link is incorporated (i.e., personal names and photos), contrast effects are obtained, even for memory participants (Experiment 1). Moreover, personal contexts (that favor actor-trait links) lead to contrast effects, while situational contexts (that favor the mere activation of behavioral labels) result in contrast effects (see Bassili & Smith, 1989^b), these effects being independent of type of instructions. Notice that, the fact that under some circumstances, memory participants show contrast effects and impression participants show assimilation effects cannot be accounted by an awareness-based explanation (Moskowitz & Roman, 1992).

This study is crucial in its implications. If we accept for Stapel et al.'s (1996) interpretation, we support the notion that trait inferences occurring under memory instructions are merely behavior descriptions. On the other hand, if we claim that conditions in which contrast effects are observed under memory instructions (i.e., when a photo, or a supporting personal context are included) (Stapel, et al., 1996) make participants aware of making trait inferences, then it means that trait inferences do not occur spontaneously under these conditions. Either interpretation is controversial to spontaneous trait inference research.

Again, despite the fact that most authors within the spontaneous trait inference domain claim that spontaneous trait inferences occur without perceiver's awareness, empirical evidence regarding this question is scarce and contradictory. While in some studies participants report no awareness (Lupfer et al., 1990; Uleman & Moskowitz,

1994; Winter & Uleman, 1984; Winter et al., 1985), in others trait inference awareness is reported (Moskowitz, 1993a; Uleman et al., 1986); while in some studies no correlation is found between reported awareness and trait-cued recall efficacy (Lupfer et al., 1990; Moskowitz, 1993a; Uleman et al., 1992; Winter & Uleman, 1984), in other studies correlations are found to be significant (Uleman & Moskowitz, 1994; Uleman et al., 1986; Winter et al., 1985). Maybe these disparities are explained by specific differences between paradigms which, if being true, would indicate that awareness is a feeble characteristic of the trait inference process. Given the difficulties underlying self-report measures (Nisbett & Wilson, 1977), the use of alternative unobtrusive measures is important. However, also in this case, the existing evidence is limited and weak (Moskowitz & Roman, 1992).

2.3.3. Efficiency

A cognitive process is efficient when it doesn't require attentional resources (Bargh, 1994). If inferring a trait from a behavior is an efficient process, it should not be dependent on available cognitive resources, and should not be interfered by concurrent tasks (Logan, 1979).

Winter et al. (1985) presented behaviors as distractors between trials of a digit-recall task, and manipulated the difficulty of the digits. Results showed that presenting easy or difficult series of numbers had no influence in the magnitude of spontaneous trait inferences (for similar results, see Lupfer et al., 1990). In the same way, no other recall-cued condition (semantic, gist, and no-cue) was affected by the load manipulation. This might be explained because all inferences or associations are efficient, or because the load manipulation was not effective. Actually, although participants rating difficult series as more difficult, and recalling less digits in the difficult series conditions, no direct measure of spare capacity was applied.

In a subsequent study, Uleman and collaborators (Uleman et al., 1992) tried to clarify whether the load manipulation used by Winter et al. (1985) was in fact depleting participant's resources. They introduced two modifications to the previous paradigm. First, during the digit recall task, participants had also to monitor a white light, and press a key whenever the light appeared. If the depletion manipulation is working, participants should take more time to press the key in the high load condition. Second,

an easier load condition, formed by series of the same number (e.g., 24-24-24-24-24), was added to the other two previously used conditions (single digit condition, and multiple digit condition)

Results showed that participants take less time to respond to the probe light in the same-number condition than in the single-digit and multiple-digit conditions, with no differences between the last two conditions. The lack of differences between these two conditions used by Winter et al. (1985) suggests that cognitive load was not effectively manipulated in their study. In terms of recall, the difficulty of the digit task interfered with cued-recall performance. Specifically, trait-cue efficacy was worse as the difficulty of the digit series increased (Uleman et al., 1992). Thus, contrary to Winter et al. (1985), results indicate that trait inferences are influenced by the available cognitive resources. Based on these findings, the authors concluded that “trait inferences are not automatic by a capacity criterion” (Uleman et al., 1992, p. 89).

Data from Winter et al. (1985) and Uleman et al. (1992) are in apparent contradiction. Notice that results from Uleman et al. (1992) show that the load manipulation used by Winter et al. (1985) was not effective, as measured by the probe light concurrent task. However, the same load manipulation that didn't affect trait-cued recall in Winter et al. (1985) had an effect in Uleman et al. (1992) study. Uleman and collaborators (1992) suggested that these different results can be explained by the fact that the concurrent probe task adds an extra cognitive load. Thus, in the study from Winter et al. (1985), participants have in fact less resources in the difficulty than in the easy condition. However, in both cases, the remaining resources are still enough for the trait inference process. The impact of the difference in resources for STIs is only detectable when an additional load is introduced by the probe task.

In opposition to these results, Todorov and Uleman (2003, Experiment 1), using the false recognition paradigm, showed no differences in the magnitude of trait false recognition between a self-paced condition and a rapid sentence presentation condition. Fast stimulus presentation is generally used as a method of attentional overload (Bargh, 1994). In this study, however, the rapid condition was compared with an unusual self-paced condition, and not with the typical 5 seconds presentation condition. This might be an inadequate control condition, since it probably facilitates performance. Thus, we

cannot know whether the spontaneous trait inference magnitude differed in *both* conditions, compared with the typical 5-seconds condition.

Todorov and Uleman (2003, Experiment 3) further showed that a concurrent digit recall task didn't affect the number of false trait recognitions. In this study, participants were presented with a number before the behavior, and after the behavior they were given a number and asked to indicate whether it was the same or it was a different number from the one previously presented (by pressing different keys). Compared with the cognitive load task used by Uleman et al. (1992), this task is less resource consuming, which might explain the different pattern obtained.

Using the savings in relearning paradigm, Crawford et al. (2007) also reported no cognitive load effects on the magnitude of savings effects. Here, each behavior-photo pair was presented for 12 seconds. The load manipulation was induced by presenting a number and by asking participants to recall the number after each pair disappeared.

Finally, Wigboldus et al. (2004) put participants under high or low cognitive load while they were performing a recognition probe task. In the control condition (in which the target was described as "human") it was found that participants take more time to correctly say that the trait probe was not presented in the sentence under high than under low load conditions (Wigboldus et al., 2004, Experiment 2). However, the procedure doesn't allow us to clarify whether the cognitive load is directly affecting the trait inference process or making the task more difficult in general. A better way to test the efficiency of the trait inference process within this paradigm would be to request the recall of the number before the probe word presentation and not after, as the authors did.

Based on this set of studies, we can conclude that the spontaneous trait inference process presents some degree of cognitive efficacy (Crawford et al., 2007; Todorov & Uleman, 2003). However, the process is not totally independent of available cognitive resources, being disrupted under very high load conditions (Uleman et al., 1992).

2.3.4. Controllability

Controllability has to do with the amount of capacity to moderate or eliminate the operation of a cognitive process (see Bargh, 1994; Logan & Cowan, 1984). In order to control a process, perceivers must be aware of its influence (Strack & Hannover

1996). Awareness is, thus, a necessary but not a sufficient condition for control to be exerted. Motivation and capacity to override the process are also necessary elements (Devine, 1989)

Few studies have directly addressed the question of whether spontaneous trait inferences are controllable. In order to try to control spontaneous trait inferences participants must, in the first place, become aware of their possible influence, and then be motivated to eliminate their operation. Due to its natural features, the recognition probe paradigm (Newman, 1991, 1993; Uleman et al., 1996) might produce such conditions. Optimal performance in this paradigm implies that participants prevent trait inferences (in fact, not only trait, but any type of inferences). Evidence for trait inferences within this paradigm suggests that participants are not able to control their occurrence. However, it is also observed that participants have better performance with the progress of the experiment. For example, Uleman et al., (1996) reported evidence of spontaneous trait inferences in the first, but not in the second, half of the experiment, suggesting the possibility that that participants learned to have some control over the process throughout the experiment. In other studies, however, no performance increase was observed along trials (McKoon & Ratcliff, 1986; Wigboldus et al., 2003).

Instead of indirectly activating inhibitory processes, another way of studying the controllability issue is by explicitly instructing participants to avoid trait inferences. The only study that used explicit suppression instructions were conducted by Uleman and Blader (2001; see Uleman, Blader, & Todorov, 2005). In this study, the savings in relearning paradigm (Skowronski et al., 1998) was adapted in order to apply the PDP procedure – *Process Dissociation Procedure* (Jacoby, 1991; Jacoby, Toth, & Yonelinas, 1993; Jacoby, Yonelinas, & Jennings, 1997). Participants were presented with behavior-photo pairs under memory instructions. After different delays (no delay, 20-minute delay, and 2-day delay) the previous photos were presented, and participants were asked to provide their impressions of the actors using trait rating scales. In an inclusion condition, participants were asked to remember the behaviors, and to use them in their impressions. In the exclusion condition, participants were told that the behaviors were randomly paired with the photos, so they should remember the behaviors and not be influenced by them. In the inclusion condition, a higher probability of using the trait implied by the behavior (compared with a baseline condition in which trait rating of the

photos were provided without previous exposure to any type of information) is a function of controlled and automatic processes. In contrast, in the exclusion condition, a higher use of the implied trait will only occur when participants fail to control its influence.

Results showed that participants' trait ratings were influenced by both conscious and automatic processes in the no-delay and in the 20-minute delay conditions. With a long 2-day delay, however, the influence of controlled processes disappeared, while the contribution of automatic processes remained significant. The influence of both controlled and automatic processes diminished under the longer 2-day delay.

The controllability issue is a question open to empirical scrutiny. The recognition probe paradigm might indirectly lead to an intention to control trait inferences, and existing data suggest that some control over the process can be learned. An attempt to directly address this question was made by Uleman and Blader (2001). In this study, what was being tested, however, was whether the *influence* of a previous spontaneous trait inference can be controlled, and not if the spontaneous trait inference process itself is controllable. Results from this study suggested that participants have some capacity to control the influence of previously presented behaviors in target evaluation, under short delays. The control capacity is not total, since it doesn't eliminate the influence of automatic processes. Under longer delays, perceivers seem no longer able to retrieve and control the influence of previous information, with automatic processes still influencing responses.

It would also be interesting to explore whether attempts to control trait inferences would result in ironic hyperaccessibility effects (Macrae, et al., 1994; Wegner, 1994; Wegner & Erber 1992).

Evidence about the unintentionally, awareness, efficiency, and controllability of the STI process suggests that the trait inference process presents a considerably degree of automaticity. However, STIs seem not to be an unconditionally automatic process. There are several reasons for this. First, the spontaneous trait inference process is influenced by the type of processing goal. Low level processing goals reduce, and may even eliminate, the occurrence of spontaneous trait inferences (Uleman & Moskowitz, 1994). In order to characterize the process as automatic, the mere attention to the

behavior should result in similar levels of spontaneous trait inferences, independently of the underlying processing goal. Second, there are studies in which participants report some awareness of the trait inference process (Moskowitz, 1993a; Uleman et al., 1986) and some studies report positive correlations between reported awareness and trait-cued efficacy (Uleman & Moskowitz, 1994; Uleman et al., 1986; Winter et al., 1985). Moreover, a general problem with awareness data is that it is generally based on self-reports (Nisbett & Wilson, 1977). Third, high levels of cognitive load eliminate spontaneous trait inferences (Uleman et al., 1992). Fourth, there is evidence that participants can have some control over the influence of previous spontaneous trait inferences (Uleman & Blader, 2001). Finally, several studies have shown that individuals differ in the extent to which they engage in spontaneous trait inferences (Caldwell & Newman, 2005; Moskowitz, 1993a; Newman, 1993; Uleman et al., 1986; Zelli, Cervone, & Huesmann, 1996). Thus, contrary to Winter and Uleman's (1984) intuition, spontaneous trait inferences seem not to take place every time a behavior is observed. These results fit well with data from other domains, which shows that most cognitive processes cannot be described as totally automatic (see Bargh, 1989; 1992). If processes systematically fail to show a restricted automaticity, we can question whether testing the automaticity of a process is the most adequate research approach. A look into the ambiguities that have surrounded the definition of automaticity (for reviews see Moors & De Houwer, 2006; Wegner & Bargh, 1998) will help us to better understand how the STI process should be approached in these terms.

2.3.5. Conditional Automaticity View

Knowing whether human beings are mainly guided by automatic processes or whether they have control over their actions and mental processes, represents a great intellectual debate within philosophy (e.g., Dennet, 1984; 1991) and psychology (e.g., James, 1890; Bargh & Ferguson, 2000; Wegner & Bargh, 1998). Classic perspectives (LaBerge & Samuels, 1974; Logan, 1980; Neely, 1977; Posner & Snyder, 1975; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977) classify cognitive processing as either automatic or controlled. Automatic processes are viewed as possessing a number of different features. According to these classic views, automatic processes are initiated by the mere presence of the stimulus, and start without people's intentions;

once initiated they are difficult to alter or suppress (i.e., uncontrollable); they operate rapidly and without people's awareness; and they are not dependent on attention or available cognitive resources (because of that, they can operate in parallel with other automatic or controlled processes).

Controlled processes are defined by the opposite features. They are dependent on people's intentions, are controllable, conscious, and non-efficient. Because controlled processes are dependent on attention, they operate in a serial manner. Some authors (e.g., Bargh, 1992, 1994; Logan, 1985) argued that such dual processing distinctions suggest a mutually exclusive and unitary view of automaticity⁴: *Mutually exclusive*, because, despite interactions between the two types of processes were admitted (Logan, 1980; Shiffrin & Schneider, 1977), processes themselves were characterized as automatic or as controlled; and *unitary* because, in order to characterize a process as automatic, all its components should co-occur. This would distinguish automaticity from its component properties (Bargh, 1989, Fiske, 1989).

The dichotomy between controlled and automatic processes was crucial to the development of the socio-cognitive literature, being at the core of many prominent dual processing models that have great relevance in the explanation of diverse social phenomena (Chaiken & Trope, 1999). This distinction has had probably an incomparable impact on social psychology research, with numerous studies trying to explore whether social perception, judgment, and behavior operate in an automatic way (e.g., Bargh & Chartrand, 1999; Greenwald & Banaji, 1995; see also Bargh & Ferguson, 2000; Wegner & Bargh, 1998).

However, it was soon pointed out (e.g., Bargh, 1989, 1992, 1994; Logan, 1985; Logan & Cowan, 1984; Paap & Ogden, 1981; Zbrodoff & Logan, 1986) that mental processes should not be analyzed dichotomously as automatic or controlled. This unitary view of automaticity was criticized both on theoretical and empirical grounds. From a theoretical point of view, Zbrodoff and Logan (1986) argued that there are no bases to defend a unitary view of the automaticity concept. In fact, while it is very easy to

⁴ It is not clear, however, that attentional researchers generally endorsed a unitary view of automaticity. As Moors and De Houwer (2006), mention "a careful reading of the early writings by proponents of the capacity view reveals that its interpretation as strict all-or-none is an overstatement" (p. 299). In fact, Shiffrin and Schneider (1977), for example, admitted combinations between control and automatic features, by saying that not all control processes are available to conscious perception, and that some automatic processes require control in order to be initiated.

conceive of processes that combine properties of automatic and controlled processes, it is hard to conceptualize an example of a “pure” automatic process. In addition, seeing automatic processes as unitary could be misleading in research terms. Researchers could think that it was enough to study one of the features in order to characterize a process as automatic, in what Bargh (1992) called an “automaticity by-default” approach. Moreover, if a process satisfies one of the automaticity criteria, researchers could infer that the other features were also present, since automatic components were said to co-occur. Finally, if processes could only be characterized as automatic or controlled, how would processes that satisfy some of the automaticity criteria, but not others, be defined? Both the unitary and the mutually exclusive properties underlying the automatic/controlled dichotomy seemed not to capture the full complexity of reality.

Experimentally, evidence consistently confirmed the idea that cognitive processes tend to combine characteristics typical of controlled processes with characteristics of automatic processes. For example, despite the fact that stereotype activation was initially seen as an inevitable process (Devine, 1989; Purdue & Gurtman, 1990), it was shown that it is dependent on the level of individuals egalitarian beliefs (Fazio, Jackson, Dunton, & Williams, 1995; Lepore & Brown, 1997; Moskowitz, Salomon e Taylor, 2000), availability of cognitive resources (Gilbert & Hixon, 1991; Spencer, Fein, Wolfe, Fong, & Dunn, 1998; but see Bargh, 1999), and that it varies with previous expectancies (Blair & Banaji, 1996, but see Bargh, 1999), and processing goals (Macrae, Bodenhausen, Milne, & Calvini, 1999; Macrae et al., 1997).

Even processes that were viewed as classic examples of pure automatic processes were shown to present characteristics of controlled processes. Very well learned automatized processes, such as walking, typing, or reading, need a certain level of conscious intention to be initiated, and are highly controllable (Logan, 1985; Logan & Cowan, 1984; Tzelgov, Henik, & Leiser, 1990; Zbrodoff & Logan, 1986). Encoding of alphabetic letters, viewed as a prototypical automatic process, despite being unintentional requires resources and interferes with secondary tasks (Ogden, Martin, & Paap, 1980; Paap & Odgen, 1981). Another example is the Stroop effect (Stroop, 1935; for a review see MacLeod, 1991), usually seen as highly uncontrollable and unintentional, was shown to require some level of attention (e.g., Francolini & Egeth, 1980; Kahneman & Chajczyk, 1983; Kahneman & Henik, 1981; Kahneman &

Treisman, 1984), and, recently, evidence for the controllability of the effect was also presented (Tzelgov, Henik, & Berger, 1992; Tzelgov et al., Leiser, 1990). These examples clearly support the idea that automaticity is not an all or none phenomenon (for other examples see Bargh, 1989, 1992).

Considering these arguments, Bargh proposed an alternative view of the automaticity concept (Bargh, 1989; 1992; 1994). Two central ideas can be extracted from Bargh's proposal. First, according to Bargh, automaticity is a multidimensional concept. The four "horsemen of automaticity", as Bargh (1994) described them, are: efficiency, intentionality, controllability, and awareness. Contrary to previous conceptions, these four components of automaticity were viewed as being highly independent from each other, and were generally described as functioning in a continuous way (see also Moors & De Houwer, 2006). The implication of this view is that the different automaticity features can be independently implied in the execution of a mental process. Thus, different combinations of the four dimensions can be found across different mental processes. The second central idea in Bargh's proposal is that all automaticity is conditional. The occurrence of any process, even of processes previously taken as examples of automaticity, is always dependent on certain enabling conditions. These enabling conditions can vary from process to process. Some processes may require stimulus awareness; others may need some level of attention; and others may be dependent on specific processing goals. Naturally, as Bargh (1989, 1992) stated, the fewer the requirements, the more recurrent the process will be.

Based on his view of automaticity, Bargh (1989, 1994) recommended that researchers study each component of automaticity separately, in order to establish the minimum conditions for the occurrence of a process. The focus should be, not on testing the existence of pure automatic processes, but on exploring the conditions of their occurrence. This approach has a greater explanatory value. For instance, different processes previously grouped under the same "automaticity umbrella", vary in the necessary conditions for their occurrence. Thus, used without distinction, the label "automatic," when applied to these different processes, can cast a shadow over the (more relevant) underlying mechanisms.

According to Bargh (1989, 1992), by knowing the minimum conditions for the occurrence of a cognitive process, one could then group the different processes in terms

of their automaticity requirements. Bargh (1989, 1992) tried to formulate such a classification. He analyzed diverse social processes that were studied in terms of automaticity, and proposed that most of them could be classified in three automaticity categories: pre-conscious automaticity (processes that not require conscious awareness of the stimulus), post-conscious automaticity (processes that require attentional awareness of the stimuli), and goal-dependent automaticity (processes that require both conscious awareness of the stimulus and the existence of a specific processing goal). Critical to our discussion, Bargh (1989; 1992) included the STI process in the *goal-dependent automaticity* category. Based on the Uleman and Moskowitz (1994) findings, the STI process was described as being conditional upon a meaningful processing of behaviors. In fact, we presently know that STIs are not only conditional upon a meaningful processing goal, but are also dependent on other factors. As previously mentioned, research has been demonstrating that STIs occur under certain conditions, but are undermined under other conditions (Uleman et al., 1992; Wigboldus et al., 2003).

Although we agree that STI can possibly be described (from an automaticity perspective) as a conditional process, we would propose that more crucial than defining the STI process with a label (as being *goal-dependent automatic*, for instance; Bargh, 1989, 1992) is developing theories that describe the process. Such theories would naturally incorporate conditions for the occurrence of the process, as well as inform us about its pervasiveness. The explanatory inefficacy of automaticity labels (Bargh, 1989, 1992) is reflected in the more recent work of Bargh himself (e.g., Bargh, 1997; 2005; Bargh & Williams, 2006), who no longer characterizes processes according to his automaticity typology, preferring to use general definitions of automaticity (see Bargh, 1997; Bargh & Williams, 2006). We should mention that, in this way, Bargh might contribute to a neglect of his own concerns about inaccurate definitions of automaticity. In fact, Bargh continuously states that “social psychological phenomena are essentially automatic” (Bargh, 1997, p.3), and argues in favor of the automaticity of daily life (Bargh & Williams, 2006) and of complex social behavior (Bargh, 2005; Bargh & Ferguson, 2000). However, no clear definition of automaticity is provided. Moreover, conditions for the occurrence of the processes, as well as the conditional nature of automaticity, seem no longer to be major concerns.

We suggest that the best research approach in order to understand the conditional nature of the STI process is by developing theories that are able to explain its underlying processing mechanisms, and not by trying to diagnose the automaticity of the process. This is especially true when we consider the ambiguities that still surround the definition of automaticity. The question about how to define automaticity is an open contemporary debate in the field (Bargh, 2005; 2006; Moors & De Houwer, 2006). For example, there is no consensus about what the components of automaticity are (e.g., should the four criteria proposed by Bargh be generally accepted?). We also don't know, at an experimental level, whether the different components are independent, or overlap with each other. In addition, although Bargh's (1989, 1994) definitions of the features of automaticity are generally applied, it is not consensual how those features should be defined. For example, although efficiency has been viewed as implying a single resource pool with limited capacity by some researchers (Kahneman, 1973), others have proposed a multiple resources view, in which different types of resources do not necessarily interfere with each other (Navon & Gopher, 1979; Wickens, 1980; for a review see Wickens, 2008). In the same vein, Moors and De Houwer (2006) proposed a distinction between different types of consciousness and controllability. Considering these points, and having in mind the fact that automaticity features are seen as varying in a gradual way (Bargh, 1989; 1994; Moors & De Houwer, 2006), the automatic/controlled dichotomy loses its meaning. It seems impossible to delineate a line that can distinguish controlled from automatic processes.

In sum, we presented theoretical and empirical arguments that illustrate how the distinction between automatic and controlled processes, although useful to stimulate research, fails to capture the nature of most cognitive processes. If spontaneous trait inferences are, just like other cognitive processes, conditional, it would seem more important to understand what are the mechanisms responsible for the process and the factors that regulate its operation. For this proposal, existing theoretical models of inferential processes underlying general comprehension might be of important value (e.g., McKoon & Ratcliff, 1992, 1995).

2.4. Spontaneous Trait Transference Effect

The tendency to attribute personality traits to individuals who are merely describing someone else's behavior has been known as the spontaneous trait transference (STT) effect. The trait implied by the behavior that is being described is *transferred* to the communicator himself. For instance, if John describes a dishonest behavior performed by an acquaintance, John himself will be perceived as dishonest. This effect was first demonstrated by Carlston et al. (1995, Experiment 4). In this study, the savings in relearning paradigm was used, but participants were told that the photos paired with the behaviors were not of the actors of the behaviors, but of communicators who were describing the behaviors. Even in that case, however, a facilitation in learning of the photos-traits pairs was observed, indicating that the traits implied by the behaviors had become associated with the communicators.

In subsequent studies, the spontaneous trait transference process has been shown to be a persistent and robust phenomenon. Spontaneous trait transference was demonstrated to occur: (a) under cognitive load conditions (Crawford et al., 2007); (b) when participants are informed about the nature of the effect, and are told to avoid it (Carlston & Skowronski, 2005, Experiment 3; Skowronski et al. 1998); (c) after a 2-day delay between the exposure and test phase (Skowronski et al. 1998) and; (d) with impression formation instructions (see also Skowronski et al. 1998, Experiment 4). This last result seems to be, however, in contradiction with results from Carlston et al. (1995, Experiment 4), who found no evidence of transference effects in an impression formation condition.

Skowronski et al. (1998, Experiment 2) also explored whether the transferred trait has real consequences to the impressions formed about the communicators. In order to test this idea, the learning and the recall tasks were replaced by a trait rating task. In this study, both self-description and acquaintance description trials were included in order to directly compare spontaneous trait transference and spontaneous trait inference effects. Results revealed that communicators were rated higher on the trait implied by the described behavior, in comparison with a condition in which they didn't describe the behavior (i.e., STT effect). Second, the magnitude of the transference effect was weaker than the magnitude of the trait inference effect. That is, trait ratings were higher when the trait was implied by a behavior performed by the actor (self-description condition)

than when it was described by another actor (acquaintance description condition) (see also Skowronski et al., 1998, Experiment 4). Third, the transference effect was only observed with the implied trait and did not generalize to other traits (see also Carlston & Skowronski, 2005; Skowronski et al., 1998, Experiments 3 and 4). Finally, it was shown that the transference effect is independent of participant's mistakenly remembering the behavior as having been performed by the communicator (see also Skowronski et al., 1998, Experiments and 4). In a last experiment (Skowronski et al., 1998, Experiment 4), it was also shown that the transference effect occurs under more natural settings, for example, when videos with the communicators describing the behaviors are used (Skowronski et al., 1998, Experiment 4).

In a subsequent study, Mae et al. (1999), using both a relearning task (Experiments 1 and 2), and a rating scale task (Experiment 3) showed STT effects even when communicators are well known celebrities. Moreover, the transference effect was independent of the traits being congruent, neutral, or incongruent with previous knowledge about the celebrities. These results show that the spontaneous trait transference effect is not limited to actors about whom participants have no previous knowledge. In the same study, it was also demonstrated that the variable need for cognition has no influence in the magnitude of the trait transference effect. Thus, the tendency to engage in more systematic processing doesn't prevent the biased transference of a trait to its informant.

In sum, spontaneous trait transference seems to be a highly persistent phenomenon. It emerges under a wide variety of circumstances, and is resistant to participant's attempts to control its occurrence.

2.4.1. Theoretical Accounts of STT

Skowronski and colleagues (Mae et al., 1999; Skowronski et al., 1998) proposed a model of spontaneous trait transference, in which the effect is generally explained by associative mechanisms. According to their model, the STT effect results from the intervention of three processes. First, perceivers activate a personality trait during behavior encoding (*trait activation*); then, the trait becomes associated with the communicator (*trait association*); and finally, the formed association has the power to implicitly influence subsequent impressions of the communicators (*trait influence*).

However, alternative explanations were also presented. Some authors argue (Bassili & Smith, 2002; Skowronski et al., 1998) that it is entirely plausible that participants attribute traits implied by behaviors to communicators because they assume that communicators and acquaintances are familiar with each other, and for that reason, should share similar personal characteristics. Alternatively, participants may reason that, if a communicator had chosen to describe a specific characteristic of an actor, it must be because he approves of it, or considers the attribute to be important for him. Subsequent data made clear that these alternative explanations cannot totally explain the trait transference effect occurrence.

Skowronski et al. (1998), for example, showed spontaneous trait transference effects, even when participants are informed that photos and behaviors were randomly paired. Under these conditions, there is no implicit personal theory that can justify the attribution of traits implied by the behaviors to the persons in the photos. The transference effect was shown also to be independent of participants' perception of approval or disapproval of the behavior by the communicator, and of the perceived authenticity of the video (Skowronski et al., 1998, Experiment 4). Additionally, transference effects were also observed with inanimate objects, in relation to which we have no personal theories (Bassili & Smith, 2002). Thus, alternative explanations of spontaneous trait transference based on implicit assumptions that participants may have about the personality of the communicators are not adequate to explain the effect.

Another alternative interpretation for spontaneous trait transference effects is that participants erroneously attribute traits to communicators because they confuse self-describing and acquaintance-describing trials (Carlston & Skowronski, 2005). As a result, they can mistakenly assume that communicators perform a behavior that they were merely describing. This error may be the result of participants being inattentive to the self or other nature of each trial during encoding (i.e., inaccurate encoding) or because participants erroneously recall other-descriptions as having been self-descriptions during retrieval (inaccurate retrieval).

In order to make clear that the descriptions were not self-descriptions, in all existing studies, the gender of the actor of the behaviors was opposite to the gender of the communicator. Additionally, several results contradict a confusion processing hypothesis. Carlston and Skowronski, for example, doubled the encoding time to 20

seconds in one of the studies (Carlston & Skowronski 2005, Experiment 1), and incorporated a self pace presentation in another (Carlston & Skowronski 2005, Experiment 2). If transference effects are dependent of an inaccurate encoding, the effect should be reduced under these conditions. Encoding time manipulations, however, had no effect in the spontaneous trait transference magnitude.

The transference effect is also observed when descriptions are manipulated between-participants (Carlston & Skowronski, 2005, Experiment 2; Crawford et al., 2007b). Since participants only received self or other-descriptions, there is little room for confusions in this case. Finally, in order to rule out any hypothesis that participants may forget the nature of the descriptions, Carlston and Skowronski (2005, Experiment 2) solicited the trait rating immediately after each sentence. Again, a transference effect was observed.

Transference effects seem not to be explained by an inaccurate encoding, or inaccurate retrieval of the identity of the actor. Instead, results seem to be better explained by an associative mechanism (Mae et al., 1999; Skowronski et al., 1998). Consistent with this view, no savings effect evidence to the narrator is obtained with explicit impression formation instructions (Carlston et al., 1995, experiment 4).

2.4.2. Limits on STTs

In most of the studies that report spontaneous trait transference effects, the behavior is presented with a photo of the communicator. This procedure may facilitate the creation of an association between them. We can question what would happen if a photo of the actor himself was incorporated in the encoding context. In this case, the photo of the actor may capture the participant's attention, preventing the establishment of an associative link between the communicator and the behavior.

Crawford and his colleagues (Crawford et al., 2007; Crawford et al., 2008) provided support for this hypothesis. They presented two photos on each trial. In some conditions, one of the photos was of the *actor* of the behavior and the other photo was of someone for whom the actor was describing his own behavior (i.e., a *bystander*). In others conditions, one of the photos was of an *informant* (i.e., a person that is describing a behavior performed by someone else) and the other photo depicted the *target* of the behavior described by the informant. A savings effect was observed only for those who

performed the behaviors (actors and targets), and not for bystanders and informants. This pattern was shown both in a relearning task (Crawford et al, 2007, Experiment 1) and in a trait rating task (Crawford et al, 2007, Experiment 2).

These results are in line with findings from Todorov and Uleman (2004) using the false recognition paradigm (but see Bassili & Smith, 2002). Todorov and Uleman (2004) presented pairs of faces with each behavior. One face was of the actor and the other of a bystander. Results showed a larger false recognition effect when the implied traits were paired with the actors' faces than when they were paired with bystander faces (Todorov & Uleman, 2004, Experiments 1 and 2). The same pattern was replicated when the recognition test was presented one week after the study phase, despite being weaker⁵ (Todorov & Uleman, 2004, Experiment 3).

Thus, the presence of the actors of the behaviors during the encoding context seems to eliminate, or reduce, the transference effect observed in previous studies. How can the elimination of the effect under these conditions be explained? Crawford et al. (2008) tested an attentional hypothesis, according to which the presence of the actor captures participant's attention, preventing the association between traits and communicators.

In the first study, a probe attentional task was applied. After each trial, a probe appeared in the screen, either in the position occupied by the actor, or in the position of the other photo (the bystander or the informant, depending on the condition). If the attention is captured by the actors of the behaviors, participants should take less time to respond when the probe appears in the actor or target position, than when it appears in the informant or bystander position. In the second study, a trait rating task was included in place of the relearning task, and an attentional eye tracking measure that allowed the recording of the continuous movement of participant's eyes was used. In both studies, a savings effect was observed for actors and targets of the behaviors, but not for informants or bystanders. This pattern confirms the idea that presenting the actor of the behavior eliminates the transference effect. However, the elimination was not explained by differences in visual attention. In the first study, using the probe attentional task, no differences in response times were observed. In the second study, participants made

⁵ In this case performance was at chance level, which may indicate the use of guessing as a response strategy.

more saccades and spent more time looking to actors than to bystanders, in the self-informant condition. In the other-informant condition there was no differences in the amount of time and number of saccades between informants and targets. However, the attentional pattern observed does not mediate trait ratings.

Todorov and Uleman (2004) also excluded an attentional explanation for their results. In this case, they tried to equalize the amount of time participants look to each one of the photos. Two sentences were presented with two different actors, with each one of the sentences describing a behavior performed by one of the actors. Higher levels of false recognitions were observed when the implying traits were paired with the corresponding actors than when traits were paired with the control faces (i.e., the other actor). Todorov and Uleman (2004) concluded that the effect cannot be totally explained by the amount of attention, although they recognized that attention has some level of influence since the effect was reduced under these conditions.

Based on these findings, both Todorov and Uleman (2004) and Crawford and collaborators (Crawford et al., 2008) concluded that the elimination of the transference effect caused by the presence of the actor of the behavior is due, not to a greater visual attention directed to actors, but to the occurrence of a greater attributional processing. According to this explanation, the presence of actors stimulates a deeper attributional analysis, in comparison with a situation in which actors are not present. This attributional processing prevents associations between traits and other individuals that happen to be present during the encoding context. However, exactly how and why attributional processes would disrupt the associative link between communicators and behaviors was not clarified.

It must be also noticed that measures of visual attention give information about the orientation of participant's eyes, but is not informative in relation to what participants are thinking, or about the contents of their working memory (Crawford et al., 2008). It is possible that participants think more about the actors of the behaviors, independently of their visual orientation. This would imply that the pattern can be explained by stronger associative links between actors and behaviors, without necessarily postulating that a qualitatively different attributional processing is occurring for actors.

Independently of the theoretical explanation, results prove that spontaneous trait transference effects are not inevitable. Consequently, associative processes underlying transference effects cannot be characterized as totally automatic and passive. The presence of the actor impedes the creation of an association between the communicator and the behavior, otherwise established. Why exactly the association is not formed under these conditions is not clear. But the most plausible interpretation seems to be that the creation of an associative link between other individuals present in the context and the behaviors requires some minimal level of processing (probably not captured by visual attentional measures).

Only one other study reported the elimination of the STT effect. Carlston and Skowronski (2005, Experiment 3) asked some participants to recall the other or self nature of the descriptions made by informants before they made trait ratings. For these participants, the STT effect was not observed. Since warnings have no effect on STTs, this result was explained by a specific retrieval mechanism. According with the authors, participants opted to be more cautious in their ratings because they lack memory for informants describing their behaviors. The more conservative responses impeded the influence of previously formed associations. Hence, transference effects can be eliminated either by presenting the actor of the behavior during encoding or by making participants question the bases of their judgments.

2.5. Other Spontaneous Inferences

Despite the fact that research has been mainly focused on spontaneous trait inferences, other types of meaningful social inferences probably can also occur spontaneously. There are a number of meaningful inferences that can be extracted, including gist, situational, predictive, or emotional inferences. All of them are highly informative, and are essential to correctly capture the regularities of the social world.

The same arguments that are used to sustain the mechanization of the trait inference process can be applied the occurrence of these other inferences (Uleman et al., 1996). For example, we may argue that inferring an individual's personality traits is fundamental because it allows us to predict how others will act. However, if we don't develop a corresponding capacity to take into account the situation, the individual's goals, or the emotions they express, our ability to infer the dispositional rules of the

social word will be compromised. In fact, if trait inferences were the only inferences acquiring some level of automation, probably they would be highly unsuitable for the individual's adaptation. Perceivers would be constantly inferring dispositional traits, independently of the situational constraints, individual history, or the specific goals individuals pursue. We can imagine this as more harmful than beneficial for individuals, more chaotic than stable. In line with this reasoning, a number of studies have shown that other social inferences may also occur spontaneously.

2.5.1. Gist Inferences

Behavior gists are words that capture the general meaning of the behavior. Essentially, they are words that summarize the action performed by the actor of the behavior. For example, the gist word of the behavior "steps on his girlfriend's feet" is "dancing", while the inferable trait is "clumsy" (for other examples, see Winter et al., 1985). Pre-tests of the gist words for behaviors explicitly asked participants to report the word that best characterizes the behavior activity, and results show that reported words are different from trait words. Following the operationalization used in previous studies, gist inferences describe what the individual is doing, while trait inferences describe the meaning of the individual action in personality terms.

The first study showing spontaneous gist inferences was conducted by Winter et al. (1985). They showed that trait and gist cues were equally effective in retrieving behaviors. Later studies also reported the occurrence of spontaneous gist inferences (Lupfer et al., 1990; Uleman & Moskowitz, 1994; Uleman et al., 1992), with gist cues in these cases being even more effective retrieval cues than trait cues.

Other characteristics of the gist inference process were reported. Uleman et al. (1992) demonstrated that gist cues became less effective with increasing levels of cognitive load, although showing less interference than trait cues. Thus, gist inferences seem not to be a totally cognitively efficient process. Uleman and Moskowitz (1994) reported, first, that gist inferences tended to be either positively correlated or uncorrelated with trait inferences, which means that they do not compete with each other. Second, gist cues were generally more effective than trait cues in retrieving actors. The gist of the behavior seems to become more strongly associated with the actor than do the personality traits. Third, opposite effects of processing level were

observed for gist and trait cues. While trait inferences increase with more semantic processing goals, gist inferences show the reverse pattern, decreasing as the processing goal becomes more semantic. In fact, evidence for gist inference is not obtained under semantic goal conditions (i.e. analyzing the gender of each word). While trait inferences tended to increase the more semantic is the processing goal, gist inferences tend to decrease as the processing goal involves greater semantic analysis. In fact, evidence for gist inference is not obtained under semantic goal conditions (i.e. analyzing the gender of each word). This result is difficult to interpret and is inconsistent with other results reported in the same article showing that instructions to ignore the meaning of behaviors reduced the level of spontaneous gist inferences. Contrary to the previous data, these results suggest that gist inferences are somehow dependent on semantic processing.

Existing studies that explored the occurrence of gist inferences have generally obtained evidence for their spontaneity (Lupfer et al., 1990; Uleman & Moskowitz, 1994; Uleman et al., 1992; Winter et al., 1985). These findings fit well both with the three stage model (e.g., Gilbert et al., 1988) and with the identification-inference model (Trope, 1986) according to which the initial, more automatic, stage of behavior encoding is its identification or categorization. Because inferring the gist of the behavior may be seen as an intrinsic part of its identification it may be not surprising that evidence for the spontaneity of the process is generally obtained.

2.5.2. Situational Inferences

One of the first studies exploring the existence of spontaneous situational inferences was conducted by Lupfer et al. (1990). In this study, it was found that situational cues were not effective in retrieving behaviors. However, this is not surprising since behaviors were taken from Winter et al. (1995) and pre-tested to elicit dispositional, rather than situational inferences. More remarkable was the effect of adding situational or trait background scenarios. Dispositional background information stressed the importance of characteristics of the actor to explain the behavior, while situational background information highlight the importance of external constrains to explain the behavior. Results showed that background information that encourages situational inferences significantly increases the magnitude of spontaneous situational inferences, but has no impact on spontaneous trait inferences. In contrast, dispositional

background information boosts trait inferences, with no impact on situational inferences. This suggests that spontaneous situational and trait inferences work independently, and can both co-occur.

In this study, spontaneous situational inferences were shown to be affected by cognitive load, but only in the situational background condition. In this case, situational inferences were more likely to occur with unconstrained cognitive resources than under cognitive load conditions. In general, participants reported almost no awareness of making situational inferences and no correlation was observed between situational trait awareness and situational-cue efficacy.

Since the behaviors in Lupfer et al.'s (1990) study were developed to elicit trait inferences, direct comparisons between spontaneous trait and situational inferences are not possible. Sentences specifically pre-tested to imply situational inferences were applied in a study from Lupfer, Clark, Church, DePaola, and McDonald (1995; cited by Uleman et al., 1996). In one of the experiments (Lupfer et al., 1995, Experiment 2), a recognition probe task was incorporated between digits presentation, in an allegedly digit memory task. The probability of participants incorrectly indicating that the probe was part of the previous sentence was higher when trait probes followed trait implying sentences, and when situational probes followed situational implying sentences, compared with control conditions. In another experiment (Lupfer et al., 1995, Experiment 3), trait implying or situation-implying sentences were presented as distractors in a lexical decision task. Participants took less time and were more accurate in identifying the word correspondent to the inferable situation after reading situation implying sentences, and were better in identifying trait words after reading trait implying sentences.

In the previous mentioned studies, spontaneous trait inferences and spontaneous situational inferences were shown to occur for different behaviors. Ham and Vonk (2003) demonstrated that the two types of inferences can co-occur for the same behaviors. After observing critical sentences (e.g., John gets an A on the test) participants have more difficulty rejecting implied traits (e.g., smart), and also to reject implied properties of the situation (e.g., easy), than unrelated probe words (Ham & Vonk, 2003, Experiment 1). Evidence for the simultaneous occurrence of spontaneous trait inferences and spontaneous situational inferences was also obtained using an

adaptation of the savings in relearning paradigm (Ham & Vonk, 2003, Experiment 2). The observed savings effect was shown to be independent of the type of processing goal. Specifically, instructions to form an impression about the actor, about the situation, or about the event had no impact on the trait and situational savings effects. Thus, spontaneous situational inferences occur regardless of participants having a trait dispositional goal, and spontaneous trait inferences occur regardless of participant's being focused on the situational aspects of the event.

In sum, existing evidence shows that people spontaneously infer not only traits, but also situational aspects of the events they observe. This data contradicts Gilbert three stage model according to which people first make an automatic dispositional inference, and only then correct it by considering the situational constraints of the behavior (e.g., Gilbert et al., 1988).

Also, according to the reported studies on spontaneous situational inferences, spontaneous situational and trait inferences are not mutually exclusive (i.e., the occurrence of one type of inference does not eliminate the occurrence of the other). Thus, despite the intuitive notion that people use traits and situations as alternative explanations of behavior, empirical data do not support such notion. This pattern is also consistent with findings from the attribution literature (Krull, 2001; Miller, Smith, & Uleman, 1981; Solomon, 1978).

2.5.3. Goal Inferences

Inferring other's goals and intentions is a fundamental ability of a social being. If someone approaches us, it is crucial that we are able to quickly determine his intentions. Is it his intention to harm us? Is it his intention to be kind to us? (Fiske et al., 2007; Heider, 1958; Reeder, 2009). Because others' actions are driven by their internal goals, inferring those goals allow us to understand and predict others' behavior. As Moskowitz and Grant state (2009, p.6): "goals provide for us meaning about why other people act the way they do". It is not surprising that several theoretical approaches view behavior has mainly explained by reference to intentions and motives (Malle, 1999; Read, 1987; Reeder, 2009; Schank & Abelson, 1977).

Goals have long been recognized as essential constructs in person perception (Lewin, 1951; Heider & Simmel, 1944; Heider, 1958). Heider, for example, considered

intentionality to be the central factor in personal causality (Heider, 1958). According to Heider, the first stage in person perception consists in exploring whether the actor had the intentional goal to perform the observed action (personal causality) or whether the action was made accidentally, or due to environmental forces (impersonal causality). Only after defining the intentional nature of the action are deeper levels of causal analysis explored. Despite Heider's considering goal inferences as a basic step of person perception, subsequent research mainly focused on traits. As Malle (2008, p. 165) notice: "Although textbook entries... consistently laud Heider as the father of attribution work, apparently the father's words were not heeded". In a similar vein, intentionality is viewed as a pre-condition to correspondent inferences according with the Correspondence Theory (e.g., Jones & Davis, 1965). If the behavior is performed without intention, there is no basis to assume that it reflects something about the personality of the actor. Recently, the multiple inference model developed by Reeder (Reeder et al., 2004; Reeder, 2009) also puts emphasis on motives. The model assumes that the perceiver typically make multiple inferences about behavior, including inferences of motives and traits. Importantly, although the model assumes multiple inferences are drawn, intentional behavior is primary perceived in terms of underlying motives, according to the model.

According with the vision that goal inferences have a predominant role in person perception, several studies have been consistently supported the notion that people tend to conceptualize actions in terms of goals. For example, in a classic study, Heider and Simmel (1944) showed that participants extract goal intentions even from the observation of geometric figures movements. Authors conclude that "as soon as we ascribe a certain movement to a figural unit and consider this unit as an animated being, perception of motive or need is involved (Heider & Simmel, 1944, p.257).

Developmental studies confirm the early emergence of the ability to infer an actor's goals. It was shown that 9-month-old babies already show the capacity to interpret behaviors in goal terms (Csibra, Gergely, Biro, Koos, & Brockbank, 1999; Gergely, Nadasdy, Csibra, & Biro, 1995; Premack & Premack 1997; Woodward 1998). Similar results were obtained with chimpanzees (Uller, 2004; Uller & Nichols, 2000), suggesting that the mechanism to identify goals occurred early in evolutionary terms.

Direct evidence that people spontaneously infer goals from behaviors was recently provided by Hassin et al. (2005; see also Aarts, Dijksterhuis, & Dik, 2008; Van der Cruyssen et al., in press). Using a cued-recall paradigm, Hassin and collaborators (Hassin et al., 2005) demonstrated that goal cues are more effective in retrieving goal sentences, comparing with control sentences (Experiment 1). The effect occurred even when it was clear in the sentence that the goal was impossible to attain (Experiment 2). The spontaneity of the goal inference process was also demonstrated in a recognition probe paradigm (Experiment 3) and in a lexical decision task (Experiment 4). In all studies, participants reported no intention or awareness of making goal inferences. Additionally, Aarts et al. (2004) showed that goals inferred unconsciously influence subsequent behavior, making participants more likely to act in agreement with the goal (i.e., goal contagion).

Additional evidence that goal inferences can occur spontaneously comes from the text comprehension literature. Studies in this domain have demonstrated that readers make online goal inferences as they read action statements (Long & Golding, 1993; Long, Golding, & Graesser, 1992; but see McKoon & Ratcliff, 1992 for a minimalist view of comprehension inferences). These findings will be discussed in detail in the text comprehension section.

In sum, existing studies show that goal inferences occur as automatically as trait inferences (Hassin et al., 2005). In addition, some authors suggest that goal inferences are not only as automatic as trait inferences, but may well precede trait inferences. According with this view, inferences about goals play a central role in trait inferences (Jones & Davis, 1965; Idson & Mischel, 2001; Read & Miller, 1993; Read & Miller, 2005; Reeder et al., 2002; Reeder et al., 2004; Reeder, 2009). Read, Jones, and Miller (1990) provide support to this notion by showing that the extent to which a behavior is related with a goal associated with a trait predicts both the extent to which the behavior is rated as prototypical of the trait, and the probability of making trait inferences from the behavior. Actually, the idea that trait inferences are dependent on inferences about actors' goals is at the center of the recently developed Social Dynamics Model (Read & Miller, 2005) that states that, in order to make sense of others people's actions, perceivers construct narratives based on actor's goals. The same idea is expressed in the multiple inference model of Reeder (Reeder et al., 2004; Reeder, 2009), according to

which inferences concerning the motives that drive the behavior of the actor guide the trait inferences that are made about that actor.

Thus, as some have suggested (Aarts et al., 2008; Read & Miller, 1989; Read, et al., 1990) it might be the case that goal inferences occur in an early categorization stage of behavior processing. This is in agreement with Heider and Simmel's (1944, p.257) statement that "often the attribution to motive adds something and differentiates between different interpretations of actions". Thus, there are both empirical and theoretical reasons that support a recurrent occurrence of spontaneous goal inferences. However, even goal inferences are not unconditional and are probably constrained by coherence requirements (McKoon & Ratcliff, 1992).

2.5.4. Others Spontaneous Inferences

Besides those previously mentioned, other social inferences have been shown to occur in the absence of deliberative intentions. For example, McKoon and Ratcliff (1986) provided evidence that people make predictive inferences online during text comprehension. Uleman et al. (1996, Experiment 3) confirmed these findings using a recognition probe paradigm. They compared trait and predictive inferences, and showed that both inferences occur in similar degrees.

Counterfactual inferential thinking was also shown to come to mind in a spontaneous way (McEleney & Byrne, 2006; Roese & Olson, 1997; Roese et al., 1995; Sanna & Turley, 1996, 2000). Counterfactuals refer to thoughts about how an outcome could have been different (Roese, 1997). It was shown, for example, that people spontaneously report counterfactual thoughts, especially after controllable and negative outcomes (McEleney & Byrne, 2006). However, in this study a thought report measure was applied. This type of measure is not conclusive about the spontaneity question, since participants may report thoughts that would not occur to them spontaneously.

Roese and Olson (1997) made use of a more indirect measure. In this case, participants were given positive or negative feedback about their performance in an anagram task. After the anagram task, participants responded to a series of filler questions in which a counterfactual statement related with the previous task was included (e.g., "my score could have been much different"). Results showed that participants take less time to respond to the counterfactual statement after receiving

negative feedback, which suggests stronger activation of counterfactual thinking. Thus, it seems that counterfactuals can come to mind spontaneously, but the spontaneity is constrained by the controllability and valence of the outcome

There is also research showing that people spontaneously infer emotional states of observed targets (Ford & Milosky, 2008; Gygax, Oakhill, & Garnham, 2003; Gernsbacher, 1994; Gernsbacher & Robertson, 1992; Gernsbacher et al., 1992; Gernsbacher, Hallada, & Robertson, 1998). Gernsbacher et al. (1992) instructed participants to read scenarios that implied emotional states, and showed that participants took significantly less time to read a subsequent sentence that included the implied emotional word than a sentence that included a mismatched emotional word (Gernsbacher et al., 1992, Experiment 1 and 2; see also Gernsbacher & Robertson, 1992). In order to rule out the possibility that emotional activation could be occurring only when subsequent sentences were being read, an obtrusive measure was included in another experiment (Gernsbacher et al., 1992, Experiment 3). In this case, participants were simply asked to pronounce words after reading the stories. Results showed that words correspondent to the implied emotion were read much quicker than words correspondent to the emotion opposite to the one implied by the stories.

The extraction of emotional states during behavior comprehension was shown to be a process immune to cognitive load manipulations (Gernsbacher et al., 1998, Experiments 1-3). Additionally, it was observed that participants take more time to read sentences that included mismatch emotional sentences than match sentences, independently of the implied emotion being explicitly mentioned in the text or being merely implied by the story (Gernsbacher et al., 1998, Experiments 4). This set of results led authors to conclude that readers automatically activate emotional knowledge during text comprehension.

Murphy, Wilde, Ogden, Barnard, and Calder (2009) adapted the paradigm used by Gernsbacher and colleagues (Gernsbacher & Robertson, 1992; Gernsbacher et al., 1992) to study the occurrence of spontaneous moral processing. It was predicted that, if participants make moral inferences as part of behavior comprehension, they should take more time to read moral matching than moral mismatching sentences presented after reading paragraphs with moral outcomes. Results confirmed the authors' hypotheses. In

addition, differences in reading times between match and mismatch conditions were not affected by load manipulations, suggesting the efficiency of the process.

In line with these findings, Ham and Van den Bos (2008) obtained evidence of occurrence of a specific type of moral inferences, justice inferences. The recognition probe paradigm was applied in this study. Results showed that participants have more difficulty correctly rejecting justice related words (e.g., just, fair) as not being presented in previous sentences, after reading justice implying sentences than after reading control sentences. Spontaneous justice inferences were shown to be especially strong for personally relevant and unjust outcomes (Ham & Van des Bos, 2008).

2.5.5. The Multiple Social Inferences View

These studies suggest that spontaneous trait inferences are not the only kind of social inferences that occur spontaneously. Existing studies show that other social inferences can also occur in the absence of explicit intentions, and might function in a similar way to trait inferences. Furthermore, studies that explored the simultaneous occurrence of different kinds of spontaneous social inferences (Ham & Vonk, 2003; Uleman & Moskowitz, 1994; Uleman et al., 1996, Experiment 3) have shown that multiple social inferences can co-occur, without interfering with each other.

The view according to which multiple inferences can co-occur is not new, and has been incorporated in several theoretical approaches. Uleman (1999), for example, suggested that multiple, and even contradicting, social inferences may be activated until the context disambiguates the relevant inference. Ham and Vonk (2003) presented a similar theoretical perspective. Using the three stage models of attribution (e.g., Gilbert & Malone, 1995; Gilbert et al., 1988; Krull & Erickson, 1995; Trope, 1986; Trope & Gaunt, 1999) as a framework, Ham and Vonk proposed that multiple concepts and inferences are activated in a first categorization stage. This multiple activation occurs quickly and in a non-selective way, allowing the activation of unrelated or inconsistent concepts. In a second stage, perceivers draw intentional inferences that are guided by their current processing goals. Here, some of the previously activated concepts are inhibited, while others are more strongly activated. Only in this stage, relevant concepts are linked to specific actors about whom perceivers are forming impressions. In the third stage, previous inferences are adjusted or corrected. The idea that various potential

meanings are simultaneously activated, with the activation levels being subsequently adjusted, is also a postulate of several approaches within the word comprehension domain (e.g., Gernsbacher & Faust, 1991; Kintsch, 1988; Kintsch & Mross, 1985).

In the same way, the Automatic Causal Inferences framework (Aarts et al., 2008; Hassin, 2005; Hassin et al., 2005; Hassin, Bargh, & Uleman, 2002; see also, Aarts & Hassin, 2005) asserts that different spontaneous inferences, such as trait (Winter & Uleman, 1984), goal (Hassin et al., 2005) and predictive inferences (McKoon & Ratcliff, 1986), are included in a more general category of causal inferences. People automatically infer several social inferences, as long as they are relevant to the causal understanding of the event.

Finally, the multiple inference model (Reeder et al., 2004; Reeder, 2009), despite not explicitly concerned with spontaneous inferences, state that the perceiver make multiple inferences from observed intentional behavior, and those inferences are integrated into a coherent impression about the actor.

Under certain conditions, different kinds of inferences can occur spontaneously. Thus, we may consider whether different mechanisms and processes should be used for explaining the occurrence of the different spontaneous inferences, or whether the different spontaneous inferences should be explained within a common comprehension framework (see Leddo & Abelson, 1986; Leddo, Abelson, & Gross, 1984; Read, 1987).

2.6. Open Questions

Since the initial article published by Winter and Uleman (1984), spontaneous trait inferences research has experienced an enormous growth, providing results that challenge the traditional view of the person perception process. However, despite the exciting results provided by this new program of research, a number of questions are still open to research query. For example, do the paradigms prove that trait inferences occur during the encoding stage? Are spontaneous trait inferences descriptions of the actor, or merely descriptions of the behavior? Should the process underlying spontaneous trait inferences be described as inferential or associative?

These open questions suggest there is still a considerable ambiguity surrounding trait inference research. Next, we will describe in detail these different debates, will

look into the relevant empirical findings, and will discuss the implications of those questions for the evaluation of the spontaneous trait inference program of research.

2.6.1. Inferences at Encoding Stage

The literature on spontaneous trait inference assumes that perceivers spontaneously infer personality traits from behaviors during the encoding stage. However, according with suggestions of several authors (Corbett & Doshier, 1978; McKoon & Ratcliff, 1981, 1986, 1992; Wyer & Srull, 1989), results taken as evidence that trait inferences are unintentionally made during encoding can be explained by the operation of processes taking place only during retrieval. The principal concern is that tasks can be contaminated by explicit processes.

2.6.1.1. The Contamination Problem

A great problem of developing tasks in order to study spontaneous trait inference occurrence is the *contamination problem* (Jacoby, 1991). Within an implicit memory framework (for reviews see Richardson-Klavehn & Bjork, 1988; Roediger, 1990; Schacter, 1987; Shimamura, 1986), STI researchers have to face the difficulty of developing conceptual implicit tests (Roediger, 1990). Implicit memory tests are, by definition, those that are facilitated by previous experiences, without the participant's conscious recollection of those experiences. Implicit memory tests contrast with explicit memory tests, in which participants are overtly instructed to recollect previous experiences. Implicit tests can be perceptual (i.e., when they are dependent on perceptive processing), or conceptual (i.e., when they are dependent on semantic processing).

However, the fact that a test is named "implicit", as Hourihan and MacLeod (2007) note, does not ensure that the test is being performed in an implicit way; which is in line with Jacoby's claim that memory tests are not "pure" (Jacoby, 1991). Thus, in order to accurately describe a test as implicit, it should be assured that performance is not being contaminated by conscious recollection processes. Implicit conceptual tests are particularly vulnerable to this problem because the same semantic processing that facilitates performance on conceptual implicit tests *also* benefit explicit memory retrieval (e.g., Jacoby & Dallas, 1981). The danger is that participants become aware of

the relation between study and test phases, and use explicit recollection in order to facilitate performance. If that happens, the test can no longer be considered implicit.

The contamination has been considered the major problem of conceptual implicit memory tests (Butler & Berry, 2001; Hourihan & MacLeod, 2007). In a recent analysis, Butler and Berry (2001) considered that all conceptual implicit memory tests are questionable because they suffer from a general problem of contamination. They argued that it is almost impossible to completely rule out the possibility that intentional or conscious processes are intervening in such tasks.

Naturally, participants will only use conscious retrieval strategies if they realize that such strategies facilitate their performance in the implicit test at hand. Unfortunately, that is the case in most of the implicit conceptual tests, including those applied in the STI domain. Using intentional retrieval strategies is generally functionally-relevant for task performance. The solution would be to develop implicit conceptual tests in which participants had no benefit of explicitly retrieving the previous material (for a recent example of such a test, see Hourihan & MacLeod, 2007). Next, we'll see how the problem of alternative retrieval accounts, and the possibility of contamination, are reflected in the different paradigms used in STI research.

2.6.1.2. Encoding vs. Retrieval

In cued-recall paradigms, it was suggested (McKoon & Ratcliff, 1986; Wyer & Srull, 1989) that the trait cue efficacy can be explained by the fact that participants can use the trait probe (e.g., friendly) as a cue to generate other semantic related cues (e.g., How may someone be friendly?), including the generation of typical friendly behaviors. The generated associations, by its turn, would be likely to cue the previously presented behavior. Thus, effects may be explained by the use of a backward associative mechanism (Singer, 1979) that mainly reflects the associative strength between generated new cues and sentences (McKoon & Ratcliff, 1986). Notice that it may well be the case that participants are functioning in the same way as in behavioral pre-tests. That is, they search for behaviors that are the best representatives of the given traits. This made some authors conclude that “cued-recall experiments could not distinguish inferences generated at recall from inferences generated at encoding” (McKoon & Ratcliff, 1992, p.459).

However, there are results that support the notion that encoding processes have at least some influence in the trait-cued paradigm. For example, several encoding manipulations were shown to affect trait-cue efficacy as, for example, varying encoding processing goals (Uleman & Moskowitz, 1994) and depleting cognitive resources (Uleman et al., 1992). However, these findings do not prove that retrieval processes are irrelevant to the observed pattern.

Findings from other STI paradigms are also open to retrieval explanations. Concerning the recognition probe paradigm, McKoon and Ratcliff (1986; see also Keenan, Potts, Golding, & Jennings, 1990), asserted that the fact that participants have more difficulty rejecting a trait implied by a previously presented sentence can be due to processes taking place only after the probe word has been presented. The fact that probe traits are highly compatible with previous sentences (still in short term memory) can make the decision more difficult, independently of encoding processes. Specifically, participants may use a context-checking procedure (Forster, 1981) at the time probes are presented. By using this procedure, probes that are incompatible with the context of the sentence are easily rejected, while probes that match the context are more difficult to reject.

Uleman et al. (1996, Experiment 3) tried to rule out a retrieval account. They assumed that, if participants are using a context-checking mechanism, when they find a semantic match they will compare the probe word with the verbatim content of the sentence. Because longer sentences would take more time to check, a positive correlation was expected between sentence length and response times. Results, however, showed no correlation between sentences length and response times, which led the authors to conclude that a retrieval-checking account was not responsible for their findings. However, it is not certain whether participants are actually comparing the words of the sentence with the probe word. Given that quick responses are required, it is more likely that they base their responses in checking mechanism based on the general meaning of the sentences.

Results from the savings in relearning paradigm can also be explained, not by encoding, but by retrieval processes. In this paradigm, it is typically observed (e.g., Carlston & Skowronski, 1994) that trait-photo pairs are learned more easily when behaviors implying those traits were previously paired with the photos. This facilitation

learning effect might be explained, not by previous trait inferences, but by the fact that participants retrieve the presented behaviors during the learning phase (and possibly also during the cue-recall task). To take an example, imagine that participants see a photo of “John” paired with the behavior “Visited a sick friend at the hospital”. When they see John’s photo again, they might remember that John was the person who visited a friend at the hospital. Supported by the provided trait (“friendly”), an explicit trait inference may occur during this phase. Additionally, when John’s photo is again presented during the cued-recall task, the behavior may be again retrieved.

Aware of this fragility of the paradigm, Carlston and Skowronski (1994, Experiment 2) tried to rule out a retrieval alternative explanation. In order to do that, they included an interval of two days between exposure and learning phases and a behavior recognition task at the end of the experiment. The level of correct behavior recognition was relatively high (about 60% of correct responses). However, the savings effect was observed even when behaviors could not be recognized, and it was independent of the interval of time. A savings effect was also observed, although weaker, after an interval of 7 days between exposure and learning phases (Carlston & Skowronski, 1994, Experiment 3).

Finally, results from the false recognition paradigm (e.g., Todorov & Uleman, 2002) can also be explained by the intervention of behavioral retrieval and context-checking mechanisms (Forster, 1981). Specifically, when trait-photo pairs are presented for participants, and they are asked to indicate whether the trait was included in the behavior previously paired with the photo, participants might retrieve the behavior. Since the presented trait matches the general context of the behavior that was retrieved, false recognitions are likely to occur. Thus, results can be explained without postulating that a trait inference had occurred during encoding.

Contradicting a retrieval account, Todorov and Uleman (2002, Experiment 5) showed a false recognition effect when 120 behaviors were presented at test; a condition that would make behavior recall more difficult. Additionally, in order to test whether the false recognition effect is moderated by correct behavioral recall, Todorov and Uleman (2002) included both a recognition task (Experiment 5) and a cued recall task, in which photos were presented and participants were asked to recall the behaviors (Experiments 6). Similar to Carlston and Skowronski (1994, Experiment 2), the false

recognition effect was observed even when the corresponding behaviors were not recalled.

However, results from both Carlston and Skowronski (1994), and Todorov and Uleman (2002) showing independence between explicit recall and both the savings effect (Carlston & Skowronski, 1994) and the trait false recognition effect (Todorov & Uleman, 2002) are open to two alternative interpretations. It could be the case that both implicit tasks (i.e., the savings in relearning task, and the false recognition task) are immune to contamination aspects. This is the interpretation favored by Carlston and Skowronski (1994) and Todorov and Uleman (2002), and it implies that STI occur during encoding. The alternative interpretation poses that the observed dissociation between recall (an explicit measure) and the applied implicit measures reflect a difference in sensitivity of tests. Notice that obtaining an effect in one dependent measure, but not in other, represents a single dissociation. It is widely assumed (e.g. Shallice, 1988) that single dissociations may be explained, not by qualitative differences in underlying processes, but by quantitative differences in the sensitivity of tests. In the present case, it is possible that implicit tests used to measure trait activation are more sensitive, or produce more stable results, than tests applied to measure the behavior memory (see McDermott, 1996, for an example of how time interval increases false memories, but decreases accurate recall). If this is the case, the fact that in both studies (Carlston & Skowronski, 1994; Todorov & Uleman, 2002), additional manipulations made accurate recall of behaviors particularly difficult (a 2-day delay in the first study, and the inclusion of 120 behaviors in the second) may have made more apparent the less sensitive, and less stable, nature of explicit measures. This second interpretation is not theoretically relevant to STI research.

Proving that explicit and implicit measures are being based on different processes would require, not demonstrations of single dissociations, but evidences of double dissociations. That is, it would be necessary that researchers not only demonstrate that under some circumstances, effects are observed on implicit tests and not on explicit tests, but also that, under other circumstances, effects are obtained in explicit tests, with no (or inverse) effects in implicit tests.

Summing up, the most frequently applied paradigms in spontaneous trait inference research are open to alternative retrieval accounts. This issue is of critical

importance. The impact of spontaneous trait inference research resides in the polemical proposal that STIs occur independently of our choice, every time a behavior is encoded. For this reason, it is important that more studies are done that prove the encoding nature of the process. The weaknesses of the various paradigms make clear how difficult and challenging it is to prove the spontaneity of the trait inference process.

2.6.2. Inference about the actor or about the behavior

In addition to the encoding assumption, another major assumption underlying spontaneous trait inferences is that inferred traits are descriptions of the actors of the behaviors. Initial studies (e.g., Winter & Uleman, 1984; Winter et al., 1985) seemed to take this as a certain proposition. Researchers, however, started to question the nature of the generated traits (Bassili, 1989a, 1989b; Claeys, 1990; Higgins & Bargh, 1987; Newman & Uleman, 1993; Park, 1989; Uleman et al., 1993; Whitney, Davis, & Waring, 1994). Specifically, are the traits descriptions of actors (i.e., actor categorization) or mere summary descriptions of behaviors (i.e., behavior categorization)? It is completely different to infer that “John did something friendly” from inferring that “John is friendly”.

Notice that it is possible that participants identify the trait dimension of the behavior (“this is a friendly behavior”) without making dispositional inferences about the personality of the actor (“this is a friendly person”). The distinction between the identification and dispositional inference phase is well captured by Trope’s model (Trope, 1986).

Initial questions about the nature of the process were driven by the fact that trait-cues were not effective retrieval cues of actors. If traits specifically referred to actors, they should function as effective cues for retrieving the actors in the stimulus sentences. However, studies using the cued-recall paradigm consistently reported trait-cued recall of actors as lower than trait cued-recall of behaviors. Additionally, in some studies actor trait-cued recall was not superior to control cue conditions (e.g., Bassili & Smith, 1986; Uleman et al., 1986; Winter & Uleman, 1984; Winter et al., 1985), even when actor swere made salient by presenting their photos (Uleman et al., 1993, Experiments 1 and 2). In contrast, under explicit impression formation trait cues were shown to be efficient in recall of actors (Uleman & Moskowitz, 1994, Experiment 3).

Evidence is mixed concerning the issue of whether trait inferences are mere behavior labels or dispositional descriptions of actors. Some evidence supports the notion of traits as characterization of behaviors. For example, Bassili and Smith applied a trait fragment completion task after behavior presentation (1986, Experiment 2) and showed that trait fragments were more likely to be correctly completed when accompanied by the actor of the sentence. However, the effect was only observed under impression formation conditions, and not with memory instructions. They interpreted this result as indicating that a direct link between the actor and the trait is established only under explicit impression conditions, and no such link is formed under memory conditions.

Claeys (1990) replicated the pattern of Winter and Uleman (1984), regardless of the fact that actors were not even mentioned in the sentences. In this case, sentences started with a verb (e.g., “Watches the stars from the garden”). Since the same pattern of trait cue efficacy was observed, even when there was no possibility to make a specific person inference, Claeys concluded that spontaneous trait inferences are a characterization of the behavior, and not of the actor. However, this result doesn’t prove that, when actors are included, inferences about actor’s personalities are not drawn (see Moskowitz, 1993a).

Thus, evidence from the trait cue paradigm provides no evidence for a specific link between the actor and the inferred trait. Trait cues were shown to be effective in retrieving actors only for individuals with a high need for structure, presumably likely to engage in stronger categorization processes (Moskowitz, 1993b).

Stronger evidence in favor of an existing link between actors and traits is provided by the savings in relearning paradigm (e.g., Carlston & Skowronski, 1994) and by the false recognition paradigm (e.g., Todorov & Uleman, 2002). In fact, these methods were specifically developed with the intention of providing better tests of the hypothesized link between actors and traits. In the savings in relearning paradigm, it is typically observed that old “photo-trait” pairs are easier to learn than old photos paired with traits paired with other photos during the initial learning phase, and also than new “photo-trait” pairs (Carlston & Skowronski, 1994, Carlston et al., 1995). These findings suggest a specific learning facilitation of old photos-traits pairs, indicating the existence of a link between previous photos and implied traits.

In a similar way, results from the false recognition paradigm show that trait false recognition is higher for match trials (faces paired with traits implied about them), then for mismatch trials (old faces randomly paired with traits implied by sentences about other faces) (Todorov & Uleman, 2002). If the inferred trait was merely a description of the behavior, pairing it with the right or wrong photo should not influence trait false recognition. Again, it seems that this pattern is due to the specific link between a certain actor and a correspondent trait.

Existing evidence provide mixed findings concerning the question about the nature of the trait label. While evidence from the cued-recall paradigm cannot clarify this issue, data from savings in relearning and false recognition paradigm is more convincing. However, a major problem of both these paradigms is that photos may automatically trigger explicit impression processes. The fact that effects are relatively unaffected by processing goals supports such a hypothesis. Additionally, findings obtained with the savings in relearning paradigm show that traits can be associated with persons that are not the actors of the behavior (Carlston et al., 1995), which put into question the specificity of the actor-trait link. For all these reasons, new studies are needed to explore the question of whether inferred traits are incorporated in the mental representation of actors, or not.

2.6.2. Association or inference?

Despite the amount of existing research, the nature of the processes underlying spontaneous trait inferences occurrence is still unclear. A critical question concerns whether the process responsible for the trait-actor link is inferential or associative (Brown & Bassili, 2002; Carlston et al., 2005; Carlston et al., 2007; Carlston et al., 2008; Todorov & Uleman, 2004; Uleman, 1999). Empirical results showing links between actors and traits may be equally explained by the intervention of an inferential process (that mentally represents the trait as a characteristic of the actor) or by the operation of a relatively blind associative process (that merely links the activated trait with any actor that happens to be salient in the context).

Surprisingly, this question was apparently put aside by most of the literature. While implicitly it was assumed that trait-actors links result from inferential processes,

others were aware, and seem to support, an associative explanation of the process (Uleman, 1999). In this respect, Uleman stated that:

Spontaneous impressions are linked to actors by mere association (if they are linked at all), whereas intentional impressions are correctly linked to the logically appropriate actor. Trait concepts activated spontaneously by one actor's behavior may become associated with another person in that setting, or may not even be associated with any particular person. (p.147)

The ambiguity surrounding this topic is probably related to the fact that it is not only a complex conceptual question (e.g., what is the difference between an inference and an association), but also because it represents a challenging issue in empirical terms. For instance, the inferential or associative nature of the process is related to the question of whether traits are characterizations of behaviors or of actors. Specifically, seeing activated traits as being mere behavior labels that became associated with any element present in the retrieval context favors an associative account. In contrast, seeing inferences as being descriptions of the actors presumes the operation of “true” inferential processes. However, it should be noticed that empirical findings supporting the existence of a link between actors and traits (e.g., Carlston & Skowronski, 1994, Todorov & Uleman, 2002) are necessary, but not sufficient, to prove the inferential nature of the process. Stronger links between actors and traits may result from stronger associations being established between actors and traits (after all, actors are the more salient elements in the context). In other words, while showing that traits are mere behavior labels clearly supports an associative view, results that provide evidence for links between traits and actors are open both to inferential and associative accounts.

The debate about the associative or inferential nature of STIs was only explored in a systematic way when studies appeared showing trait transference effects (Carlston et al., 1994). Carlston and colleagues (Carlston & Skowronski, 2005; Crawford et al., 2007; Crawford et al., 2008') showed that traits can become associated with other individuals presented during the encoding context, besides the actors of the behavior, giving support to Uleman's intuitions (Uleman, 1999). However, spontaneous trait transferences are usually weaker than spontaneous trait inferences effects. A savings

effect for irrelevant actors was observed even when participants were explicitly told that photos and behaviors were randomly paired (Skowronski et al., 1998). Furthermore, in a provocative study, Brown and Bassili (2002) showed that savings effects were observed even for pairs of traits and inanimate objects (for example, bananas). These results clearly indicate that findings from the savings in relearning paradigm can occur in the absence of the intervention of inferential processes. Even more relevant, they raise the possibility that spontaneous trait inferences can be generally explained by the operation of automatic associative processes, which operate independently of more relevant person perception mechanisms.

Based on these findings, two general perspectives about the processes underlying STI and STT occurrence were proposed. On the one hand, Bassili and colleagues (Bassili, 1989b; Bassili, 1993; Bassili & Smith, 2002) argued that both STI and STT may be explained by the intervention of associative processes. On the other hand, according to Carlston and collaborators (Carlston & Skowronski, 2005; Crawford et al., 2007; Crawford et al., 2008) STTs and STIs are distinct processes, explained by the operation of different mechanisms. While STIs are based on inferential processes, STTs result from associative processes.

Carlston and collaborators (Carlston & Skowronski, 2005; Crawford et al., 2007) further specified the nature of the distinction between associative and attributional processes. They proposed that associative processing results from spatial and temporal contiguity, forming *unlabeled* associative links. This means that the existing relation between the two associated constructs is not designated (for example, the trait friendly can be associated with John, without the trait being marked as a property of John). Attributional processing, in contrast, involves deeper processing and the activation of attributional knowledge. However, attributional processing does not necessarily have to be a conscious or effortful process. Links created from attributional processing are stronger than those resulting from associative processing, and the existing relation between the two constructs is specifically labeled (e.g., friendly is targeted as a property of John).

Based on this distinction, Crawford et al. (2007) presented three existing differences between STIs and STTs. First, STIs involve attributional processing and STTs involve only associative processing. Second, as a result of different processing, an

associative link is created in STTs, and an inferential link is created in STIs. Third, in later judgments, participants directly retrieve the prior attribution in STIs, while in STTs they construct an inference, this process being implicitly influenced by the existing association.

In order to sustain this perspective, authors (Carlston & Skowronski, 2005; Crawford et al., 2007) reported several studies showing differences between STIs and STTs. First, STI effects are usually stronger than STT effects (Bassili & Smith, 2002; Skowronski et al., 1998). Second, the fact that random pairing of behaviors and photos reduces the STI effect was taken as evidence that other processes (i.e., attributional) besides associative processes intervene in the effect (Crawford et al., 2007; Skowronski., 1998). Third, it was argued that in the case of STIs, people apply implicit theories of personality, using the created inferential links to generalize to other trait dimensions. Accordingly, several studies (Carlston & Skowronski, 2005; Crawford et al., 2007a; Crawford et al., 2007b; Skowronski et al., 1998) showed that trait generalization effects are more likely for actors than for communicators of the behaviors. Fourth, negativity effects (i.e., greater impact of negative behaviors on impressions) are observed for self-descriptions, but not for other-descriptions (Carlston & Skowronski, 2005; Crawford et al., 2007a; Crawford et al., 2007b). Since authors argued that negativity effects are characteristic of attributional processes, this was taken as further evidence that attributional processing is present in STIs, but absent in STTs. Fifth, a concurrent inferential task in which participants had to detect if communicators were lying or telling the truth reduced STI magnitude, but had no effect on STTs (Crawford et al., 2007b; Skowronski et al., 1998). It was assumed that the lie detecting inferential task competes with the inferential processing, but not with associative processing.

The differential effects of these manipulations were taken as supporting the notion that spontaneous trait inferences reflect attributional processes, while spontaneous trait transference reflect associative processes (Carlston & Skowronski, 2005, Crawford et al., 2007a; Crawford et al., 2007b). However, any of these differences might also be explained by different degrees of associative strength. In this purpose, we agree with Crawford et al. (2007) in that:

...the case that those processing differences exist is by no means conclusively made (...) differences in association strength, and not the postulated differences in processes underlying STI and STT (inferential vs. associative), might account for these effects. (p.679)

Furthermore, asserting that spontaneous trait inferences involve *attributional* process is resuming the question about whether trait inferences and causal attribution should be seen as distinct processes. However, this question was apparently solved. The importance of distinguishing between trait inferences and causal attribution was amply supported, both by theoretical and empirical bases (Erickson & Krull, 1999; Hamilton, 1988, 1998; Hewstone, 1989; Smith & Miller, 1983). The point is clearly made by Smith and Miller (1983), who demonstrated that trait judgments take as much time to be made as gender judgments, while causal judgments take longer times. They concluded that dispositional inferences occur in a quick way as part of behavior comprehension, whereas causal attribution does not occur automatically. The distinction between causal attributional and dispositional inferences is clearly made by Hamilton (1998). According with this author, while causal attributional processes refers to considering the causes of the behavior (i.e., Why that happened?), dispositional trait inferences involve going from a behavior to a trait inference about the actor (i.e., “He did something intelligent so he must be an Intelligent person”). In other words, trait inferences may occur independently of considering the reasons that conduct to the behavior. This is especially true for *spontaneous* trait inferences. In this case, the causal factors that drove the behavior are unlikely to be considered. In addition, there are no empirical reasons to think that STIs might involve the consideration of causal factors. In sum, the distinction between the inferential and attributional processes seems well established. Thus, asserting that STIs involve attributional processes in order to distinguish them from STTs (e.g., Crawford et al., 2007) and taking the variables known to affect causal attribution in order to explore whether they differently affect STIs and STTs is a debatable research approach.

As a conclusion we should say that the question about whether spontaneously inferred traits became part of the representation of the actor, or are merely the result of

accidental associative processes, is not answered yet. Further studies are certainly needed to explore the mechanism underlying STI and STT occurrence.

2.7. Summary of the Chapter

In 25 years of research we learned a lot about the trait inference process. Initial evidence that perceivers infer traits from behaviors, without intentions to form impressions, was provided by Winter and Uleman (1984). Since this initial study, different paradigms were refined to study the conditions under which spontaneous trait inferences occur. Four paradigms have been mostly applied: the cued-recall paradigm (Winter & Uleman, 1984); the recognition probe paradigm (e.g., Newman, 1990; Uleman et al., 1996); the savings in relearning paradigm (e.g., Carlston & Skowronski, 1994); and the false recognition paradigm (e.g., Todorov & Uleman, 2002).

The automaticity question was one of the main questions explored within this literature. Existing data, however, show that the spontaneous trait inference process is not unconditionally automatic, being eliminated, for example, by shallow processing goals (Uleman & Moskowitz, 1994) and high levels of cognitive load (Uleman et al., 1992). Thus, it is not the case that spontaneous trait inferences occur every time a behavior is processed. Instead, the process seems to have some cognitive flexibility.

In line with this view, other social inferences were found also to occur spontaneously, such as gist inferences (e.g., Winter et al., 1985), situational inferences (e.g., Lupfer et al., 1995), goal inferences (Hassin et al., 2005), predictive inferences (McKoon & Ratcliff, 1986), and others (see for example, Gernsbacher et al., 1992; Ham & Van den Bos, 2008; Murphy et al., 2009). This research suggests, first, that trait inferences are not necessarily more relevant to social perception than other types of social inferences, and second, that the occurrence of these different social inferences might be governed by common cognitive principles. Thus, giving greater relevance to trait inferences, and then trying to prove its automaticity, might not be the best approach if we want to understand the principles by which spontaneous trait inferences are governed.

Besides the automaticity issue, other assumptions about the nature of the STI were found difficult to prove empirically. The hypothesis that guided the STI literature was that perceivers spontaneously infer personality traits about actors during behavior

encoding (Winter & Uleman, 1984). Based on the literature, we verify that the three assumptions implied by this statement (i.e., STIs occur during encoding, are specifically linked to the actor, and involve inferential processes) remain unclear.

First, evidence that STIs occur during encoding may be generally explained by the intervention of processes taking place only during retrieval (e.g., Corbett & Doshier, 1978; McKoon & Ratcliff, 1981, 1986, 1992; Wyer & Srull, 1989), despite some findings supporting the intervention of encoding mechanisms (e.g., Uleman & Moskowitz, 1994).

Second, it is not certain whether trait inferences are characterizations of the actor, or merely behavior labels. Stronger evidence for an existing link between actors and traits comes from the savings in relearning paradigm (e.g., Carlston & Skowronski, 1994) and from the false recognition paradigm (e.g., Todorov & Uleman, 2002). However, the possibility that explicit impression formation processes intervene in these procedures cannot be completely ruled out.

Finally, it remains open whether links between traits and actors are due to the accidental co-occurrence in working memory of traits and any actor that happen to be salient in the context, or whether specific dispositional processes are involved (e.g., Brown & Bassili, 2002; Carlston et al., 2005; Carlston et al., 2007; Carlston et al., 2008). The fact that traits may be transferred to individuals that are merely describing others behaviors (Carlston et al., 1995), and even to inanimate objects (Brown & Bassili, 2002), is a strong argument in favor of the fact that associative mechanism have at least some influence in the process.

How can we clarify all these controversial debates in order to give a clearer picture of the principles governing spontaneous trait inferences? We think that an analysis of the text comprehension literature may be useful in understanding the trait inference process.

CHAPTER III

TEXT COMPREHENSION

The text comprehension literature may provide us critical inputs into the understanding of the STI process. Crucial to our proposal, an analysis of the text comprehension research gives support to a more flexible and contextually-dependent view of the spontaneous trait inference process. But why is the text comprehension literature so crucial for STI research? The reasons can be summarized in two major points.

First, one central assumption underlying the text literature is that comprehension involves inferential thinking (e.g., Johnson, Bransford, & Solomon, 1973; Sharkey, 1986). As a consequence, a major focus of the field has been on exploring how, and which, inferences are made online during text comprehension (Graesser & Bower, 1990; Graesser et al., 1994; McKoon & Ratcliff, 1990a; 1990b; 1992; Rickheit & Strohner, 1985). Thus, if there is any literature that can provide us with valuable information about “inference generation,” it is without doubt the text comprehension literature. Second, the text comprehension literature represents an enormous and highly elaborated field. As a result, several stimulating reviews have been published throughout the years (Balota, Flores d'Arcais, & Rayner, 1990; Britton & Graesser, 1996; Goldman, Graesser, & van den Broek, 1999; Graesser, Gernsbacher, & Goldman, 2003; Graesser, Millis, & Zwaan, 1997; McKoon & Ratcliff, 1992; Rickheit & Strohner, 1985). Implications of text comprehension research are important, not only for the understanding of language comprehension but also for the process of understanding in general. Researchers from this domain soon realized that, in order to understand the type of inferences that are made during comprehension, it was also crucial to understand the general principles that guide human understanding. As Shank (1975) noted:

It has been apparent to researchers within the domain of natural language understanding for some time that the eventual limit to our solution of that problem would be our ability to characterize world knowledge...Thus we would extend our previous view of language analysis to the problem of understanding in general. (p. 117)

The multidisciplinary effort of psychologist, linguistics, and artificial intelligence researchers has made the text comprehension literature reach an uncommon

level of sophistication, both at experimentally and theoretically. Especially remarkable is the existence of a high number of theoretical models (e.g., “constructionist theory” by Graesser et al., 1994; “construction-integration model” by Kintsch, 1988; 1998; “minimalist model” by McKoon & Ratcliff, 1992; “knowledge access network model” by Sharkey & Mitchell, 1985; “causal network model” by Trabasso & Magliano, 1996; Trabasso & Sperry, 1985). Because they are sufficiently explicit about underlying representations and processes, these models allow the formulation of specific predictions that can be submitted to empirical scrutiny. This is particularly relevant when we consider that STI research has progressed in a strange theoretical silence, almost in the absence of integrative theoretical accounts.

Given the emphasis of both STI and text comprehension literatures on inferential processes, it is not surprising that researchers from both domains frequently face the same kinds of difficulties. In fact, from almost every research problem that we find in the STI domain, we can hear an echo from text comprehension researchers trying to solve exactly the same question. For example, both fields have been trying to find effective paradigms in order to test the occurrence of online inferences (e.g., Keenan et al., 1990; Uleman et al., 1996), and both have been faced with the problem of existing alternative retrieval accounts (see Carlston & Skowronski, 1994; Corbett & Doshier, 1978; McKoon & Ratcliff, 1992; Todorov & Uleman, 2002; see also Johnson, Bransford, & Solomon, 1973). It is difficult to understand how the two literatures have developed in such parallel lines, even with some voices arguing for a stronger integration between the two fields (e.g., Read & Miller, 1993; Read & Miller, 2005). While in methodological terms, several paradigms from the text comprehension literature have been imported to the study of STIs, the theoretical consequences of the text comprehension research to the study of STIs have not generally been considered.

In the present work we claim that an examination of the text comprehension literature can be very important for generating hypotheses about the nature of the STI process. In the present chapter, we summarize those aspects of the text comprehension literature that are simultaneously essential for describing the nature of the field and also are relevant to the study of STIs. We start by introducing a major problem in simulating artificial systems of comprehension: the “explosion of inferences” problem (Schank, 1987). Then, two general comprehension frameworks will be described: the script

approach (Schank & Abelson, 1977) and the situational model approach (Graesser, Swamer, Bagget, & Sell, 1996). Both of these frameworks attempt to describe how the mental system is able to reach comprehension while avoiding the explosion of inferences problem. However, they are too general in order to sustain specific predictions about how and which inferences are generated during comprehension. Later in the chapter, we will describe two approaches that were developed with the specific aim of explaining inference generation: the constructionist theory (Graesser et al., 1994) and the minimalist approach (McKoon & Ratcliff, 1992). According to the constructionist theory, readers generate online those inferences that are necessary to the construction of a coherent and meaningful situational model of the event. Recently, however, the minimalist approach has rejected a constructionist view of inference generation. Minimalist researchers claim that readers only generate automatically inferences that are easily available and inferences necessary for local coherence. The minimalist approach is the one that presents greater specificity in terms of processing mechanisms, and has generally been supported by empirical evidence. We continue the chapter by describing the general implications of the text comprehension literature to STI research, and end by specifying how the minimalist framework can be used to generate principles about the conditions under which the STI process occurs. As we will see, these principles support a more flexible view of the STI process.

3.1. The Explosion of Inferences

According to *top-down* or constructivist perspectives, previous knowledge is an essential element of the comprehension process (Anderson, 1990; Emmott, 1997). In opposition, *bottom up* perspectives (e.g., Johnson-Laird, 1987; Perfetti & McCutchen, 1986) claim that words possess a literal, “zero-context” meaning that is independent of world knowledge and context. This view claims that the comprehension of a sentence is reached through the sum of the analysis of its constituent’s parts (for a relevant discussion, see Anderson, 1990).

However, independently of the fact that words might potentially activate stereotyped senses, construction of *meaning* seems to generally imply much more than that (Anderson, 1990; Emmott, 1997). Most researchers endorse the view that existing knowledge is intrinsically incorporated within the comprehension process (Anderson,

1990). Evidence for this argument comes from different lines of evidence. For example, when we read a sentence or narrative, there are ingredients essential for their understanding that are not explicitly mentioned (e.g., Bach, 1994). However, those non-explicit aspects can be easily inferred. Taking as illustration a sentence from Schank and Abelson (1977, p. 9), if someone says “I like apples” we immediately understand that the person is talking about eating, and it is superfluous, even apparently redundant, to explicitly mention to it (i.e., “I like eating apples”). In addition, there are many indirect, metaphorical, and figurative sentences whose meaning cannot be reached by summing up the meaning of their components. If we read that “Beth drank the whole bottle” (Anderson, 1990), we know that Beth drank the contents of the bottle, and not the bottle itself. In agreement with this, and contradicting the notion that people always compute both the literal meaning (i.e., meaning of the sentence independently of extra-linguistic knowledge, see Gibbs, 1984) and the utterance meaning of sentences (i.e., what the speaker means in a particular context, see Gibbs, 1984) suggested by some authors (e.g., Clark & Lucy, 1975; Searle, 1979), several studies have shown that the literal meaning of sentences is not always accessed, and that people do not have to first compute the literal meaning of indirect sentences in order to reach their utterance meaning (Gibbs, 1979; 1983; 1984; 1986; 1994; see also Kintsch, 2000). Finally, discourse is filled with words that have a number of potential meanings (Granger, Eiselt, & Holbrook, 1986). In these cases, people must be able to disambiguate those words, selecting the interpretation that is more appropriate to the specific context in which the word is being mentioned.

These examples show that knowledge is so pervasively implied within comprehension that sometimes we don't even notice its intervention. The crucial role of knowledge in filling the gaps of text and discourse was recognized by language computational researchers, who were generally convinced that in order to create artificial comprehension systems, background knowledge should be incorporated into those systems (e.g., Charniak, 1977; 1978; DeJong, 1979; Shank & Abelson, 1977). However, initial attempts to create such systems were marked by serious difficulties. The major problem was that the number of inferences that could be drawn by knowledge-based systems was very large, resulting in an *explosion of inferences* (Rieger, 1975; Schank, 1975). Even simple sentences could potentially lead to the

generation of an unrestricted number of inferences (Ram & Hunter, 1992). Of course, this unlimited ability to generate inferences is pernicious to an effective comprehension system. Not all inferences are equally relevant in a given situation. A mind capable of computing all possible inferences in any given moment would get “lost in thought”, and would be incapable of extracting coherent meaning from the information that came in into the system. Somehow, these systems should be able to control the inferential process in order to restrict, or select, those inferences that are most valuable in a given circumstance.

Thus, despite being evident that it was necessary to incorporate knowledge into comprehension systems, the specific mechanisms by which stored knowledge interacted with online comprehension were much less clear. One possible solution would be to incorporate more specific knowledge, such that it can be more functional to the system. In line with this view, Schank and Abelson (1977) suggested that the inferential explosion problem could be solved by incorporating script-based knowledge into comprehension computer programs.

3.2. *Scripts Approach*

A script is a “predetermined, stereotyped sequence of actions that defines a well known situation” (Shank & Abelson, 1977, p.41). Scripts can be viewed as a type of *schematic* knowledge structure (Abelson, 1981; Alba & Hasher, 1983). The concept of schema has a very long tradition within psychology, originating with works from Piaget (Piaget, 1926), and Bartlett (1932). A key feature that differentiates scripts from other kinds of cognitive structures is the fact that they represent the temporal sequence of actions within a specific event. Thus, scripts provide us not only with information about the typical actions that occur in a particular context, but also about the *order* in which those actions occur.

The emphasis on schema-based knowledge structures was reinitiated by a movement brought about by psychologists and artificial intelligence researchers at Yale University (Schank & Abelson, 1977; Schank, 1982; Galambos, Abelson, & Black, 1986). These researchers were interested in understanding how knowledge is organized in the human mind, and in the type of knowledge structures that allow people, and

computers, to understand social information. According to these researchers, people represent specific world knowledge about familiar contexts in the form of scripts.

The script notion emerged, and gained strength, due to its functionality for computer modeling (“SAM” – Script Applier Mechanism – was the first computer program implementing the script concept, see Cullingford, 1978). Incorporating script-based knowledge into comprehension systems would be a potential way of resolving the inferential explosion problem (Charniak, 1977; 1978; Shank & Abelson, 1977). A system of knowledge organized by scripts would not have to compute all potential inferences, but only those that are relevant to the script that is active in a particular moment. Using an example taken from Sharkey (1986), such a comprehension system would not mistakenly interpret the word “serve” as the initial step of a tennis match if the “restaurant script” is active. In such a context, “serve” would be rightly interpreted as “presenting the food”. The ability of activating the right meaning, in the right context, represents a greater advantage to any comprehension system.

The benefits of conceptualizing knowledge as being organized in scripts structures is linked to their function of generating *expectations* about what is normal to happen in a given familiar situation, and about the order in which events typically occur. These expectations have a number of consequences that are highly economical to an effective comprehension system (Schank & Abelson, 1977, 1995; Shank, 1982; Galambos, Abelson, & Black, 1986). First, script-based expectations facilitate the interpretation of events, simplifying processing of information that matches the script content and alerting for the presence of deviations from the script. We will take the very well known “restaurant script” (Shank & Abelson, 1977) to illustrate this point. People tend to share a script about the sequence of actions that are typical to occur when we go to a restaurant (i.e., enter into the restaurant, sit, receive the menu from the waiter, choose the food, eat, pay, and leave the restaurant). Thus, scripts provide us with knowledge about the sequence of actions, about the objects, and about the actors that are likely to be observed in the context of a restaurant. Because of that, any event that is part of the script is promptly interpreted. We don’t have to waste time, for example, thinking about the reasons why the waiter is giving us a menu. Notice that script activation also explains why we use the definite article “the” to mention “waiter” in a restaurant context. The presence of “the waiter” is contemplated by the restaurant script.

As previously said, script activation is also useful to detect the presence of inconsistencies. For example, we find it normal when the waiter gives us the menu in the context of a restaurant, but if we get on a bus and the driver gives us a menu, that will be very surprising. The reason is that this behavior completely violates our script about “taking the bus”.

Second, script-based expectations guide people’s behavior (Abelson, 1981; Schank & Abelson, 1995). The activation of the restaurant script, for instance, allows people to know exactly how to act and react within the context of a restaurant. Cognitive resources can then be saved for other, more relevant, tasks.

Third, interaction and communication between people is facilitated by the existence of script-based expectations (Abelson, 1981; Schank & Abelson, 1995; for effects of communication facilitation due to script activation on pre-schoolers, see Nelson & Gruendel, 1979). If two people share the same script, they don’t have to spend effort explaining why they are acting as they are. The meaning of their actions is implicitly provided by the script. This explains why writers and communicators omit certain details from their narratives. They have an implicit understanding about the type of knowledge that their readers and listeners have, and avoid filling the discourse with unnecessary details.

The script concept is theoretically sound. The activation of a script has the potential to facilitate the interpretation process, direct people’s behavior, and fill in spaces left empty by communicators. Existing evidence gives support to the intervention of script-based knowledge in the process of comprehension. Few studies, however, have provided data about the specific role of scripts in inference generation.

3.2.1. Empirical Evidence for the Role of Scripts

The notion that scripts play an important role in the comprehension process has been generally supported. Bower, Black, and Turner (1979), for example, showed that people tend to share a common knowledge about the sequence of actions that compose typical events. Bower et al. (1979) further demonstrated that participants tend to reorganize their recall of previous scrambled stories, according to their script structure.

Existing empirical data also support the role of scripts in interpretation. It was shown that vague stories are better remembered (Bransford & Johnson, 1972; Dooling

& Lachman, 1971) and comprehended (Bransford & Johnson, 1972) when participants are provided with the topic or theme of the story before reading it.

Another major line of evidence has explored how scripts may result in memory bias. Several studies demonstrated that participants tend to falsely recognize not mentioned script-consistent information (Bower et al., 1979; Dewhurst, Holmes, Swannell, & Barry, 2008; Graesser, Woll, Kowalski, & Smith, 1980; Lampinen, Faries, Neuschatz, & Toglia, 2000; Walker & Yecovich, 1984), with the script's central concepts being more often misrecognized than peripheral concepts (Dewhurst et al., 2008; Lampinen et al., 2000; Walker & Yecovich, 1984). In addition, while false recognition of peripheral concepts tends to increase with the number of text implications, recognition errors for central actions were shown to be independent of the number of text references implying those actions (Walker & Yecovich, 1984). This pattern was interpreted as an indication that scripts selectively activate information according to their relevance, with central actions being stronger activated (Walker & Yecovich, 1984).

Miller and Gazzaniga (1998) found a pattern of false recognition similar to the one observed in previous studies, but using pictures of familiar scenes. It was observed that the level of false recognition of items highly predicted by the events was as high as the level of accurate recognition, a pattern similar to the one obtained with lists of words in the DRM paradigm (e.g., Roediger & McDermott, 1995). In the DRM paradigm participants are presented with list of semantic associates of a critical nonpresented word. Results show that critical items are recalled as much as words presented in the middle of the list, and are more likely to be falsely recognized than presented words (Roediger & McDermott, 1995). Neuschatz and collaborators reported a pattern similar to Miller and Gazzaniga (1998) using videotaped material (Neuschatz, Lampinen, Preston, Hawkins, & Toglia, 2002).

In contrast, script research has shown that participants tend to exhibit better memory discrimination for script-*atypical* information (Bower et al., 1979; Lampinen et al., 2000). Atypical script memories are also more likely to include vivid details, and to lead to more subjective feelings of “remember”, when the remember/know procedure is applied (Tulving, 198; see also Rajaram & Roediger, 1997). These findings are in line with the role of scripts in alerting the system for inconsistencies, and are well

incorporated by Script Pointer Plus Tag models (Graesser, Gordon, & Sawyer, 1979; Shank & Abelson, 1977), and by the Dynamic Memory model (Schank, 1982). Both of these classes of models predict that inconsistent information is more likely to be indexed episodically as an exception to the script.

The fact that script consistent information tends to be falsely recognized was interpreted by some authors as indicating that inferences are made during encoding (e.g., Dewhurst et al., 2008). That would cause participants to have difficulty distinguishing presented information from information that was merely inferred. However, false recognition results can equally be explained by processes taking place only during retrieval (Singer & Leon, 2007).

Therefore, previous studies do not clarify whether script-based inferences are made during encoding. Nevertheless, the way script structures were usually conceptualized suggests that they may have an important role within the inference process. Abelson (1981), for instance, proposed that scripts function by pre-activating script-relevant information. In a similar way, the Knowledge Access Network (KAN) model (Sharkey, 1986; Sharkey & Mitchell, 1985) describes scripts as networks of script-relevant concepts linked to a central concept (i.e., *script subnode*). The network is assumed to be regulated by parallel spreading activation principles (Collins & Loftus, 1975; Anderson, 1983), such that the activation of one concept node depends on the strength of association with other activated concepts. According to Sharkey, this model has the advantage of representing script knowledge as being interrelated with the functioning of the rest of our memory. In fact, according to the model, even information that is only remotely related with the script may receive residual activation.

The pre-activation framework developed by Sharkey (Sharkey, 1986; Sharkey & Mitchell, 1985) accommodates well several data, including the fact that priming of one action of the script results in faster recognition of actions belonging to the same script (Anderson, cited by Abelson, 1981); findings showing that priming of the script name leads to faster recognition of script-related information (den Uyl & van Oostendorp, cited by Anderson, 1981); and the faster lexical decisions for script related words (compared with unrelated or neutral words) after participants were presented with sentences evoking the script context (Sharkey, 1986).

These results support the view that script instantiation results on the decrease in the threshold for script-relevant information. Apparently, the mind becomes predisposed to process certain pieces of information. The implication for inferential processes is that scripts have the potential of orienting cognitive processing to those inferences that are more appropriate to the script, with script-based information having a higher probability of being involved in inferences than unrelated information.

Some authors suggested that script information is not merely pre-activated, but is actually inferred and treated as if it had been mentioned (Shank & Abelson, 1977). This argument is in line with the notion that one of the most important properties of scripts is their *gap-filling* function. According to this perspective, scripts allow readers to make inferences beyond what is explicitly referred in the text (Dewhurst et al., 2008; Schank & Abelson, 1977). For example, if we know that some actions of the script took place (e.g., someone entered into a restaurant, ordered food, and left the restaurant) we may infer that other script-actions had also occurred (e.g., that the person also had read the menu, eaten, and paid the bill) (Dewhurst et al., 2008).

The problem is that it is not clear how much of the script is actually inferred (see Abelson, 1981). In fact, as some authors note (Rickheit, Schnotz, & Strohner, 1985), it would be very cognitively consuming to infer all the details of an experienced script-situation. Aware of this fact, Shank and Abelson (1977) raised the possibility that, instead of filling the mind with superfluous details, a mark is established in order to indicate that the script occurred, such that script-information may be used later, if necessary.

In sum, the script concept has an extremely important role in clarifying how social reality is mentally represented. The nature of knowledge is one of the “big” questions within psychology, and script research has provided useful contributions concerning this issue. Some authors, however, claim that script theory suffers from a general problem of lack of specificity that prevents precise predictions about how scripts affect text processing (Rickheit et al., 1985). Maybe due to this fact, script research provides us with few empirical data about what type of script-based inferences are made, when those inferences occur, and whether script-based inferences have an automatic or optional nature. These questions have been more closely addressed by other researchers (e.g., McKoon & Ratcliff, 1992).

3.3. Situation Models

Many discourse comprehension researchers endorse the view that comprehension involves the construction of a *situation* model (van Dijk & Kintsch, 1983) or *mental* model (Johnson-Laird, 1983) about the state of affairs of what is conveyed in the text (for a review, see Zwaan & Radvanski, 1998). A situation model (or mental model) can be described as a “lifelike mental representation of the people, settings, actions, goals, and events explicitly mentioned, or inferentially suggested by the text” (Graesser et al., 1996, p.12). It was generally assumed that situation model representations coexisted with lower levels of text representation (i.e., surface and textbase propositional representations) (Graesser et al., 1996; Graesser et al., 1997; Johnson-Laird, 1983; van Dijk & Kintsch, 1983; Wyer, 2004).

The distinction between the situational model concept and the script concept should be clarified. While a script is a stereotyped representation of an event, a situation model is a representation of a specific event occurrence. This distinction can be seen as reflecting a difference between a semantic general representation of an event (i.e., script) and an episodic event representation (i.e., situational model). Zwaan & Radvanski (1998) illustrated the distinction using the restaurant script. While a script stores information about the sequence of typical events that occur in a visit to a restaurant, a situation model describes one specific episode about going to a restaurant. In fact, the situation model framework was developed, in part, by the idea that comprehension should not be conceptualized as the mere activation, and application, of the right script structure. However, script-knowledge, as well as other structures of knowledge, is assumed to be applied when a situation model is constructed (Graesser et al., 1994; van Dijk & Kintsch, 1983; Zwaan & Radvanski, 1998).

While the notion that perceivers represent narratives using situational models is endorsed by several authors, less agreement exists about the specific nature of the situation models. What exactly is a situational model is something that researchers have been trying to clarify. One general approach to this question characterizes situation models as multidimensional structures (Gernsbacher, 1990; Johnson-Laird, 1983; Zwaan, Langston, & Graesser, 1995; Zwaan & Radvanski, 1998). According to this perspective, situational models contain information about five different dimensions: time, space, causation, intentionality, and protagonist. The Event-Index Model

developed by Graesser and collaborators (Graesser et al., 1995) postulates that, as readers comprehend a story, they monitor these five dimensions. Events that share dimensions with previous information (for example, events that are contiguous in terms of time or location) are easier to integrate into the current situational model. In opposition, events whose indices mismatch previous information are difficult to integrate in the current situation model, and a new situation model is likely to be initiated (these notions are also incorporated within the Structure-Building Framework, see Gernsbacher, 1990). For example, if we are reading a story about a protagonist in a specific context, and the next paragraph introduces a different protagonist acting in a totally different context, a new situational model is likely to be created in order to comprehend the new episode, and later the two episodes can then be integrated in a global model (see Zwaan & Radvanski, 1998).

In support of the model, studies have shown that participants take more time to read events that present discontinuities in terms of temporal, causal, goal-related, and protagonist information (Rinck & Weber, 2003; Zwaan, Magliano, & Graesser, 1995; Zwaan, Radvansky, Hilliard, & Curiel, 1998). These studies found weaker evidence for the role of space discontinuities. Most of the other existing studies, however, have explored the role of each dimension separately. This research has generally confirmed that readers tend to monitor space (Ehrlich & Johnson-Laird, 1982; Glenberg, Meyer, & Lindem, 1987; Haenggi, Kinstch, & Gernsbacher, 1995; Morrow; Bower, & Greenspan, 1989; Morrow, Greenspan, & Bower, 1987; Rinck, Williams, Bower, & Becker, 1996), time (Bestgen & Vonk 1995; Claus & Kelter, 2006; Ohtsuka & Brewer 1992; Mandler, 1986; Speer & Zacks, 2005), causality (Keenan, Baillet, & Brown, 1984; Deaton & Gernsbacher, 1997; Klin & Myers 1993; Magliano; Baggett, Johnson, & Graesser, 1993; Millis & Just, 1994; Myers, Shinjo, & Duffy, 1987; Singer & Ferreira, 1983; Singer, Halldorson, Lear, & Andrusiak, 1992), intentionality (Dopkins 1996, Long et al., 1992; Lutz & Radvansky, 1997; Suh & Trabasso 1993), and properties of the protagonist (Albrecht & O'Brien 1993; Carreiras, Garnham, Oakhill, & Cain, 1996; Hakala & O'Brien 1995; Myers, O'Brien, Albrecht, & Mason, 1994; O'Brien & Albrecht 1992), while reading narratives.

The situational model framework, and the event index model in particular, are useful in providing important information about the general principles of generating

coherent mental representations of narratives. Nevertheless, they are silent about specific processing mechanisms. For example, it is assumed that readers construct a situation model during comprehension that incorporates explicit information, as well as inferences that are relevant to the understanding the event. However, which inferences are generated online and incorporated into the mental representation of the text is something that is not specified. Are inferences generated continuously during comprehension and do they include every detail of the situation? Or are inferences only computed when they are necessary to fill a gap, and guarantee local coherence?

These are critical questions. While the field is consensual in relation to the fact that text representation includes not only the information mentioned but also extra inferences, there is less concordance about the type of inferences that are systematically generated during comprehension. In this regard, there are two main general approaches: the constructionist approach (Graesser et al., 1994) and the minimalist approach (McKoon & Ratcliff, 1992, 1995). These two approaches make contrasting predictions about the type of inferences that are usually generated online. While the constructionist view claims that inferences are made both to reach local and global coherence, the minimalist approach argues that only a limited number of inferences are encoded online.

3.4. Constructionist Approach

The constructionist theory was developed in order to account for the type of knowledge-based inferences that are generated during the construction of a situation model. Knowledge-based inferences were defined as resulting from the activation of long term memory contents. These contents are assumed to be incorporated into the mental representation of the text (Graesser et al., 1994; Graesser et al, 1996). If information is activated without being incorporated into the mental representation, then it is said that an inference did not occur, but only a transient activation of information (Graesser et al, 1996).

The basic principle of the constructionist theory is that text processing is guided by a *search after meaning*. Graesser et al. (1994) specified this principle in three distinct assumptions. First, the construction of the situation model is in line with reader's processing goals (*reader goal assumption*). Second, the *coherence assumption* assumes that the reader tries to reach text coherence at both local (i.e., organization of adjacent or

of a small number of sentences) and global levels (i.e., integration of chunks of information in higher order elements). Third, according to the *explanation assumption*, the reader tries to explain why actions, events, and states are being referred to in the text.

Essentially, the theory assumes that the reader tries to construct a meaningful and coherent representation of the text. This representation will be as deep as possible, including the information available, the goals of the reader, and reader's previous knowledge (Graesser et al., 1996). The only occasions in which a meaningful representation of the text is assumed not to be constructed are when: a) the text is unstructured or incoherent; b) the reader does not have the required background knowledge; and c) the reader has superficial processing goal (e.g., proofreading) (Graesser et al., 1996).

In this process of search after meaning, constructionist theory predicts that there are a number of inferences that are usually generated online: *referential* (i.e., when a word is referential of another word explicitly mentioned in the text), *superordinate goals* (i.e., the goal or intention underlying actions), *causal antecedents* (i.e., factors that caused the events that are mentioned in the text), *global thematic* (i.e., inferences about the global meaning, or gist meaning, of chunks of information) and *character emotional* inferences. In opposition, inferences about *causal consequences* or future episodes, *subordinate goals* (i.e., inferences about details of how actions were executed), *instruments*, and *state* inferences (i.e., that include character traits, properties of objects, and spatial relationships among entities) are predicted to *not* usually occur online (Graesser et al., 1994; Graesser et al., 1996). Graesser and colleagues (see Graesser et al., 1996) argued that inferences about subordinate goals and states imply unnecessary levels of detail, which are usually beyond the reader's goals. On the other hand, inferences about causal consequences were regarded as being cognitive useless, given the multiplicity of possible outcomes. According to the model, there are only two conditions in which such elaborative inferences are said to occur online. First, when the reader has specific goals that imply the generation of those inferences. Second, in order to satisfy "convergence and constraint satisfaction" requirements (Graesser et al., 1994, p.377). That is, when inferences are highly activated by multiple sources, or highly constrained by the context.

Constructionist theory authors stressed the fact that their aim is to distinguish between online and offline inferences, and not necessarily on the automatic character of those inferences. Also, they believe that the online-offline distinction should be regarded as a continuum, and not as a rigid dichotomy. The nature of inferences may vary depending on reader's goals, previous knowledge, and experimental task (Graesser et al., 1996).

Graesser and collaborators (1994) reviewed some evidence in favor of their theory. The occurrence of referential inferences, such as anaphoric inferences (i.e., inferences that relate a word that refers back or substitutes a previous word, with their referent) has been consensually demonstrated in the literature (Bever & McElree, 1988; Clark & Sengul, 1979; Dell, McKoon & Ratcliff, 1983; Duffy & Rayner, 1990; Gernsbacher, 1990; Haviland & Clark, 1974; McKoon & Ratcliff, 1992; O'Brien, Duffy, & Myers, 1986; Ratcliff & McKoon, 1978; Sanford & Garrod, 1981). For example, O'Brien et al. (1986) tested the occurrence of anaphoric inferences using both a probe recognition task (Experiment 1) and a naming task (Experiment 2). They found faster responses to a probe word when the last sentence of a previous paragraph included an anaphor that made reference to it. In another study, Swinney, Ford, and Bresnan (cited by Swinney & Osterhout, 1990) presented evidence in favor of the fact that pronouns tend to activate their referents. In this study, participants heard sentences and had to make lexical decisions about referents at different moments of sentence presentation. Facilitation decision effects were found only when referents were presented after the correspondent pronouns, and not after other parts of the sentence. In a similar vein, Swinney and colleagues (Swinney, Ford, Frauenfelder, & Bresnan, 1987, cited by Swinney & Osterhout, 1990) presented sentences like: *the policeman saw the boy that the crowd at the party accused of the crime*. Results revealed faster lexical decisions after the verb "accused" for words related with the referent (boy), but not for words related with "crowd", or control words.

Existing data also supports the occurrence of superordinate goal inferences. Long and colleagues (Long et al., 1992; Long, Golding, Graesser, & Clark, 1990), using different online measures, demonstrated that inferences of superordinate goals are more likely to occur online than subordinate goal inferences. Evidence for online causal antecedent inferences was also presented by different authors. Magliano et al. (1993)

used a lexical decision task and showed that causal antecedent inferences occur online, contrary to causal consequence inferences, which are not generated online. Long and colleagues (Long et al., 1990) compared the online nature of causal antecedent inferences and state inferences, using a lexical decision task, and verified that lexical decision times were much shorter for causal antecedent words than for state words. Potts, Keenan, and Golding (1988) applied a naming latency task after participants read two sentence excerpts. Depending on the text condition, the critical word could be a causal antecedent, a causal consequent, or a word not related with the previous excerpt. Time to name the word was faster in the causal antecedent condition, but did not differ between the causal consequent and not-related word condition.

Finally, a critical point concerns the online occurrence of global inferences. According to constructionist theory, received information can be linked to previous contents no longer in working memory, even if the text is locally coherent. As we will see later, a minimalist approach does not predict the occurrence of global inferences, as long as the text is locally coherent. In this respect, Singer (1993) conducted a study which compared fragments of texts that imply a global inference (global inference condition), with fragments with similar words, but not implying the global inference (control condition). For example, the first sentence of a text could be either: a) “Valerie left early for the birthday party” (global inference condition), or a’) “Valerie left the birthday party early” (control condition). Three sentences intervened between the first sentence and a test question. In this example the sentences were: “b. She checked the contents of her purse/ c. She backed out of the driveway/ d. She headed north on the freeway/ e. She exited at Antelope Drive/ f. She spent an hour shopping at the mall.” After all excerpts were presented, the test question involving the global inference was made (in this case: “Do birthday parties involve presents?”). Because three sentences intervened between the first “activating goal” sentence (a) and the sentence the “reactivating goal” sentence (f), the first sentence would no longer be available in working memory. In addition, no coherence break in the story is present in order to instigate a global inference, as a minimalist approach would propose. Thus, while a constructionist approach would predict that a global inference would link the two sentences, a minimalist account would not predict such an inference. Results showed that participants take significantly less time to answer the question in the goal inference

condition than in the control condition. Also, similar response times were observed in the goal inference condition and in an explicit condition, in which the goal was explicitly mentioned in the last sentence (i.e., f. She spent an hour shopping for a present at the mall).

These results, however, are open to alternative accounts involving deliberative retrieval strategies. More convincing data were reported by Suh and Trabasso (1993). In this case, an online recognition test of an initial “goal statement” was applied in different moments of a story. Illustrating with an example, in the hierarchical version of the story, Jimmy wants to buy a bike. He asks his mother to buy the bike, but she refuses. In the sequential version of the story, the mother agrees and buys the bike. The continuation of both stories includes actions that, in the hierarchical version, reflect a new attempt to buy the bike (e.g., Jimmy wanted to earn some money, Jimmy asked for a job at a nearby grocery store, etc). These actions should reinstate Jimmy’s initial goal of buying a bike in the hierarchical, but not in the sequential version. It was shown that participants take less time to correctly recognize the goal statement (i.e., Jimmy wanted to buy a bike) after sentences that reinstate the goal (e.g., Jimmy wanted to earn some money) in the hierarchical version, than after the same sentences in the sequential version.

Finally, in support of the fact that inferences about the emotional character of protagonists are made online, research by Gernsbacher and colleagues was cited. In this study, it was shown that participants take less time to name implied emotional state words than inappropriate emotions (Gernsbacher et al., 1992, Experiment 3).

Thus, previous studies suggest the online occurrence of referential, superordinate goals, causal antecedents, global theme and emotional inferences. In contrast, there are data that support the notion that causal consequences (Magliano et al., 1993; Potts et al., 1988), subordinate goals (Long & Golding, 1993; Long et al., 1992; Long et al., 1990), and instrumental inferences (McKoon & Ratcliff, 1981; Singer, 1979) do not occur online.

Essentially, according to the constructionist theory, the reader is actively engaged in an effort after meaning. Inferences are made in order to fulfill the reader’s goals, to continually explain why events occurred in the text, and to integrate information at both local and global levels. Recently, however, a minimalist framework

developed by McKoon and Ratcliff (1992) rejected the constructionist notion that inferences are generated in order to sustain the construction of a situational model. According with McKoon and Ratcliff, an analysis of the exiting literature on inference generation supports the view that readers are minimal inference encoders that only generate online (a) those inferences that are easily available, and (b) inferences that are necessary to establish local coherence between statements.

3.5. Minimalist Approach

The minimalist approach was developed in order to account for the class of inferences that are automatically, as opposed to strategically, encoded during text comprehension. McKoon and Ratcliff (1992) defined inference as “any piece of information that is not explicitly stated in the text” (p. 440). This broad definition was purposely used in order to stress the variety of existing inferences (i.e., from more simple connecting and bridging inferences, to more complex elaborative inferences).

A major claim of the minimalist approach is that, contrary to constructionist’s proposal, readers *do not* construct a “life-like” representation of the text. In contrast, only two types of automatic inferences are assumed to be mentally encoded: inferences that are needed for local text coherence and inferences based on easily available information. Local coherence is achieved, according to McKoon and Ratcliff (1992, p.441) by making sense of “those propositions of a text that are in working memory at the same time, in other words, propositions that are no farther apart in the text than one or two sentences.”

The information easily available may arise from existing general knowledge, or from previous text information. In addition, previous text information may be in working memory, or be easily retrievable from the text representation already in long term memory.

McKoon and Ratcliff (1992) emphasize the fact that a minimalist hypothesis does not reject the existence of strategic inferences. On the contrary, if readers adopt special strategies or have specific processing goals, then strategic inferences may be incorporated into the mental representation of the event. In fact, because readers generally pursue specific goals, elaborative inferences that fulfill those goals are likely

to be encoded. However, strategic inferences are assumed to be built upon a basic level representation that includes automatic inferences.

The distinction between automatic and strategic inferences made by the minimalist approach is consistent with other approaches within discourse literature. For example, Swinney and Osterhout (1990) distinguish between perceptive inferences and cognitive inferences. Perceptive inferences are seen as automatic, mandatory and independent of contextual conditions. By contrast, cognitive inferences are described as neither automatic nor mandatory. Cognitive inferences are under the cognitive control of the subject, and are constrained by word knowledge, plausibility, and pragmatic requirements. Because of that, cognitive inferences “cannot be immediately, or universally, generated” (p. 21). Swinney and Osterhout (1990) endorsed the view that the reader is equipped with only a very limited set of perceptive inferences (mostly, anaphoric inferences), while the majority of inferences studied in the literature should be classified as cognitive inferences.

In a detailed analysis, McKoon and Ratcliff (1992) presented several arguments in favor of a minimalist account. First, McKoon and Ratcliff argued that many previous studies that had explored the online nature of inferences have methodological problems (McKoon & Ratcliff, 1990b, 1992, 1995; see also Ratcliff & McKoon, 2008). A major difficulty in a lot of studies is that it is impossible to distinguish between inferences that were made during encoding from inferences made only during test. For example, evidence from cued-recall tasks could be due to the fact that an association between the cue and the retrieved sentence was established only during retrieval. Strong evidence for this was reported by Corbett and Doshier (1978). They showed, for example, that the word “hammer” was an effective retrieval cue for the sentence “the workman pounds the nail”. However, it was also demonstrated that “hammer” is equally effective when the sentence explicitly mentioned some *other* device (not a hammer) was used (i.e., “the workman pounded the nail with a rock”). Thus, the probe word was an effective retrieval cue, even in a condition when it was not encoded. Another methodological problem pointed out by McKoon and Ratcliff (1995) is that, in many previous studies, participants could have used explicit, elaborative or goal strategies in order to perform the task. Because this is a major difficulty in exploring the recurrence of certain types of

inferences, they stressed that automatic inferences can only be studied in the absence of any special processing goals.

As a second aspect in favor of the minimalist approach, McKoon and Ratcliff (1992) argued that findings showing the occurrence of certain inferences (e.g., referential and causal inferences), using more appropriate online measures, can be easily accommodated by the minimalist account. According to a minimalist approach, inferences necessary to establish local coherence are automatically encoded, which naturally includes both referential inferences and causal inferences based on information that is easily available, or necessary to guarantee local coherence to the text (see McKoon & Ratcliff, 1992, for several examples of studies showing the automatic nature of these inferences).

Third, McKoon and Ratcliff (1992) presented evidence showing that global inferences that link separated parts of the text are not usually generated during comprehension. In one study (McKoon & Ratcliff, 1992, Experiment 1), participants were presented with texts that were all locally coherent. Using a recognition probe task, it was shown that the level of activation of the general goal was the same in a control condition (in which the initial goal was no longer pursued), and in conditions in which the actor was pursuing the same goal by different means (i.e., “try again” condition and “substitution” condition). These results are in agreement with the minimalist prediction that, as long as texts are locally coherent, general goal inferences are not required. In opposition, results do not fit with a constructionist account. Constructionists would predict a higher level of goal activation in the “Try again” and “Substitution” conditions, because in these cases the protagonist is still trying to reach the goal.

Using a similar paradigm, McKoon and Ratcliff (1992, Experiment 2) presented further support for the fact that global inferences are not normally generated, unless the text presents locally incoherencies. It was further shown that global inferences are not activated when texts are locally coherent, even if there are global inconsistencies. In this study, global inconsistencies were introduced by presenting a new goal inconsistent with a goal stated in the introduction of the story (for example, going to a restaurant instead of going to a picnic). On the other hand, results showed that global inferences are activated when texts are locally incoherent, but can be made coherent by activating global information. Again, these results are in line with the minimalist view, according

to which global inferences should occur in the second case, but not in the first. However, results contradict a constructionist approach that would predict global inferences in both conditions (see Experiment 3 and 4, for further evidence that global inferences are only generated in the presence of local coherence breaks).

Finally, McKoon and Ratcliff (1992) reviewed data showing that certain elaborative inferences are *not* automatically generated. This evidence supports a minimalist view, but is in opposition with a constructionist account. According with the minimalist view, elaborative inferences should not occur automatically because they are not usually necessary to ensure local coherence. In opposition to the constructionist notion that readers construct real-life mental representations, many nonminimal inferences should occur.

For example, McKoon and Ratcliff showed that instrumental inferences do not occur automatically. They argued that, according to a constructionist account, the sentence "Mary stirred her coffee" would involve the inference "spoon". In fact, it is reasonable to assume that a real-life representation of the event would incorporate the instrument. However, different studies have failed to provide evidence for instrumental inferences. Doshier and Corbett (1982), for example, predicted an increase in interference when the instrument is presented as a test item in a Stroop task. However, no Stroop interference was found, suggesting that comprehension of the event did not entail an inference of the instrument. Notice that, according to a minimalist account, an inference about the instrument should occur, but only if the accessibility level is increased to a sufficiently high level (i.e., if it is easily available). This prediction was confirmed by McKoon and Ratcliff (1981, Experiment 1). In this study, the instrument was explicitly referred to in the first sentence of a text (e.g., hammer). The last sentence of the text was the instrument-implying sentence (e.g., "Then Bobby pounded the boards together with nails"). After the last sentence, a probe recognition task was applied with implied instruments as critical probes. In this case, a facilitation effect was obtained for the implied instrument, confirming the minimalist hypothesis that easily available inferences are likely to occur automatically.

Further evidence for this notion was provided in a study conducted by McKoon and Ratcliff (1986, Experiment 3), in which they explored the occurrence of predictive inferences. It was demonstrated that predictive inferences are mentally represented only

when those inferences are made available experimentally. In this study, predictive sentences like “the actress fell from the 14th story” (that would lead to the prediction “dead”) were used. A speeded probe recognition delayed test was applied, with the particularity that a prime was presented before the critical probe word (e.g., dead). The prime could be a neutral prime (i.e., ready) or a word from the sentence (i.e., actress). Decreases in accuracy were observed only when the probe word was preceded by a relevant prime (i.e., actress), and not when they were preceded by neutral primes. McKoon and Ratcliff interpreted this result as evidence that there are certain inferences that are only *minimally encoded* (McKoon & Ratcliff, 1986, 1989a, 1990a). Again, this pattern of results contradicts the notion that highly predictable events are clearly encoded, as would be assumed by a constructionist account.

3.5.1. Minimalist Approach Development

Despite the existing evidence supporting a minimalist view of the inference generation process, there are some aspects in the model that needed further clarification. For example, one class of inferences that are assumed to be automatically generated are those that are easily available. McKoon and Ratcliff (1995) themselves pointed out that it is not totally clear what “easily available” might mean. That is, the mechanisms by which general knowledge and previous text information are made available from long term memory were not specified by the authors. In order to clarify this point, McKoon and Ratcliff proposed a *memory-based text processing* mechanism (McKoon, Gerrig, & Greene, 1996). This mechanism is well captured by a *resonance* metaphor (see also O’Brien & Myers, 1999; Myers & O’Brien, 1998). Specifically, it is assumed that, by a resonance mechanism, each element of the text (i.e., each word and proposition) cues information from all memory in a passive, quick, parallel way (McKoon et al., 1996; McKoon & Ratcliff, 1995; Ratcliff & McKoon, 2008). Contents currently in working memory function as signals to the entire contents of long term memory. Long term memory concepts highly associated with working memory contents, or long term memory concepts frequently referenced in the text, are more likely to be reactivated, and to become “easily available” for inference generation. This proposal imports principles of global memory models (Gillund & Shiffrin 1984; Hintzman 1986, 1988; Murdock 1982).

The possibility that contents from long term memory are activated without consuming resources has also been integrated in other prominent text comprehension models. For example, Ericsson and Kintsch (1995) added principles to the construction-integration model in order to incorporate this notion. The highly restricted capacity of working memory was seen as not being adjusted to an efficient discourse processing. In order to explain how the mind circumvents this limitation, Ericsson and Kintsch (1995) proposed that information from working memory can cue long term memory contents in an effortless manner. In this way, information from long term memory may become available, allowing an extension of the capacity of working memory and forming what the authors have called a “long term working memory” (Ericsson & Kintsch, 1995).

McKoon et al. (1996; see also Greene, Gerrig, McKoon, & Ratcliff, 1994) presented data that corroborate the passive activation of long term memory contents. They presented participants with a story of two actors making reference to a third person. Results showed that when the two actors are reunited, the third person is highly accessible, despite not being present, and not being recently mentioned. These findings are consistent with a memory-based processing mechanism, according to which the two actors serve as cues to evoke associated concepts.

3.5.2. The Gradual view of Inferences

The minimalist approach represents a challenging, and innovative approach, to the problem of inference generation. A minimalist view clearly contradicts the notion that readers construct mental models about the text situation. Instead, readers are seen as “minimally encoders” who, in the absence of special goals, only represent automatically those inferences that are essential for local coherence, and inferences that are based on highly available information. Such view is convergent with other researchers’ proposal that readers generally construct “good-enough” representations of the text (Ferreira, Bailey, & Ferraro, 2002; Klin, Guzmán, Weingartner, & Ralano, 2006; Sanford & Graesser, 2006).

The minimalist notion elegantly explains why readers do not face themselves with “explosions of inferences” problems during text processing. It also has the advantage of depicting the reader as an efficient cognitive resource manager who does

not expend unnecessary effort in processing a text, unless some special motivation is in place.

Further, the fact that the model had incorporated a “resonance mechanism” to explain how long term memory contents become available enhanced the theoretical sophistication of the minimalist approach (McKoon et al., 1996; McKoon & Ratcliff, 1995; Ratcliff & McKoon, 2008). This development has two major advantages to the study of inferential processes. First, it incorporates cognitive mechanisms that explain how background knowledge and received information interplay in order to sustain inference generation. This has been considered one of the greatest “mysteries” within discourse comprehension literature (e.g., Schank & Abelson, 1977). Moreover, this account naturally integrates inference generation within the larger function of the human mind. Inference processes are seen as working in line with general principles of memory and knowledge activation. Second, according with a resonance explanation, long term memory activation is seen as gradual. This allows the operationalization of inferences in terms of strength, instead of an all or none view of inference generation. Thus, instead of simply asking whether an inference had occurred or not, it should also be explored the degree in which an inference had been encoded. In fact, there is evidence showing that inferences may vary in terms of encoding strength; with some inferences being only minimally encoded (Klin et al., 2006, McKoon & Ratcliff, 1986, 1989a). Importantly, this gradual perspective of inference encoding highlights the fact that the likelihood of an inference to become visible during retrieval is dependent on the specific conditions of retrieval (McKoon & Ratcliff, 1986, 1989a). Given the same encoding conditions, inferences might be visible under some retrieval settings, but look as if they had not occurred under other retrieval contexts. As McKoon and Ratcliff (1989a) notice, if an inference is strong enough it will be visible under a number of different retrieval contexts, while if it was only minimally or partially encoded, then it must only be revealed under favorable conditions. McKoon and Ratcliff (1986) provided supporting for this view. Using a delayed probe recognition method, they found evidence for predictive inferences, only when the probe word was preceded by a prime word of the previous text, but not when the probe word was preceded by a neutral prime (i.e., “ready”). Based on this data, McKoon and Ratcliff (1986) concluded that predictive inferences are minimally encoded into memory and became visible only

when additional information is provided that matches the encoding information (see also McKoon & Ratcliff, 1988; 1990a).

Researchers within text comprehension have been faced with the necessity of incorporating a gradual view of the inference generation process. An all-or-none view seems to be limited to explain the flexibility of the inferential activity underlying comprehension. In addition, such a view is also inconsistent with general principles of memory and cognition that are endorsed by several prominent cognitive models (Gillund & Shiffrin 1984; Hintzman 1986, 1988; Murdock 1982). Within discourse research, besides McKoon and Ratcliff's proposal (1986, 1989a), Kintsch's model also easily integrates the notion that some inferences might be inferred only to some degree (Kintsch, 1988, 1998). According to the construction integration model (Kintsch, 1988; 1998), during an initial *construction* phase, the text is represented by a network of activated concepts. The network results both from explicit text information and from general knowledge (that is assumed to be activated in parallel, as a result of level of association with explicit concepts). Initially, this representation includes several plausible concepts, including concepts that were "incorrectly" activated, as a result of parallel activation mechanisms. In a second phase, an *integration* process is used in order to eliminate incorrect meanings and select those elements that are more contextually appropriate. This process functions in continuous recycles of activation during text processing. As a result, information that is only weakly represented in the network will vanish, if it is no longer referenced. In contrast, if many subsequent concepts are associated with a concept from the network, posterior recycles of activation will capture this information, and the concept will become more strongly represented. Importantly, inferences that are only weakly associated (i.e., potential inferences) can become highly active over time, as long as they are referenced by multiple sources in the text. Thus, during the process of integration, those inferences that are more adequate to the context are naturally selected, while inferences that are contextually inappropriate are deactivated.

A more gradual view of inference generation has also been endorsed by constructionist researchers. For example, Graesser and collaborators state that "some inferences may slowly emerge as text is received rather than discretely popping in when a particular statement is comprehended" (Graesser et al., 1997, p.184), and further

mention that “the degree to which an inference is encoded might be strengthened or attenuated as more information is received” (Graesser et al., 1994, p.376).

In sum, the minimalist framework proposes a continuum view of inference encoding. Such a view has also been endorsed by other researchers within the field. As a consequence, issues about strength of inference encoding, and about the variables that govern the inference continuum are more and more considered to be critical issues within discourse processing research.

3.6. General Implications for STI Research

The text comprehension literature explores a number of questions that are relevant to the study of the trait inference process, including the conditions necessary for inference occurrence, the clarification of the mechanisms by which background knowledge and text information interplay, and the way inferences can be represented (in some cases, *minimally* represented).

There are several aspects that might be important to be considered within the STI domain. First, text literature growth was made largely through the development of theoretical models that allow predictions about when inferences are likely to occur. For instance, the minimalist approach predicts that elaborative inferences will occur when they are highly available, or when the reader is faced with local incoherencies. STI research, in contrast, has been developed largely in the absence of theoretical models that allow specific prediction about inference occurrence conditions (but see Reed & Miller, 2005).

Second, text comprehension researchers tend to conceptualize the inferential process as integrated within the larger functioning of the comprehension cognitive system. As a consequence, the explanation of the mechanism underlying inference generation has always been linked to more general assumptions about memory functioning and knowledge organization. Exploring whether trait inferences function works in agreement with more general cognitive principles could be important to an accurate characterization of the trait inference process.

Third, discourse researchers’ formulations about inferential processes tend to be sensitive to the “flexibility that must be present (...) in human information-processing in general” (Ratcliff & McKoon, 1997, p. 41). For example, Graesser et al. (1994) posit

that inferences are guided by coherence principles, and point out that the online-offline distinction is a continuum that may vary depending on various factors, including previous knowledge and experimental tasks. McKoon and Ratcliff (1986, 1989a) describe inferences as being made in order to guarantee local coherence, or based on information more easily available. Kintsch' model (1988, 1998) assumes that the system naturally selects those inferences that are more appropriate to the context. Thus, inference encoding is not conceptualized as a hermetic automatic process, but as working in line with contextual and coherence requirements. In the STI domain, studies are needed in order to explore the flexible nature of the spontaneous trait inference process. Such data might reveal important venues in terms of new theoretical formulations.

Finally, a gradual view of inference encoding was proposed by minimalist authors. Some inferences might be only weakly encoded, while others are strongly encoded. While the former inferences are visible only under a restricted number of retrieval conditions, the latter are generally detectable, and have stronger consequences for future processing. Strong inferences are assumed to be mentally represented as part of the event occurrence. In contrast with a continuum view of inference encoding, STI research has examined inferences occurrence according to an all or none view. In addition, STI methods have neglected the importance of exploring whether inferences are mentally represented, and thus fail to understand not only the degree to which trait inferences are encoded, but also the consequences of trait inference activation for future information processing.

3.7. Consequences of a Minimalist Approach to a Flexible View of STIs

In addition to the general implications of text comprehension research to STIs previously outlined, the minimalist framework developed by McKoon and Ratcliff (1992) allow the generation of specific predictions about spontaneous trait inference occurrence.

There are a number of different reasons why the minimalist framework is especially suitable for making predictions about the occurrence of STIs. First, the minimalist approach is the one with most empirical support, and is the framework that has been tested under more precise and detailed methodological conditions. Second, the

minimalist proposal is a processing model that defines the conditions that tend to promote inference encoding, and not a taxonomic model that identifies the classes of inferences that occur online. Models that propose taxonomic distinctions, as the constructionist theory (Graesser et al., 1994), might easily be accused of arbitrariness. Third, the minimalist approach is the only one that explicitly focused on the automaticity of inference encoding and thus, is the one with greater direct implications for STI research, a field that has been concerned with exploring whether trait inferences occur in the absence of conscious intentions.

Grounded on the minimalist framework, we may establish three basic principles underlying spontaneous trait inference occurrence that support a more flexible view of the STI process:

1. Spontaneous trait inferences are more likely when easily available. This principle follows from the minimalist assumption that inferences tend to occur automatically when they are based on information from long term memory that becomes easily available from the text. This would suggest that, under conditions in which traits are highly implied, and are thus easily activated from the text, spontaneous trait inferences are likely to be observed. In contrast, under conditions in which trait concepts are not so highly implied by the text, spontaneous trait inferences would be less likely to occur;

2. Spontaneous trait inferences work in line with local coherence requirements. This principle is based on the minimalist assumption that automatic inferences occur in order to guarantee text local coherence. We specify the assumption by establishing that spontaneous trait inferences can be *facilitated* if they guarantee greater local coherence, but can also be *inhibited* if they lead to local incoherencies;

3. Spontaneous trait inferences vary in a continuum of strength. This principle incorporates the notion that inference encoding should be conceptualized in terms of degree, and not in an all or none way. From this principle, it would be predicted that the encoding strength of trait inferences varies with encoding conditions. In the stronger level of inference encoding, inferences are assumed to be accessed in subsequent tasks and to influence future information processing.

In the present work, we propose that the functioning of the STI process works in line with these three processing principles. These principles imply a more contextual

dependent-view of the STI process. Based on the outlined processing principles, specific hypothesis can be generated about trait inference occurrence. These hypotheses will be detailed in the next section, and will be tested in a set of empirical studies.

3.8. Summary of the Chapter

The text comprehension literature has exhibited enormous progress throughout the years. One of the major challenges of this literature has been on resolving the tension between the absolute necessity of the comprehension system to be based on inferential processes, on one hand, and the extraordinary ability of human cognitive system to only generate those inferences that are most adequate in a specific context, on the other hand.

The script framework (Shank & Abelson, 1977) was developed in order to circumvent the explosion of inferences problem (Rieger, 1975). The concept of script is highly valuable in order to clarify the nature of social knowledge structures, and their influence within the comprehension process. However, the script concept suffers from a problem of lack of specificity that makes it difficult to make exact predictions about inference generation. An exception is provided by the knowledge access network model (Sharkey, 1986; Sharkey & Mitchell, 1985). According to this model, we would predict that script-central actions are the more likely to be inferred, given the activation of a script. Even so, ambiguities exist about what types of scripts exist, the circumstances in which central-script actions are inferred, and the automatic nature of those inferences.

The inferences that are usually generated during the comprehension process were more extensively explored by other researchers. This field has been marked by the opposition between the constructionist theory (Graesser et al., 1994) and the minimalist approach (McKoon & Ratcliff, 1992, 1995). Constructionists maintain that the inferences that are generated online are those that support the creation of a meaningful and coherent situation model of the event. A constructionist approach to inferential generation was, however, criticized by McKoon and Ratcliff (1992, 1995). According to these authors, a complete mental model about the situation is not created, unless the reader has specific processing goals. McKoon and Radcliff also assert that the criteria used by constructionist authors to define the class of inferences that occur online seem arbitrary. For example, despite the constructionist theory (Graesser et al., 1994)

postulate that inferences are not inferred online, McKoon and Ratcliff (1992) claim that it is reasonable to assume that the sentence “stirring coffee” would lead to the inference “spoon”, if it is presumed that a situation model of the situation is created.

The minimalist proposal (McKoon & Ratcliff, 1992, 1995) argues for a restrictive view of the inference process generation. According to this view, readers only generate automatically inferences to guarantee local coherence, and those inferences that are easily available from knowledge activation. The information that is easily available was operationalized by a memory-based text processing (McKoon et al., 1996), by which text information activates background knowledge in a passive, fast, and unrestricted way.

It should be noticed that considerable debate still exists between the constructionist and minimalist approaches concerning some important questions. For example, the question about whether individuals make global inferences is controversial, with some evidence favoring the occurrence of global inferences (Suh & Trabasso, 1993), and other studies showing that individuals only produce global inferences in case of local coherence breaks (McKoon & Ratcliff, 1992).

However, existing empirical data tend to support a minimalist view of inference encoding. While evidence is straightforward in proving the automatic occurrence of inferences that are necessary for local coherence, studies tend to show that elaborative inferences are not likely to be automatically encoded.

Based on the minimalist account, we generate three assumptions about spontaneous trait inferences occurrence conditions. According with these principles, STIs will be more likely to occur when they are highly implied by text information, STIs work in line with local coherence demands, and STIs vary in terms of encoding degree. Incorporating the basic principles underlying the minimalist framework into the study of STIs supports a more malleable view of the STI process.

CHAPTER IV

PRESENT PROPOSAL:

A Flexible View of Spontaneous Trait Inferences

4.1. State of the art

The way we infer the personality traits that characterize other people based on the observation of their external behavior has been a prominent theme of research within social psychology. Seen as an essential part of the process of “making sense of people” (Kunda, 1999), trait inferences were differently conceptualized throughout the years. In the first two chapters of the present work, attribution research (Chapter I) and spontaneous trait inferences literature (Chapter II) were reviewed. The examination of these two domains provides a clear picture about how trait inferences have been defined and investigated within the domain of social psychology.

Initial attributional researchers conceived the social perceiver fundamentally as a rational agent. As a consequence, trait inferences were viewed as occurring only in a late stage within the person perception process, being conditional upon a deliberative analysis of the causes that had promoted the behavior occurrence (Jones & Davis, 1965, Kelley, 1967). Thus, inferring a personality trait about an actor was seen as a consequence of a causal attributional analysis of his behavior.

Several authors have pointed out, however, that trait inferences and causal attributions were better defined as separate cognitive processes (Hamilton, 1988, 1998; Ericsson & Krull, 1999; Krull, 2001). The basic argument was that trait inferences are sometimes drawn (for example, “he showed a friendly behavior, so he must be a friendly person”), regardless of considerations about the causes of the behavior (e.g., “why did he act in a friendly way?”). Empirical data gave support to this distinction (Bassili, 1989; Erickson & Krull, 1999; Johnson et al., 1984; Smith & Miller, 1983; see also Hilton, Smith & Kim, 1995; Reeder & Spores, 1983). For example, it was shown that the time to make a dispositional trait inference is much quicker than the time to make a causal attributional analysis (Smith & Miller, 1983). This perspective is incompatible with the previous view according to which trait inferences only take place after an analysis of the causes of the behavior. The implication in terms of research was that trait inferences could be investigated without being necessarily incorporated within an attributional framework, as had previously been the case. Therefore, conditions were met for the development of an independent field of research that became known as the *spontaneous trait inference literature*.

The main goal of this emergent literature was to explore whether trait inferences could be drawn without any intention or conscious awareness on the part of the perceiver. This was a radically different way of approaching the trait inference process. While attributional theorists viewed trait inferences as occurring only after a careful causal analysis, now trait inferences were seen as indistinguishable from the process of behavioral comprehension itself, and were portrayed as a largely unavoidable process. By showing that even when participants are only instructed to memorize lists of behaviors, providing implied traits as cues helps behavioral retrieval, Winter and Uleman's (1984) initial findings were interpreted as evidence that trait inferences occur spontaneously during behavior encoding. From then on, employing a variety of methods and procedures, STI research was mainly oriented to test the automaticity of the trait inference process (e.g., Todorov & Uleman, 2003; Winter et al., 1985).

4.2. The Present Proposal: A Flexible view of the STI process

Contrary to an unconditional perspective of the STI process, in the present proposal we suggest a flexible view of the spontaneous trait inference process. We based our proposal on the analysis of the spontaneous trait inference literature (Chapter II), as well as on the examination of how inference processes have been investigated within the text comprehension domain (Chapter III). Next, we will briefly outline the major arguments of these literatures in favor of a flexible view of the STI process.

4.2.1. Evidence from STI research

Spontaneous trait inference research has provided a large amount of empirical data that help to characterize the STI process. Although the dominant emphasis has been on exploring whether the trait inference process fulfills all the criteria for automaticity (Todorov & Uleman, 2003; Uleman et al., 1985; Uleman et al., 1996; Winter & Uleman, 1984), some of the data give support to a more flexible view of the STI process.

It was shown, for instance, that the magnitude of the STI process may vary depending on the study conditions. Uleman and Moskowitz (1994) demonstrated that STIs are dependent on the type of processing goal. They have shown that STIs are weaker, the less the processing goal of the perceiver involves a semantic analysis of the

behavior. Moreover, no evidence of STIs was found under certain shallow processing conditions (i.e., graphemic goal conditions, Uleman & Moskowitz, 1994). These results led to characterizing STIs as conditional upon a meaningful processing of behaviors (Bargh, 1989, 1992). In addition, it was also shown that STIs are undermined when high load tasks are performed concurrently (Uleman et al., 1992), contradicting the view that STIs are entirely independent of the availability of cognitive resources.

Besides this data showing that STIs are undermined under certain study conditions, a number of different studies have consistently reported evidence for the spontaneity of other types of inferences, including gist (Winter et al., 1985), situational (e.g., Lupfer et al., 1995), goal (Hassin et al., 2005), predictive (McKoon & Ratcliff, 1986), and emotional (Gernsbacher et al., 1992) inferences. If this evidence is taken as showing that all these inferences are automatically drawn under all circumstances, then we will have a very ineffective picture about how our comprehension system works. It seems more effective for a comprehension system to assume that the spontaneous occurrence of the different types of inferences depends on specific processing conditions. It is likely that the more the behavior and the context imply a certain inference (whether a trait, gist, goal, situational, predictive or emotional inference) the more likely it will be that the inference will be drawn in a spontaneous way.

In addition to the points outlined, one of the strongest lines of evidence in favor of a flexible view of the STI process comes from a study conducted by Wigboldus and collaborators (Wigboldus et al., 2003). In this study, it was shown that associating a stereotype with the actor of the behavior affects the magnitude of the STI process. Given the crucial nature of this evidence in favor of a more flexible view of the STI process, as well as its relevance to our own program of research, this study will be described in detail.

4.2.1.1. STIs are affected by stereotypes: A strong argument in favor of the STI flexibility

Despite the large amount of empirical evidence, STIs have been investigated in what Wigboldus and colleagues have called a “social vacuum” (Wigboldus et al., 2003, p.471). In fact, studies within the STI domain typically present actors as abstract entities about whom participants have no additional knowledge besides the behavior presented.

No information is provided, for instance, about the social categories of the actors. However, as Wigboldus and colleagues noted, this approach does not reflect the social reality outside the laboratory. People usually have knowledge about the social groups to which the actors belong, at least about the more perceptively salient social categories, such as their age, gender, and racial ethnicity. In the few STI studies in which information about a social group was provided, it was purposely pre-tested in order to be irrelevant for the interpretation of the behavior (Winter & Uleman, 1984). This methodological aspect may have led to an erroneous picture about the conditions for the occurrence of the spontaneous trait inference process. Crucial to our point, it may have led to a possible overestimation of the pervasiveness of the process.

In order to circumvent this aspect, Wigboldus et al. (2003, see also Wigboldus et al., 2004) associated a category label with the actor that could be either stereotype-consistent or stereotype-inconsistent with the behavior. For example, the behavior “wins the science quiz” (i.e., a behavior that would lead to the spontaneous inference “Smart”, under the usual conditions in which STI were previously studied) was presented with either a stereotype-consistent label (i.e., professor) or with a stereotype-inconsistent label (i.e. garbage man). The recognition probe paradigm was applied in these experiments. A series of behaviors was presented to participants and immediately after each behavior a probe word appeared on the screen. On critical trials the probe word was the trait implied by the behavior, not included in the previous sentence. Participant’s task was to indicate whether the probe word was part of the previous sentence, as fast and accurately as possible.

Results from five different experiments (Experiments 1-5) consistently showed that participants take significantly less time to indicate that the trait was not included in the behavioral sentence when the stereotype is inconsistent with the behavior than when the stereotype is consistent with the behavior. Thus, spontaneous trait inferences are less likely to occur when behaviors are inconsistent with the stereotype associated with the actor. The same pattern of results was observed when the stereotypic label was subliminally presented (Experiments 3 and 4).

In order to rule out an alternative explanation of the results based on mere stereotype activation effects, in one of the experiments the same stereotypic labels were paired with neutral behaviors (e.g., “The professor cycles through the street”)

(Experiment 2). Results showed that participants took significantly less time to reject the trait (e.g., smart) when the stereotypic label (e.g., professor) was presented with a neutral behavior than when the same stereotypic label was presented with a consistent behavior. Thus, the previous pattern of results cannot be explained by mere stereotype activation effects. If that was the case, the results should be the same regardless of whether the behavior was stereotypically neutral or consistent. Instead, results seem to be due to a specific interaction between stereotype activation and spontaneous trait inferences made from behaviors.

Wigboldus and colleagues were also interested in understanding whether the observed results were due to facilitation of the trait inference process in the stereotype-consistent condition or to an inhibition of the trait inference process in the stereotype-inconsistent condition. In order to clarify this issue, a baseline condition was included (Experiments 3-5), in which the same behaviors were presented with a neutral category label (e.g., The human wins the science quiz). Results showed that participants took less time to reject the trait probe in the stereotype-inconsistent condition than in both the stereotype-consistent and stereotype-neutral conditions. However, response times didn't differ between stereotype consistent and stereotype-neutral conditions. Thus, the results seem to be explainable by the inhibition of the trait inference process in the stereotype-inconsistent condition and not by the facilitation of the trait inference process in the stereotype-consistent condition.

Wigboldus and collaborators explained the obtained pattern of results based on trait accessibility effects of stereotype activation. According with this account, the activation of a stereotype leads to the temporary accessibility of stereotype-consistent traits and to the temporary inhibition of stereotype-inconsistent traits (see Dijksterhuis & van Knippenberg, 1996). While the greater accessibility of consistent traits does not interfere with trait encoding of stereotype-consistent behaviors, the inhibition of inconsistent traits make encoding of inconsistent behaviors in trait terms more difficult.

By demonstrating that STIs are inhibited by the activation of an inconsistent stereotype, results from Wigboldus et al. (2003) clearly show that spontaneous trait inferences should not be regarded as an unconditional automatic process. Even behaviors that are pre-tested to clearly imply a personality trait *do not* lead to spontaneous trait inferences when performed by certain social actors. Since in real life

at least some of the actor's social categories are known, these results indicate that previous STI research might have overestimated the prominence of the process.

4.2.2. Evidence from Text Comprehension research

Evidence in favor of a flexible view of the spontaneous trait inference process comes also from an examination of the text comprehension literature. Despite the fact that spontaneous trait inferences research and text comprehension research address similar research questions, conclusions reached by the two fields have been quite different. While STI research is directed toward testing the generality of the trait inference process, within the text comprehension domain the tendency is to endorse a minimalist view of the inferential activity (e.g., McKoon & Ratcliff, 1992).

According with a minimalist view, automatic inferences only occur when necessary for establishing local coherence of the text, or when are easily available. In fact, even constructionist approaches (e.g., Graesser et al., 1994) in which higher levels of on-line inferential activity are admitted, trait inferences are viewed as an elaborative type of inference that do not occur on-line, unless specific processing goals are present.

Based on the minimalist approach of automatic inferences, we generated three principles underlying STI occurrence. According with these principles: (a) STIs are more likely to occur when they are easily available, (b) STIs work in line with local coherence requirements and, (c) STIs vary in a continuum of strength. These principles suggest a more flexible view of the spontaneous trait inference process. In the present research project, studies were developed in order to test each one of these general principles.

4.3. First general principle: STIs are more likely when easily available

According with this principle, trait inferences occur spontaneously as long as text information makes the traits highly available. For example, when behaviors result in a very high activation level of a specific personality trait, the trait will be spontaneously inferred. Based on this principle, we predicted that, while for behaviors that highly imply a personality trait STIs will be generally observed, for behaviors that do not unequivocally illustrate a specific personality trait evidence for the STIs will be weaker. The fact that previous STI research had used behaviors pre-tested to be highly

illustrative of a specific personality trait, and these behaviors usually being presented in the absence of any context, may explain why evidence for spontaneous occurrence of trait inferences has been typically obtained.

One way of testing the assumption that STIs are more likely when the corresponding traits are easily available is by demonstrating that they are less likely to occur when traits do not become so easily available from sentence information. In order to test this prediction, in our first experiment (Experiment 1) we applied the probe recognition paradigm with ambiguous behaviors. Ambiguous behaviors were pre-tested so they could be interpretable in terms of two different personality traits. In addition, in this experiment the same ambiguous behaviors could be presented either with a neutral social category, with a social category that was consistent with one of the implied traits, or with a social category that was consistent with the other implied trait. According to our principle according to which trait inferences will occur spontaneously as long as they are easily available, we predicted that when a social category is provided in order to disambiguate the interpretation of the behavior, evidence for STIs will be obtained. By contrast, when no category is provided that helps to disambiguate the trait interpretation of the behavior, evidence for STIs will be weaker.

4.4. Second general principle: STIs work in line with local coherence requirements

According with our second principle, STIs occurrence is guided by local coherence principles. Local coherence designates the integration of adjacent or of a small number of sentences (Graesser et al., 1994). Local coherence is regarded by many discourse researchers as a basic need underlying text comprehension (Graesser et al., 1994; Kintsch, 1988; 1998; Lorch & O'Brien, 1995; McNamara & Kintsch, 1996; Tapiero & Kinstch, 2007). In line with this view, comprehension is described by some authors as the construction of a coherent representation of the text (Gernsbacher, 1990). The minimalist approach (e.g., McKoon & Ratcliff, 1992) specifies that automatic inferences occur when they are needed to establish text local coherence. Based on this assumption, we predicted that STIs will be facilitated when they would provide greater local coherence to the text, and will be inhibited when they would provoke text local incoherencies.

Local coherence is a key concept in the text comprehension literature. Standards of local coherence are viewed as playing a determinant role in the inferential activity (Kintsch, 1988; McKoon & Ratcliff, 1992; Van den Broek, Risdén, & Husebye-Hartmann, 1995). However, because in our studies we are mainly concerned with coherence of information *about the actor*, the notion of local coherence can be seen as equivalent to the notion of consistency, as it is usually defined in the impression formation literature. Therefore, in the present experiments, talking about local coherence or incoherence is equivalent to talk about consistency or inconsistency of information about the actor, respectively. Thus, we will adopt the terms consistency/inconsistency (terms more familiar in the impression formation domain) and coherence/incoherence (terms more familiar in the text comprehension literature) interchangeably in the present studies. Accordingly with this nomenclature, we generally predict that STIs will be facilitated when they are consistent with previous information about the actor (i.e., provide greater local coherence), and STIs will be inhibited when they are inconsistent with previous information about the actor (i.e., contribute to local incoherencies).

In order to test whether STIs work in line with text local coherence requirements, a set of different experiments was developed (Experiments 2-6). While in some conditions trait inferences resulted in greater local coherence of the text, in other conditions trait inferences increased text local incoherence. We predicted that trait inferences would be facilitated in the former case, and will be inhibited in the later case.

There are different ways by which text local coherence can be manipulated. The manipulation used by Wigboldus and colleagues (Wigboldus et al., 2003), for example, may be viewed as an information coherence manipulation. When the stereotype is consistent with the behavior, the trait inference process is in line with a coherent picture of the event. By contrast, when the stereotype is inconsistent with the behavior, the trait inference accentuates the incoherence of the event. In line with this reasoning, Wigboldus et al. (2003) showed that STIs occur in the first case, but are inhibited in the later.

In our experiments, local coherence was manipulated by a different means. In some of them (Experiment 2-4), situational continuations were added to the behaviors. In addition, behaviors were presented with a social category that could be either

consistent or inconsistent with the behavior. We predicted that both the occurrence of trait spontaneous inferences and of situational spontaneous inferences (see Ham & Vonk, 2003; Lupfer et al., 1995) would be dependent on guaranteeing text local coherence. Specifically, we predicted that when the social category of the actor is inconsistent with the behavior, spontaneous trait inferences will be less likely to occur than when the social category of the actor is consistent with the behavior. In contrast, spontaneous situational inferences will be more likely to occur when the behavior is inconsistent with the social category of the actor than when the behavior is consistent with the social category of the actor.

In another set of experiments (Experiments 5 and 6), local coherence was manipulated by presenting pairs of behaviors about the same actor. In this case, a target-behavior was preceded either by a behavior that implied the same personality trait or by a behavior that implied the opposite personality trait. Again, our general hypothesis was that the pattern of spontaneous trait inferences will work in line with coherence principles. Specifically, when the target behavior is preceded by a consistent behavior, spontaneous trait inferences will be more likely than when the same behavior is preceded by an inconsistent behavior. In the first case, the trait inference is in agreement with a coherent representation of the information received while in the second case the trait inference represents an incoherent way of interpreting the behavior given the previous behavior of the actor.

4.5. Third General Principle: Spontaneous trait inferences vary in a continuum of strength

This principle states that spontaneous trait inferences vary in terms of encoding strength. The way STI research and text comprehension literatures regard this question is largely different. While STI research usually investigates trait inference occurrence in an all-or-none way, text comprehension literature endorses more and more a gradual view of inference encoding (McKoon & Ratcliff, 1986, 1989a).

According with this third principle, trait inferences can be encoded in a weaker or stronger way. Weaker trait inferences are visible only under a restricted number of retrieval conditions, while stronger trait inferences are more generally detectable. Moreover, we postulate two characteristics that differentiate weaker from stronger trait

inferences. First, strong inferences may be deliberately used in subsequent tasks. Second, strong trait inferences work as expectations in future processing of information about the same actor.

The encoding strength of spontaneous trait inferences was explored in two experiments (Experiments 7 and 8). In Experiment 7, we compared whether participants under memory and impression formation instructions differ in the extent to which they use previously inferred personality traits as retrieval cues for behaviors recall. In Experiment 8, we explored whether trait inferences under memory and impression formation instructions work as expectations that guide subsequent processing of information about the same actor. It was reasoned that if inferred traits become mentally represented as dispositional characteristics of actors, traits should act as expectations and guide subsequent information processing. As a consequence, the time participants will take to read subsequent inconsistent behaviors will be longer than the time they will take to read subsequent consistent behaviors performed by the same actor. The reason is that inconsistent behaviors are more difficult to integrate with the expectation created by the inferred personality trait.

PART II

EMPIRICAL STUDIES

SECTION I

**SPONTANEOUS TRAIT INFERENCES ARE MORE LIKELY
WHEN EASILY AVAILABLE**

EXPERIMENT 1

Probe Recognition with Ambiguous Behaviors

According with the principle of local coherence, trait inferences occur spontaneously when they are highly activated by text information. For instance, if a behavior is highly illustrative of a personality trait, the trait will be spontaneously inferred. By contrast, if the behavior does not imply the trait in such a strong way, STIs should be less likely.

An important issue within STI research is the way behaviors are usually selected. Behavioral material is chosen to consensually and unambiguously imply a specific personality trait. It is reasonable to question whether this methodology does not involve a certain degree of circularity. Based on pre-tests, behaviors most likely to consistently elicit certain personality traits are chosen. These behaviors are then presented to participants without any contextual information, but this time without asking them to make an explicit trait inference. Data from these studies is then taken as evidence that trait inferences occur spontaneously. We may question whether this evidence should be taken as an indication that people spontaneously infer traits from behaviors in a generalized way, or whether spontaneous trait inferences occur only for behaviors that are already known to strongly elicit dispositional trait inferences.

In our view, because previous studies had made use of highly trait-implicating behaviors, the flexible nature of STIs was difficult to demonstrate. However, we do not claim that previous evidence does not prove that trait inferences occur spontaneously. On the contrary, we simply contend that conditions from previous studies were the most favorable for STIs occurrence. And we certainly agree that this is the most adequate approach to an initial study of spontaneous trait inferences occurrence. However, extending this research to more complex study conditions is crucial if we want to better understand the nature and generality of the process.

In our first experiment, we applied the recognition probe paradigm (McKoon & Ratcliff, 1980, 1986; Uleman et al., 1996), but instead of presenting behaviors highly illustrative of one personality trait, ambiguous behaviors were presented. The same ambiguous behavior could be interpretable by two different personality traits. For example, the behavior “His contacts with other people are rather limited” may be interpret as implying the trait “shy” or the trait “independent”. Based on pre-tests, ambiguous behaviors were selected so that each of them implied two personality traits. Importantly, the degree to which the two underlying traits were implied by each

ambiguous behavior was considerably weaker than the degree to which traits are usually implied by non-ambiguous behaviors in STI research (see *Stimulus Material* section for details about the behavioral material).

In line with Wigboldus et al. (2003), in the present experiment a relevant social category was associated with the actors of behaviors. The social categories served to disambiguate the trait interpretation of the behavior. Depending on the condition, the category label could favor one or the other trait interpretation. For example, the behavior “His contacts with other people are rather limited” was presented with a label consistent with one of the trait interpretations (e.g., the label “librarian” favored the trait inference “shy”); a label consistent with the alternative trait interpretation (e.g., the label *businessman* favored the trait inference “independent”), or with a neutral category label (e.g., Person).

In this experiment, the typical recognition probe paradigm procedure was followed (McKoon & Ratcliff, 1980, 1986; Uleman et al., 1996). Participants were presented with a series of ambiguous behaviors. Immediately after each behavior, a probe word appeared on the screen. Participants’ task was to indicate as fast and as accurately as possible whether the probe word was included in the preceding sentence. On critical trials, the probe words were the two personality traits implied by each ambiguous behavior. These trait words were never included in the sentences. Depending on the trial, each behavior was combined with a different social category and tested with one of the two trait probes.

Following the principle that spontaneous trait inferences are more likely to occur when traits are easily inferable by the sentences, it was predicted that STIs would be more likely when the category label helps to disambiguate the trait interpretation of the behavior than when the behavior remains ambiguous by being presented with a neutral category label (e.g., He). In other words, according to our hypotheses participants would have greater difficulty to indicate that the trait probe was not included in the sentence on “match trials” (i.e., when the category label favored the trait interpretation tested by the probe word – i.e., “librarian” tested with the trait probe “shy” and “businessman” tested with the trait probe “independent”) than on “neutral trials” (i.e., when the category label was neutral in relation to any of the two trait interpretation). Moreover, we also predicted that it would be more difficult for participants to indicate that the trait was not

in the sentence on match trials than on “mismatch trials” (i.e., when the category label does not disambiguate the behavior in the direction of the tested trait probe – i.e., “librarian” tested with the trait probe “independent” and businessman tested with the probe word “shy”).

In addition, in order to rule out an explanation of the results based on mere category activation effects, we added a control condition in which the category labels were presented with neutral behaviors, and tested with the same traits presented on match trials. For example, the behavior “the businessman bought stamps at the post office” was tested with the trait “independent”. If differences in the difficulty of rejecting the trait probe words are observed independently of the behavior presented being an experimental or a neutral behavior, then effects can be attributed only to category activation effects. In contrast, we predicted that it would be harder for participants to reject the trait in the match condition, than in the control condition. This finding would prove that results are specifically due to an interaction between category activation and STIs made from behaviors provided.

Method

Participants and Design

Participants were 26 undergraduate students at University of California, Santa Barbara (21 women and 5 men). The critical trials of the experiment formed a 3 level (type of trial: match vs. mismatch vs. neutral) within-subjects design.

Stimulus Material

Experimental Trials. Seven ambiguous behaviors were used in this study. Some of the sentences were adapted from previous studies with ambiguous behaviors (Higgins, Rholes, & Jones, 1977; Sedikides, 1990), and others were taken from previous pre-tests with ambiguous behaviors (Ramos, 2006). Ambiguous behaviors could be interpreted by two personality traits (for example, the behavior *Only rarely did he change his mind even when its better if he had* may exemplify either the trait *Persistent* or *Stubborn*). Behaviors originally in English from Higgins et al. (1977) and Sedikides (1990) were translated and pre-tested with Portuguese participants. Behaviors were

presented to 23 individuals that evaluated them in two 9-point scales, corresponding to the two traits implied by the behavior. Participants were asked how applicable the trait presented in the scale was to the behavior. Responses could vary from 1 (not applicable at all) to 9 (totally applicable). Results showed that differences between evaluation means on the two trait scales were not significantly different, ($M^{TA} = 5.01$ and $M^{TB} = 5.49$), ($p > .05$). For a list of the stimuli used in this experiment, see Appendix A.

For each behavior, two category labels were selected. Each one of the category labels was applicable to each one of the two trait interpretations of the behavior. For example, the category “old man” is applicable to the interpretation of the behavior “Only rarely did he change his mind, even when it might well have been better if he had” in terms of the trait “stubborn”. In contrast, the category “manager” is more applicable to an interpretation of the behavior in terms of the trait “persistent”. In a pilot study, 36 participants were asked to evaluate on a 9-point scale whether each one of the two personality traits was applicable to describe the behavior, when performed by each one of the two actors. Results confirmed that one of the trait interpretations was judged as more applicable to describe the behavior when it was performed by the consistent actor ($M = 6.34$) than when it is performed by the alternative actor ($M = 5.48$), $t(39) = 4.03$, $p < .001$. In the same way, the alternative trait interpretation was judged as more applicable to describe the behavior when it is performed by the consistent actor ($M = 6.80$) than when it is performed by the alternative actor ($M = 6.09$), $t(39) = 3.28$, $p = .001$. For example, the trait “stubborn” was considered more applicable to describe the behavior “Only rarely did he change his mind, even when it might well have been better if he had” when it was performed by the “old man” than when the same behavior was performed by the “businessman”. By contrast, the trait “persistent” was considered more applicable to describe the behavior when it was performed by the “manager” than when the same behavior was performed by the “old man”.

At all, 42 experimental trials were created by combining the 7 ambiguous behavioral sentences with the 3 category conditions (category A, category B, and neutral category) and with the two trait probe words (Trait A and Trait B).

Filler Trials. Each of the 21 experimental sentences were also presented with 2 additional probe words that were actually included in the sentences (the verb, and other

word) in order to prevent participants from indicating only “no” responses (42 trials). In addition, in order to guarantee that any probe word was equally likely of eliciting a “yes” and a “no” response, trait versions of the sentences were developed, in which the personality traits were included. In these trait versions, the actor was described as performing a neutral behavior (e.g., “the stubborn/persistent old man parked his car near home”). The trait versions were followed either by the trait included in the sentence (42 trials) or by a word included in the sentence that could be either a verb or another word from the sentence (42 trials). For half of the participants, half of the sentences were tested with a verb and other half with another word from the sentence. For the other half of the participants, the half of the sentences that were tested with a verb were now tested with a word, and vice-versa.

Procedure

The experiment was run with the software E-prime (Schneider, Eschman & Zuccolotto, 2002). All instructions, stimuli and response measures were provided on the computer. Sessions were run in groups, but each participant worked individually in one computer. The study was presented as a “study on comprehension speed”. Instructions explained to participants that they would be presented with series of sentences (written in black) and that, after each sentence, a word would appear on the screen (written in blue). Participants’ task was to indicate, as fast and as accurately as possible, whether the word was included in the preceding sentence. Specifically, they were instructed to press the “yes” key (a red label with a “Y” on it was stuck on the “i-key” of the keyboard) if they thought that the word was in the previous sentence, and to press the “no” key (a red label with an “N” on it was stuck in the “e-key” of the keyboard) if they thought that the word was not included in the previous sentence. Participants were asked to maintain their index fingers on the keyboard throughout the experiment, in order to facilitate the task.

Each sentence was presented on the center of the screen for 2,000 ms and was followed by a blank screen presented for 500 ms. After that, the probe word appeared on the center of the screen until a response was provided by the participant. After participant’s response, a blank screen appeared for 1000 ms to discriminate between trials. Before the actual experiment started, six practice trials were presented. These

trials were composed by sentences with the same structure as the experimental sentences, but with neutral content not related with the content of any of the experimental sentences. After the practice trials, instructions informed that the experiment will be start. The first three trials of the experimental round were filler trials in order to annul start-up problems.

In total, participants went through 168 trials. We randomly divided experimental and filler trials in three groups of 56 trials. The presentation order of each group was randomized. Within each group, presentation of trials was also randomized. To avoid fatigue, a brief pause of 30 seconds was included in the middle of the experiment. At the end, participants were debriefed and thanked for their participation.

Dependent Measures

The dependent measures used were the percentage of incorrect responses (i.e., number of times participants incorrectly indicated that the trait probe was presented in the sentence), and time to give a correct response (i.e., time participants took to correctly indicate that the trait probe was not presented in the sentence).

Results

Error Rates.

In order to examine whether STIs from ambiguous behaviors are more likely when a category label disambiguates the trait interpretation of the behavior than when the behavior remains ambiguous, the number of times participants incorrectly indicated that the trait probe was in the sentence were entered in a 3 level (type of trial: match vs. mismatch vs. neutral) within subjects ANOVA. The ANOVA revealed a marginal significant effect, $F(2, 48) = 2.71, p = .07$. As expected, a planned comparison revealed that participants are more likely to incorrectly indicate that the trait was included in the sentence on match trials ($M = 5.71$) than on both neutral ($M = 4.57$) and mismatch trials ($M = 2.86$), $t(24) = 2.16, p = .02$, one-tailed. An additional planned comparison revealed no significant differences between mismatch and neutral conditions ($t(24) = 1.23, ns$, one-tailed). Results are presented in Table 1.

Table 1

Means Errors as a Function of Type of Trial

	Type of Trial		
	Match	Mismatch	Neutral
<i>M</i>	5.71	2.86	4.57

As previously outlined, each ambiguous behavior could be interpretable by two different personality traits (Trait A vs. Trait B). In the match condition category labels favored the trait interpretation tested by the probe word (Cat A–Trait A and Cat B–Trait B). In the neutral condition, the category label was irrelevant for any one of the two trait interpretations (Cat N–Trait A and Cat N–Trait B). In mismatch conditions, category labels didn’t favor the trait interpretation tested by the probe word (Cat B–Trait A and Cat A–Trait B). In order to better examine the pattern of results, comparisons were made separately for each trait interpretation of the behavior (Trait A vs. Trait B). For each behavior, each trait interpretation was randomly assigned to Trait A or Trait B conditions and separated ANOVAS were run for each Trait. The number of incorrect responses for one of the trait probe interpretations (Trait A) were submitted to a 3 level (type of category: Cat A vs. Cat B vs. Cat N) within subjects ANOVA. Planned comparisons confirmed that the number of errors was higher when a consistent category was provided ($M^{\text{CatA}} = 6.29$) than when both the category provided favored the alternative trait interpretation ($M^{\text{CatB}} = 4.57$), or when the category provided was neutral ($M^{\text{CatN}} = 4.00$), $t(24) = 1.57$, $p = .06$, being this effect marginal, $t(24) = 1.57$, $p = .06$, one-tailed.

In a similar way, we computed a 3 level (type of category: Cat A vs. Cat B vs. Cat N) within subjects ANOVA with the number of errors for the alternative probe trait (Trait B). Again, planned comparisons confirmed that the number of errors was higher when a consistent category was provided ($M^{\text{CatB}} = 5.14$) than when both the category provided favored the alternative trait interpretation ($M^{\text{CatA}} = 1.14$), or when the category provided was a neutral one ($M^{\text{CatN}} = 3.43$), $t(24) = 1.57$, $p = .04$, one-tailed.

Finally, in order to examine whether effects were due to mere category activation effects, the number of errors in the match condition were compared with the number of errors in the control condition, for both trait interpretation probes. Number of incorrect responses for one of the trait interpretation (Trait A) were submitted to a 2 level (condition: match vs. control) within subjects ANOVA. As predicted the ANOVA revealed that the number of incorrect responses was higher in the match condition ($M = 6.29$) than in the control condition ($M = 1.71$), $F(1, 24)=8,26$, $p < .01$. A similar 2 (condition: match vs. control) within subjects ANOVA was computed for the alternative trait probe (Trait B). Again, results showed that the number of incorrect responses was higher in the match condition ($M = 5.14$) than in the control condition ($M = .57$), $F(1, 24)=8,25$, $p < .01$. If results were merely due to category activation effects, the number of errors observed should be the same, regardless of the category label being presented with a trait relevant behavior or with a neutral behavior. Since the percentage of errors was much higher when a trait-relevant behavior was presented and approached zero in the neutral behavior condition, effects can only be explainable by STI effects.

Overall, consistent with our hypotheses, results indicated that when ambiguous behaviors were presented with a neutral or mismatch category STIs were weak for any of the possible trait interpretations of the behavior. Only when a category label was provided that disambiguated the trait interpretation of the behavior were stronger STIs observed. Thus, it seems much more likely that people spontaneously infer the trait “shy” from the behavior “The librarian says that his contacts with other people are rather limited” than when the same exactly behavior remains ambiguous by being performed by an unknown actor (e.g., He says that his contacts with other people are rather limited), or when the same behavior is performed by an actor that disambiguates the behavior, but in opposite direction to the “shy” trait inference (e.g., the businessman says that his contacts with other people are rather limited).

Response Times

In line with the recommendations of Ratcliff (1993), a cutoff criterion was used for responses faster than 200 ms and longer than 2000 ms. This same criterion has been applied by other authors using the probe recognition paradigm (Ham & Vonk, 2003;

Wigboldus et al., 2003). Only one participant was removed from the analysis based on this criterion.

Response times for correct responses were entered in a 3 level (type of trial: match vs. mismatch vs. neutral) within subjects ANOVA. No significant differences were revealed by the ANOVA, $F(2, 48) = 1.19, ns$. Results for correct responses were also analyzed separately for each trait probe interpretation of the behavior. For each behavior, one of the trait interpretations was randomly categorized as “Trait A” and the other trait interpretation as “Trait B”. Separate 3 level (type of category: Cat A vs. Cat B vs. Cat N) within subjects ANOVAs were run for correct response times for both Trait probe A and Trait probe B. Neither of the ANOVAs revealed significant effects, $p > .05$. Thus, no differences were found in response times. Although previous studies with the recognition probe paradigm have reported differences in response times (McKoon & Ratcliff, 1986; Wigboldus et al., 2003), there were also cases in which differences were only observed in error rates (see Uleman et al., 1996, Experiment, 1).

Finally, we compared response times for correct responses in the match condition with the control condition, for both trait implied probes. Correct response times for Trait A probe were submitted to a 2 level (condition: match vs. control) within subjects ANOVA. Consistent previous results, the ANOVA revealed that participants took significantly less time to correctly reject the trait in the control condition ($M = 829$ ms) than in the match condition ($M = 1019$), $F(1, 24) = 9.84, p < .01$. A similar analysis was computed for probe Trait B. Response times for correct responses for Trait B as probe were analyzed in a 2 level (condition: match vs. control) within subjects ANOVA. Again, the ANOVA revealed a main effect for condition, showing that participants took significantly less time to correctly reject the trait probe in the control condition ($M = 871$ ms) than in the match condition ($M = 1034$ ms), $F(1, 24) = 11.65, p < .01$. These results corroborate the pattern observed for correct responses, and confirm that the results are not only due to category activation effects.

Discussion

In this experiment, our goal was to test the principle according to which trait inferences are more likely to occur spontaneously when traits are easily inferable from behaviors. Since previous STI research has made use of behaviors highly diagnostic of

specific personality traits, evidence for the spontaneity of the trait inference process has been generally obtained. Contrary to previous studies, in the present experiment we used ambiguous behaviors. Ambiguous behaviors were pre-tested so they imply simultaneously two different personality traits. Importantly, the level to which the two personality traits were implied by each ambiguous behavior was lower than the level of trait implication of behaviors typically used in STI research.

Our hypotheses followed the principle according to which the more traits are easily inferable from behaviors, the more STIs will be likely to occur. Specifically, we predicted that evidence for STIs would be stronger when a category label is associated with the actor that helps to disambiguate the appropriate trait inference (i.e., match condition) than when the behavior remains ambiguous by being presented with a neutral category label (i.e., neutral condition). In the latter case, because the behavior is ambiguous in trait terms, participants are less likely to spontaneously infer either one of the two implied personality traits. In addition, we also predicted that when the category label that is associated with the actor favors one of the trait inferences, evidence for the spontaneity of the alternative trait inference would be less likely (i.e. mismatch condition).

In agreement with our hypotheses, our results showed that participants are more likely to incorrectly indicate that the trait was included in the previous behavior when the category label disambiguated the behavior in the direction of the trait probe word (match condition), than both when the behavior remains ambiguous by being presented with a neutral category (neutral condition), and when the category label associated with the actor favors the alternative trait inference (mismatch condition).

These results are important in revealing the flexibility of the STI process. When a behavior is ambiguous in trait terms, STIs are weaker. The underlying process is assumed to be the same in both cases. However, with behaviors highly implicative of a personality trait, the trait is strongly activated and evidence for STI is obtained, while with ambiguous behaviors both implied traits are activated to a lesser extent, and STI are less likely to emerge.

Thus, our findings support the notion that ambiguous behaviors result in the activation of trait inferences, but only to a certain degree. The STI process seems to be restrained until more information is gathered that helps to clarify which is the most

appropriate trait inference. That would make sense to the functioning of an effective comprehension system. An efficient system should not make strong inferences until information is inputted into the system that disambiguates if those inferences are appropriate for the given circumstances. Only when more information about the actor is provided, STI became stronger. In this case, the category label contributes to the occurrence of stronger STIs, from a behavior that otherwise would remain ambiguous.

Findings from the present experiment are especially relevant if we consider that many of our social behaviors are fundamentally ambiguous, and that most behaviors allow more than one trait interpretation. For example, a friendly behavior may be interpreted as false or hypocritical. An intelligent behavior may be interpreted as a lucky or effortful behavior. An adventurous behavior may otherwise be seen as an irresponsible behavior, and so on. Thus, what our results suggest is that information about the characteristics of the actor is fundamental to disambiguate the different possible trait interpretations of an observed behavior, and that this process will determine the level to which trait inferences will occur spontaneously.

An important question concerning our results is related with the mechanism responsible for the observed effects. Our goal was to examine STIs occurrence for behaviors in which the trait was not so easily inferable. In order to achieve that, we applied behaviors that imply traits to a lesser extent, than is usual the case in STI research. In addition, the behaviors used were ambiguous in the sense that they implied simultaneously two different personality traits. It should be noticed that, in this case, obtaining weaker evidence of STIs may be due to the fact that the relation between the behavior and implied traits is weak *or* because the behavior can result in multiple trait inferences. That is, weaker STIs may be merely explained by the fact that behaviors imply traits to a lower degree, independently of behaviors implying one or two traits. Alternatively, weaker STIs may result from the fact that behaviors imply simultaneously two different trait inferences, independently of traits being weakly or strongly implied by the behaviors.

A way to clarify this issue in future studies would be by manipulating these two features independently. That is, different studies should explore the occurrence of STI both with behaviors that only imply one trait but to a lower extent, and also with behaviors that strongly imply simultaneously two traits. However, it is worth to mention

that these factors are probably not independent. Ambiguous behaviors tend to exhibit relatively low levels of trait implication (see Ramos, 2006), and low trait implicative behaviors are certainly more likely to be ambiguous, and to be opened to alternative trait interpretations. So, despite our results being not totally clear about this issue, our intuition is that these are interrelated aspects.

In sum, the pattern of results from our first experiment clearly supports a flexible view of the STI process. The same ambiguous behavior can lead to different spontaneous trait inferences, depending on the actor that performs the behavior. In addition, when the ambiguous behavior is performed by an unknown actor, the two possible trait inferences are activated but to a lesser extent. These results suggest a different way of conceptualizing the STI process. Sometimes, STI seems to be regarded as an inherent property of behaviors. However, if the meaning of behaviors themselves change depending on the actor that performs the behavior (and probably depending on other contextual factors), STIs are probably best described as flexible and dependent upon the specific comprehension of the behavior.

SECTION II

SPONTANEOUS TRAIT INFERENCES WORK IN LINE WITH LOCAL COHERENCE REQUIREMENTS

EXPERIMENT 2

Probe Recognition with Behaviors and Situational Continuations

In agreement with our second principle about the conditions for STIs to occur, spontaneous trait inferences operate according to local coherence requirements. Local coherence is a basic need of the comprehension system that is assumed to guide the functioning of our inferential activity in general terms (McKoon & Ratcliff, 1992). According with this principle, we make the general prediction that spontaneous trait inferences will be facilitated when they are consistent with previous information about the actor (i.e., when they increase the local coherence of received information) and will be inhibited when they are inconsistent with previous information about the actor (i.e., when they lead to text local incoherencies).

In general, previous studies were not able to test this prediction because participants were presented with single behaviors that were highly implicative of a specific personality trait. There was no space for effects due to local coherence to become visible. As a consequence, the flexibility of the spontaneous trait inference process remained largely unexplored (but see Wigboldus et al., 2003, for an exception).

In Experiment 2, we manipulated the local coherence of the information by introducing two modifications into the behavioral sentences typically used in STI research. First, we associated a social category to the actors of trait-implying behaviors, similar to Wigboldus and colleagues (Wigboldus et al., 2003). The social category could be either consistent or inconsistent with the trait implying behavior. When the social category of the actor was consistent with the behavior we predicted that spontaneous trait inferences will be facilitated. By contrast, when the social category of the actor is inconsistent with the behavior we predict that spontaneous trait inferences will be inhibited. While in the former case the trait inference is in conformity with a coherent picture of the event, in the latter case making a trait inference from the behavior would be incoherent with the previous stereotyped knowledge about the actor.

As a second modification, we added a situational continuation to each behavioral description. This continuation provided an alternative situational explanation for the behavior. That is, by considering the situational context of the behavior, behaviors would be less likely to lead to dispositional trait inferences about the actor and more likely to lead to situational inferences. We predicted that the situational content of the sentence would have a different impact on the type of spontaneous inferences made by participants, depending on the social category that is associated with

the actor. Specifically, according to our hypothesis, spontaneous situational inferences (Ham & Vonk, 2003; Lupfer et al., 1995) will be more likely to occur when the social category of the actor is inconsistent with the behavior, but will be less likely to occur when the social category of the actor is consistent with the behavior. In line with the principle of local coherence, in the first case the spontaneous situational inference would provide greater local coherence to the information that is being processed (i.e., the situational inference is consistent with the information about the actor), while in the later case the spontaneous situational inferences would be incoherent given the previous knowledge about the actor.

Ham and Vonk (2003) provided evidence that STIs and SSIs may co-occur for the same behavior. Here, we predict that providing an inconsistent category label will inhibit STIs comparing with a condition in which a consistent category label is provided. In addition, we predict that an inconsistent category label will facilitate SSIs, in comparison with a condition in which a consistent label is provided. It is important to notice that our predictions are not in contradiction with the Ham and Vonk findings. When the trait inference and the situational inferences are equally relevant to the interpretation of the event, as in Ham and Vonk (2003) study, both inferences are likely to occur spontaneously (although probably to a weaker extent than if the event only imply a trait inference or a situational inference). In our study, this would correspond to a condition in which the sentences are presented with a neutral category label (see Experiments 3 and 4). However, in the present study we do not predict that both trait and situational inferences occur to the same degree because we assume that the magnitude of both types of spontaneous inferences will be modeled by the social category of the actor. That is, the category label provided will tend to favor (or inhibit) those spontaneous inferences that are more coherent (or less coherent) with the interpretation of the event.

The present study may be relevant for different reasons. First, as previously said, by providing information about the social category of the actor and about the contextual situation in which behaviors were performed, this study can be viewed as an initial step to examine spontaneous trait inferences occurrence in more socially meaningful contexts.

Second, in this study we could examine not only spontaneous trait inferences but also spontaneous situational inferences. Few studies have been carried out on the spontaneous occurrence of situational inferences (Ham & Vonk, 2003; Lupfer et al., 1995). This might be surprising, given the importance of the dichotomy between personal and situational forces as main determinants of behavior. As the same way as it is important to explore whether trait inferences are made spontaneously, it is also relevant to understand whether features of the situation are also spontaneously inferred.

The study of spontaneous situational inferences might also be important in order to inform us about the more general issue of whether spontaneous trait inferences are guided by the same principles as other spontaneous social inferences. Although the focus of research has been predominantly on the occurrence of spontaneous trait inferences, when other social inferences were studied evidence for their spontaneous occurrence has been obtained. Thus, an important research question is to understand whether the different social inferences follow the same processing principles, and to explore how different inferences interact with each other. In the present study, we specifically focus on the occurrence of both spontaneous trait inferences and spontaneous situational inferences. Evidence from this study may inform us about how both trait and situational inferences are constrained by the social category that is ascribed to the actor of the behavior, as well as provide indications about the relation between the two types of inferences.

As in the previous experiment, in this study we applied the recognition probe paradigm. Participants were presented with sentences composed of a trait-implicating behavior (e.g., “asked where the stars came from...”), a behavior that implies the trait “curious”) and a situational continuation (e.g., “...this being one of the homework questions”, a situation that implies the word “duty”). Each behavior description was associated with a social category that could be either consistent with the trait-implicating behavior (e.g., “the best student”) or inconsistent with the trait-implicating behavior (e.g., “the most popular student”). Each version of the sentence was tested with two different probe words: the trait implied by the behavior (e.g., curious), in order to study the occurrence of spontaneous trait inferences, and the gist of the situation (e.g., duty), to explore the occurrence of spontaneous situational inferences.

According to our predictions, sentence implying behaviors, information about the social category of the actor, and information about the situational context in which the behavior was performed will interplay with each other to guide the inferential process. Inferences that are more consistent with the information received (i.e. inferences that provide greater coherence) will be the ones more likely to occur spontaneously. Specifically, our hypothesis asserts that when the social category is inconsistent with the behavior, spontaneous trait inferences will be less likely to occur than when the social category is consistent with the behavior, in line with results from Wigboldus et al (Wigboldus et al., 2003). However, when the social category is inconsistent with the behavior, despite leading to weaker spontaneous trait inferences, it will result in *more* spontaneous situational inferences than when the social category is consistent with the behavior. Thus, the pattern of occurrence of both spontaneous trait inferences and spontaneous situational inferences will be in agreement with coherence principles.

In terms of pattern of responses within the probe recognition paradigm, this would mean that in the inconsistent category condition it would be easier for participants to correctly reject the probe trait as not being included in the sentence, but it would be more difficult to reject the situational probe (in comparison with the category-consistent condition).

Method

Participants and Design

Participants were 103 undergraduate students at University of Lisbon (91 women and 12 men). The critical trials of the experiment formed a 2 (social category: consistent vs. inconsistent) X 2 (probe word: trait, situational gist) within-subjects design.

Stimulus Material

Experimental Trials. We selected six trait-implying behaviors from norms in the Portuguese language collected by Ferreira, Morais, Ferreira, and Valchev (2005). To each sentence, a situational continuation was added, also taken from the same norms (Ferreira et al, 2005). Three of the situational continuations were slightly adapted to the

present experiment. The sentences used obeyed two criteria. First, the situational continuations added to the trait-implying behaviors result in higher situational attributions than when neutral continuations were added to the same trait-implying behaviors. Second, when situational continuations were added to the trait-implying behaviors, sentences led to lower trait attributions than when neutral continuations were added (for further details see Ferreira et al., 2005).

For each trait-implying behavior, a category label was selected that was stereotype-consistent with the behavior and also one that was stereotype-inconsistent with the behavior. In a pilot study, 44 participants evaluated each trait-implying behavior associated with each one of the category labels in a 9-point scale. Participants were asked to indicate “to what extent the behavior was stereotype-inconsistent or stereotype-consistent with the actor”. Participants should indicate the 1-point of the scale if they thought that the behavior was stereotype-inconsistent with the actor and the 9-point of the scale if they thought that the behavior was stereotype-consistent with the behavior. Results showed that behaviors were judged as more stereotype-consistent when associated with the stereotype-consistent label ($M = 5.84$) than when associated with the stereotype-inconsistent label ($M = 4.07$), $t(44) = 6.26, p < .001$.

Based on pre-tests results, 24 experimental trials were presented to participants, by varying systematically the 6 sentences, the type of category label (consistent vs. inconsistent) and the type of probe word (trait vs. situation) (see Appendix B for a complete list of the sentences and probe words used in the present experiment).

Filler Trials. In order to avoid only “No” responses, 12 filler trials were added in which the same sentences paired with each one of the category labels were tested with a probe word included in the sentence (the probes were words randomly selected from the sentences). Also, in order to prevent that participants associating trait probes with “No” responses, 12 filler trials were created in which the six experimental probe traits were included both in the consistent-category and in the inconsistent-category version of the sentences. These trials were developed by matching each category label with a neutral behavior (e.g., The curious best student of the school ate a hamburger at lunch). These same filler trials with the trait included were also tested with a probe word that was included in the sentence, in a total of 12 additional filler trials.

Procedure

The procedure was identical to Experiment 1, except for the following. Sentences were presented for 2500 ms due to the fact of being longer in length. Pilot reading times showed that the time necessary to read the entire sentence was about 2500 ms. In addition, because in this experiment the total number of trials was smaller, the presentation of the trials followed without pauses. In total, participants went through 50 trials. Trials were divided in three blocks of behaviors, with an equal number of experimental and filler trials. The three blocks of behaviors were randomly presented to each participant. Behaviors within each block were also randomly presented to each participant.

Dependent Measures

The dependent measures used were the same as in the previous study: percentage of errors and response times of correct responses.

Results

Error Rates

We started by analyzing participants' responses in terms of number of times participants incorrectly indicated probes words were included in the sentences. Proportions of errors were analyzed in a 2 (social category: consistent vs. inconsistent) X 2 (probe word: trait, situation) within-subjects ANOVA. The ANOVA yielded a main effect for social category, $F(1, 102) = 14.25, p < .04$, revealing a higher percentage of errors when the social category was consistent ($M = 6.72$) than when it was inconsistent ($M = 5.10$) with the trait implying behavior. The main effect for probe word was also significant, $F(1, 102) = 14.25, p < .001$, showing a higher number of errors for trait probes ($M = 7.93$) than for situational probes ($M = 3.88$).

More relevant to our hypothesis is the interaction between category dimension and type of probe word. According to our predictions, for trait probes the number of incorrect responses should be higher when the social category is consistent than when it is inconsistent with the behavior. By contrast, for situational probes we predicted that the number of incorrect responses should be higher in the inconsistent than in the

consistent social category condition. The ANOVA revealed that the interaction between category dimension and probe word was significant, $F(1, 102) = 4.60, p < .034$. Planned comparisons showed that the percentage of errors for trait probes was significantly higher in the consistent category condition ($M = 9.55$) than in the category inconsistent condition ($M = 6.31$), $t(102) = 2.71, p < .01$ (one tailed). For situational probes, the percentage of errors didn't differ significantly between consistent ($M = 3.88$) and inconsistent category conditions ($M = 3.91$), $t < 1$. See Figure 1.




Figure 1. Proportion of errors as a function of probe type and category dimension

Thus, as expected, for trait probe words, number of errors was higher when the social category of the actor was inconsistent with the behavior. The predicted pattern of response errors didn't emerge in terms of situational probe words.

Response Times

Participants' responses were also analyzed in terms of the time taken to correctly indicate that probe words were not presented in the sentences. Results are reported in milliseconds. A 2 (social category: consistent vs. inconsistent) X 2 (probe word: trait, situation) within subjects ANOVA was carried out on response times for correct responses. A main effect emerged for type of probe word, $F(1, 102) = 41.29, p < .001$, showing that participants took significantly more time to correctly reject trait probes ($M = 1059$ ms) than situational probes ($M = 952$ ms).

The ANOVA also revealed a marginally significant interaction between probe word and category dimension, $F(1, 102) = 2.87, p = .09$. A planned comparison revealed that participants took longer to correctly reject the situational probe word when

the category of the actor was inconsistent with the behavior ($M = 976$ ms), than when the category of the actor was consistent with the behavior ($M = 927$ ms), $t(102) = 2.46$, $p = .015$. By contrast, response times for trait probe didn't differ with category dimension ($M^{\text{cons}} = 1061$ ms and $M^{\text{incons}} = 1056$), $t < 1$. Thus, despite not being evident in the error rates, response times provided some indication that spontaneous situational inferences are more likely to occur when the social category of the actor is inconsistent with the trait implying behavior.




Figure 2. Response Times for Correct Responses as function of probe type and category dimension

Discussion

In the present study, our goal was to explore whether the local coherence of the information presented guides the occurrence of both spontaneous trait inferences and spontaneous situational inferences. In order to test these predictions, trait implying behaviors were presented to participants, similar to what is typically done in STI research. However, in this case behaviors were presented with a continuation that implied a situational inference about the event. In addition, each behavioral sentence was presented either with a consistent category label or with an inconsistent category label. According with the principle of local coherence, we predicted that in the category-inconsistent condition participants would be less likely to make spontaneous trait inferences than in the consistent-category condition. By contrast, participants in the

inconsistent-category condition would be more likely to make spontaneous situational inferences than participants in the consistent-category condition.

In line with our predictions, our results showed that participants were less likely to incorrectly indicate that probe traits were included in previous sentences when the behavior was performed by a member of an inconsistent category than when the same exactly behavior was performed by a member of a consistent category. These results replicate findings from Wigboldus et al. (2003), in this case using more complex material that included information about the situational context in which the behavior was performed. One difference between our results and Wigboldus et al. (2003) was that, while Wigboldus et al. (2003) found differences only in response times and not in accuracy, in our case the inverse occurred with differences in STIs being observed in accuracy and not in response times. It should be noted, however, that a trade off between accuracy and response times can be expected in this paradigm. Specifically, people can choose to produce fast responses with relatively low accuracy, or they can choose to give slower responses with higher levels of accuracy. If that was the case, we should observe faster responses in our experiment, comparing with Wigboldus et al (2003) experiments. However, that was not the case. Actually, response times in our study were longer than in Wigboldus et al. (2003). However, because both experiments made use of different materials, and were run with different participants direct comparisons are difficult. As an alternative explanation for the different results, we may also hypothesize that participants that provide an incorrect answer are probably the ones that would take more time to provide a correct answer. Conversely, participants with the longest times to provide a correct response would be the ones that would fail if they respond more quickly. In this sense differences in one dependent measure may obscure differences in the other measure.

In this sense, the pattern of results in our study can be seen as being actually stronger than the one found by Wigboldus et al. (2003). Differences between category-consistent and category-inconsistent conditions were revealed not only in the time participants took to indicate that the trait was not included in the sentence, but in participant's likelihood of indicating that the trait was really included in the sentence previously read.

Evidence was also obtained in terms of the magnitude of spontaneous situational inferences, depending on the social category of the actor. Specifically, we found that participants took more time to correctly indicate that the situational word was not included in the sentence in the category-inconsistent than in the category-consistent condition. In this case, however, differences were only observed in response times, and not in accuracy level. This indicates that differences between category conditions were stronger for trait than for situational probe words. Despite weaker, this pattern is in line with our predictions. That is, it suggests that when the social category of the actor is inconsistent with the behavior, participants are more likely to make spontaneous situational inferences than when the social category of the actor is consistent with the behavior.

In sum, the pattern of responses for both trait and situational probes provides indication that when the social category was inconsistent with the behavior it was more difficult to make a spontaneous trait inference, but it was easier to make a spontaneous situational inference, compared with a condition in which the same behavior was presented with a consistent social category. For instance, if someone asks where the stars come from, even if this is a homework question, people are more likely to spontaneously infer the trait “curious” and less likely to infer the situational word “duty” when the behavior is performed by the “best student of the school”, in comparison with a condition in which the same behavior is performed by the “most popular student of the school”.

These results are consistent with our view that spontaneous trait inferences have great flexibility, depending on the actor that performs the behavior, and support the idea that the occurrence of spontaneous trait inferences is dependent upon a coherent picture of the information that is processed. Since spontaneous situational inferences were also affected by the type of social category of the actor (although to a weaker extent, only in response times), the social category of the actor is a major determinant not only of spontaneous trait inferences, but also of other spontaneous social inferences that might be made from the information presented.

One aspect not clarified by our study is whether the pattern of results is explainable by facilitation or inhibition effects. Specifically, we cannot conclude whether spontaneous trait inferences were facilitated when the social category of the

actor was consistent with the behavior or whether spontaneous trait inferences were inhibited when the social actor was inconsistent with the behavior, or whether both facilitation and inhibition effects are intervening . The same reasoning is applicable to the occurrence of spontaneous situational inferences. It is not clear if spontaneous situational inferences were inhibited in the consistent-category condition or facilitated in the inconsistent-category condition, or both. In order to elucidate this question, it would be necessary to introduce a condition in which the same behaviors are presented with a neutral category label.

Another aspect of this study that deserves further consideration is the fact that differences between social category conditions were stronger for spontaneous trait inferences than for spontaneous situational inferences. In the former case, differences were evident in accuracy data, whereas in the latter case differences only arose in response times. One possible interpretation of this pattern of findings is that spontaneous trait inferences are in fact more prevalent in the comprehension of social information, compared with other types of social inferences. However, another possible factor that might be responsible for this pattern of results is related to the behavioral material used in this experiment. Trait implying behaviors were always presented at the beginning of the sentence while situational contents were always presented at the end of the sentence. This may have facilitated the finding of differences for spontaneous trait inferences. In Experiment 3 we tested this alternative interpretation, by reversing the order of presentation of situational and dispositional contents of the sentences.

Another related factor that might explain the findings of stronger differences for spontaneous trait inferences than for spontaneous situational inferences is the fact that sentences were composed of behaviors originally tested to strongly imply a personality trait (e.g., “asked where the stars came from” is a behavior pre-tested to imply the trait “curious”). Situational continuations were then added to these behaviors. However, overall the sentences might continue to be more implicative of traits than of situations. In order to test whether the stronger findings in terms of spontaneous situational inferences were due to the type of stimulus sentences, in Experiment 4 we used ambiguous dispositional-situational behaviors. In this case, exactly the same sentence could imply both a dispositional trait inference and a situational inference.

EXPERIMENT 3

Reverting the Order of Behaviors and Situations

In Experiment 3, our goal was to better clarify some aspects from our previous study. First, in the present study we were interested in ruling out social category activation effects as an alternative explanation of the results from Experiment 2. In the previous study, it was possible that social categories activate traits independently of sentence presented. For example, it might be the case that the category *the best student* activates the trait *curious*, regardless of the behavioral information provided subsequently.

As some authors suggest (Wigboldus et al., 2004) it is important to assure that stereotype activation do not explain results independently of STIs. In order to rule out this possibility, in Experiment 3 we presented the same social category labels with neutral behaviors (e.g., *The best student, the most popular student cycles through the street*). These trials were followed by the same trait probe words as in experimental trials (e.g., in this case, *Curious*). If results from the previous study are only due to effects arising from the social category, the pattern of responses to trait probe words should be the same regardless of the type of behavior presented. By contrast, if responses are dependent on an interaction between social categories and spontaneous trait inferences made from behaviors, then response differences for trait probe should be evident when behaviors imply relevant personality traits, but not when neutral behaviors are presented.

A second aim of the present study was to explore the reasons why we found stronger evidence for spontaneous trait inferences than for spontaneous situational inferences in our previous study. In fact, while differences in spontaneous trait inferences were found in for error rates, differences in spontaneous situational inferences were only observed in response times. There was one aspect in the study material that may have contributed to this pattern of results. The fact that the behavior itself was always presented at the beginning of the sentence, and the situational explanation at the end of the sentence, may have facilitated the finding of differences for trait probe words. The order in which elements of the sentences were presented might have influenced the results in different possible ways. One possibility is that the social category had a larger effect on trait inferences because behaviors were presented first. By contrast, because situational continuations are presented at the end of the sentence, the influence of the social category on spontaneous situational inferences

could be reduced. Alternatively, order may have facilitated the correct rejection of situational probes. That is, because situational continuations were presented at the end of the sentences, participants might be more likely to have the situational content still in working memory at the time they give their responses. If this is the case, the material used would be more suitable to detect differences in terms of spontaneous trait inferences than in terms of spontaneous situational inferences.

Finally, in this study we were also interested in exploring whether the moderation of the spontaneous inferences by the type of social actor is due to facilitation or inhibition effects, or both. In order to explore this issue, in addition to being presented with a consistent or inconsistent social category, the same sentences were also presented with a neutral category (i.e., “He”).

We applied the probe recognition paradigm as in the previous study. The hypotheses were also the same. In the category consistent condition, participants should be more likely to make spontaneous trait inferences but less likely to make spontaneous situational inferences (in comparison with the category-inconsistent condition). We do not make major claims about whether effects result from facilitation or inhibition mechanisms, since according to local coherence principles both facilitation and inhibition effects are predictable.

Method

Participants and Design

Participants were 39 undergraduate students at Lisbon University Institute (36 women and 3 men). The critical trials of the experiment formed a 3 (social category: consistent vs. inconsistent vs. neutral) X 2 (probe word: trait vs. situation) within-subjects design. A condition was also added in which the consistent and inconsistent categories were paired with neutral behaviors, in order to control for category activation effects.

Stimulus Material

In this study, we used the same stimuli sentences from Experiment 2, including a condition in which the sentences were presented with a neutral category. However, in this study, the order of behaviors and situational continuations was inverted (see

Appendix C). In addition, trials were included in which the same social category labels were presented with neutral sentences.

Procedure

The basic procedure and instructions were the same as those used in Experiment 2.

Dependent Measures

The dependent measures were the same as in previous studies: percentage of errors and response times for correct responses.

Results

Error Rates

The percentage of incorrect responses was analyzed in a 3 (social category: consistent vs. inconsistent vs. neutral) X 2 (probe word: trait, situation) within-subjects ANOVA. Main effects were observed both for category dimension [$F(2, 76) = 5.99, p = .003$], and for type of probe word [$F(1, 38) = 7.05, p = .012$]. The predicted interaction between category dimension and type of probe word was also significant, [$F(2, 76) = 29.13, p < .001$]. Planned comparisons revealed that, for trait probe words the number of incorrect responses was significantly higher in the consistent-category condition ($M = 18.94$) than in both the neutral-category ($M = 6.62$) and in the inconsistent-category ($M = 7.28$) conditions, $t(38) = 6.40, p < .001$ (one tailed planned comparisons). Neutral-category and inconsistent-category conditions didn't differ significantly, $t(38) = .35, ns$ (one tailed).

For situational probe words, additional planned comparisons revealed that the number of incorrect responses was lower in the consistent-category condition ($M = 5.99$) than in both the neutral-category ($M = 9.29$) and inconsistent-category ($M = 8.49$) conditions, $t(38) = 2.01, p = .026$ (one tailed planned comparisons). Neutral and inconsistent category conditions also didn't differ, $t < 1$.




Figure 3. Response Times for Correct Responses as function of probe type and category dimension

An additional ANOVA was computed in order to explore the effects of social category activation. Percentage of errors was analyzed in a 2 (social category: consistent vs. inconsistent) X 2 (sentence type: experimental vs. neutral) X 2 (probe word: trait vs. situation) within subjects ANOVA. Consistent with our hypothesis, a significant three-way interaction between category label, experimental condition, and probe type emerged, $F(1, 38) = 29.03, p < .001$, revealing a significant interaction between category dimension and probe word in the experimental, but not in the neutral, condition. Additional planned comparisons revealed no significant differences between number of errors for trait probes in the consistent-category and inconsistent-category, in the control condition, $t < 1$. Thus, presenting different category labels with neutral sentences didn't have any effect in terms of percentage of incorrect responses. Although less relevant, planned comparisons revealed that in the control condition the number of errors for situational probes was significantly different between consistent-category and inconsistent-category conditions. However, the differences observed were in the opposite direction to our hypothesis, with more errors being found in the consistent-category condition ($M = 2.41$) than in the inconsistent-category condition ($M = .85$). These results prove that our findings cannot be explained by any previous association

that might exist between situational probes and category labels. Any previous influence would function the other way around to our predictions.

Response Times

Response times for correct responses were entered in a 3 (social category: consistent vs. inconsistent vs. neutral) X 2 (probe word: trait, situation) within-subjects ANOVA. The ANOVA yield a main effect for trait probe, $F(1, 38) = 7.81, p < .01$, and also a significant interaction between category dimension and probe word, $F(2, 76) = 4.75, p = .011$. Follow-up analyses revealed that participants took more time to correctly reject the trait probe in the consistent-category condition ($M = 1001$ ms) than in both the neutral-category condition ($M = 959$ ms) and in the inconsistent-category condition ($M = 918$ ms), $t(38) = 2.92, p = .003$ (one tailed). Neutral-category and inconsistent-category conditions didn't differ significantly, $t(38) = 1.54, ns$ (one tailed).

Planned comparisons also revealed that participants took more time to correctly reject situational probes in the inconsistent-category condition ($M = 1052$ ms) than both in the neutral-category condition ($M = 1007$ ms) and in the consistent-category condition ($M = 978$ ms), $t(38) = 1.82, p = .038$, one tailed. Neutral-category and consistent-category conditions didn't differ significantly, $t < 1$.




Figure 4. Response Times for Correct Responses as function of probe type and category dimension

An ANOVA was also computed to analyze category activation effects. Response times for correct responses were submitted to a 2 (social category: consistent vs. inconsistent) X 2 (sentence type: experimental vs. neutral) X 2 (probe word: trait vs. situation) within subjects ANOVA. For error rates, the predicted three-way interaction emerged in response times, $F(1, 38) = 6.77, p = .013.$, revealing a significant interaction between category dimension and trait probe in the experimental condition, but not in the neutral control condition. Planned comparisons revealed that the time to correctly reject trait probes was not significantly different for consistent-category and inconsistent-category conditions, in the neutral sentence condition, $t(38) = 1.14, ns.$ Also, no significant differences between consistent-category and inconsistent-category conditions occurred for the time participants took to correctly reject situational probes, in the neutral sentence condition, $t < 1.$ Thus, in agreement with accuracy results, these findings show that differences in performance cannot be uniquely explained by effects resulting from social category activation.

Discussion

Results from Experiment 3 are consistent with the pattern from our previous experiment. Overall, findings support the idea that when a behavior is performed by an actor from a consistent social group, participants are more likely to make spontaneous trait inferences and less likely to make spontaneous situational inferences than when the same behavior is performed by an actor from an inconsistent social group.

Findings from this experiment made clear that results cannot be explained by mere category activation effects working independently of the type of behavior presented. When the same social categories are presented together with neutral behaviors, the expected category-based differences are not found in terms of both trait and situational spontaneous inferences. Thus, responses to probe words seem to be specifically determined by an interaction between the social category of the actor and spontaneous inferences that are made from behaviors.

A second goal of this study was to explore whether social category differences in terms of spontaneous inferences were due to facilitation or inhibition effects. In terms of spontaneous trait inferences, our study tends to support the view that trait inferences are *facilitated* when a consistent category is presented. Participants were more likely to

incorrectly indicate that the trait was in the sentence, and took also more time to correctly indicate that the trait was not in the sentence in the consistent condition, compared with the inconsistent and neutral category conditions. By contrast, no evidence for inhibition of spontaneous trait inferences was observed in this study. Accuracy levels and response times in the inconsistent-category condition were not different from those in the neutral-category condition.

For spontaneous situational differences, the pattern of results is less clear. In terms of errors, results suggest *inhibition* of situational inferences in the consistent-category condition, because in this case participants gave fewer incorrect responses than in both the inconsistent and neutral category conditions. However, pattern of response times indicates a certain degree of *facilitation* of spontaneous situational inferences in the inconsistent-category condition, since participants took more time to correctly reject the situational probes in the inconsistent than in both consistent and neutral category conditions. Thus, while inhibition effects were evident in accuracy, a slightly facilitation effect was observable in response times. This pattern seems to suggest that inhibition of spontaneous situational inferences by a consistent-category is stronger than facilitation of spontaneous situational inferences by an inconsistent-category.

There are some aspects of the present pattern of results that deserve further consideration. First, our results showing facilitation effects of spontaneous trait inferences by the presentation of a consistent-category contrast with findings from Wigboldus et al., (2003). In their case, they found inhibition of spontaneous trait inferences when an inconsistent-stereotype was associated with the actor, and no facilitation effects. According with the authors (Wigboldus et al., 2003) one reason that may explain the absence of facilitation effects is related with the type of behaviors used in their study. Because behaviors so strongly imply a personality trait, little room was left for additional effects of the social category. In our case, findings of facilitation effects may be due to the fact that behaviors were presented with a situational continuation. Thus, the resulting sentences were not so extreme in the extent to which they implied the traits, with more room being left for effects of social category activation.

The previous account explains why facilitation effects were observed in our study, and not in the study conducted by Wigboldus et al., (2003), but it doesn't

necessarily explain the reason why inhibition effects were found in Wigboldus et al., (2003) study, but not in our case. However, we can probably apply a similar reasoning. We can think that because our sentences only resulted in moderate spontaneous trait inferences, not much room was left for additional inhibition effects resulting from category presentation.

Another factor that may explain the absence of inhibition effects in spontaneous trait inferences is related with the type of category labels associated with the actors in our study and in the study from Wigboldus et al., (2003). It is possible that category labels used by Wigboldus et al. (2003) study activate stronger stereotypes than the ones used in our study. Because inhibitory mechanisms may have an important role in maintaining existing stereotypes (Wigboldus et al., 2003), stronger inhibitory effects would be expected when better established stereotypes are used. That is, the stronger the stereotype, the more participants will try to preserve it by inhibiting inconsistent trait inferences.

Finally, another goal of this study was to examine whether the order of presentation of the trait content and of the situational content within sentence would make a difference in the occurrence of both trait and situational spontaneous inferences. We specifically wanted to explore whether order of presentation could be responsible for the stronger category effects on spontaneous trait inferences than on spontaneous situational inferences observed in our previous study. This time, accuracy differences as a function of category dimension were observed both in trait and situational spontaneous inferences. The fact that differences in accuracy were also observed for spontaneous situational inferences suggests that order of presentation of situations makes a difference. Effects of category dimension on spontaneous situational inferences tend to be stronger when they appear initially in the sentence. However, as in the previous study, category differences were overall stronger for spontaneous trait inferences than for spontaneous situational inferences also in the present study. Thus, order might have some influence, but per se is not sufficient to explain the overall pattern of results.

Another aspect that might be intervening in the pattern of results is the fact that behaviors were pre-tested to imply a trait, whereas situational continuations were added *a posteriori*, and not the other way around. Because of that, in Experiment 4 ambiguous

dispositional/situational sentences were used. In this case, exactly the same sentence could imply both a spontaneous trait inference and a spontaneous situational inference. Others have considered this to be more adequate stimulus material (see Ham & Vonk, 2003). With ambiguous sentences, the simultaneous activation of trait and spontaneous inferences can be explored, without the ambiguities related with effects of presentation order of the category label, the trait-implying behavior, and the situational content.

EXPERIMENT 4⁶

Probe Recognition with Trait/Situational Ambiguous Behaviors

⁶ This study was conducted in collaboration with Kaat Van Acker

In this experiment, we followed the same procedure as in previous studies but with different stimulus materials. In this case, participants were presented with behavioral descriptions that could imply both a spontaneous trait inference and a spontaneous situational inference. For example, the description: “The person doesn’t like to go to class” might imply both the trait “lazy” as well as the situation “boring”. Using this type of material is particularly suitable to our goals because the sentences have a similar probability of eliciting a trait and a situational spontaneous inference. At the same, with this material we avoid problems related with order of presentation of the different constituents of the sentence, being impossible for participants to pay more attention to one of the parts of the sentence.

The same type of behavioral material was used by Ham and Vonk (2003). Applying the recognition probe paradigm, Ham and Vonk followed each ambiguous description either by the implied trait, the implied situational word, or by a non-related control word. Results showed that participants took more time to reject both the trait and the situational word than non-related control words. These results were taken as evidence that the same behavior may simultaneously elicit STIs and SSIs.

In our studies, the same ambiguous behavioral description (e.g., “doesn’t like to go to class”) was presented with a category that could be either consistent (e.g., “The party kid”) or inconsistent (e.g., The PhD student) with the behavior. Thus, our experiment may be seen as an integration of the studies by Wigboldus et al. (2003) and by Ham and Vonk (2003). The hypotheses were the same as in the two previous studies. When the behavior is presented with a consistent category, participants would be less likely to make STIs but more likely to make SSIs, compared to a condition in which the same behavior is presented with an inconsistent category.

Method

Participants and Design

Participants were 62 undergraduate students at University of California, Santa Barbara (52 women and 10 men). The critical trials of the experiment formed a 2 (social category: consistent vs. inconsistent vs. neutral) X 2 (probe word: trait vs. situation) within-subjects design.

Stimulus Material

Experimental Trials. . Initially, we developed series of behavior descriptions that were likely to imply both a trait and a situational inference. These sentences were added to the sentences used by Ham and Vonk (2003) and pre-tested with 42 American students (sentences used by Ham and Vonk were originally used in Dutch). Behavioral descriptions were presented two times, once followed by the implied trait and the other followed by the implied situation. In the trait condition, participants were asked “how applicable is the trait for the sentence” while in the situational condition they were asked “how applicable is the presented situation property for the sentence”. Behaviors were evaluated in a 7-point scale, ranging from 1 (“not applicable”) to 7 (“very applicable”). Some participants evaluated all the behaviors first on the trait-scale and then on the situation-scale, while others evaluated the behaviors first on the situation-scale and then on the trait-scale, in a counterbalanced way. Based on this pre-test, 10 experimental sentences were selected that lead to more similar evaluations in both trait dimensions ($M = 5.02$) and situational dimensions ($M = 5.26$), ($p > .05$) (see behavioral descriptions in Appendix D).

For each sentence, two category labels were selected that were either stereotype-consistent or stereotype-inconsistent with the behavior. In the pilot study, participants were also asked ($N = 42$) to indicate “to what extent the sentence was stereotype inconsistent or consistent with the actor”. The actor presented could be consistent, inconsistent, or neutral (i.e., “The person”) with the description. Participants were asked to indicate their responses in a 7-point scale, ranging from 1 (stereotype-inconsistent) to 7 (stereotype-consistent). The point 4 corresponded to a stereotype-neutral evaluation. Evaluation means showed significant differences in the extent to which sentences were judged stereotype consistent ($M = 5.48$), stereotype-neutral ($M = 4.03$), and stereotype-inconsistent ($M = 2.60$), in the expected direction [$F(2, 18)=37.50, p < .001$].

Based on pre-test results, 60 experimental trials were created by matching each of the 10 behavioral sentences with each category dimension (consistent, inconsistent, and neutral) and probe word (trait, situation).

Filler Trials. In order to prevent only “No” responses, the 10 experimental sentences were combined with the consistent or with the inconsistent label, and tested

with two probe words actually included in the sentences (the verb and other word), creating a total of 40 trials. Finally, in order to prevent participants of associating trait probes with “No” responses, trials were included in which the trait was included in the sentence, so participants should respond “yes”. These behaviors were composed by the category labels presented in the experimental trials and a neutral behavior (e.g., the lazy party kid/PhD student took the bus home from school). In total, 20 trials with the trait included were presented to participants.

Procedure

Stimuli were presented and responses were collected using the Direct RT software (Jarvis, 2002). We followed the same procedure as in the two previous experiments, with the following modifications. Because sentences were shorter, in this case they stayed on the screen for 1000 ms. In total, participants went through 120 randomly presented trials. Half of the trials elicit a “no” response and half elicited a “yes” response.

Dependent Measures

The dependent measures were the percentage of error, and response times for correct responses.

Results

Error Rates

The percentage of incorrect responses was analyzed in a 3 (social category: consistent vs. neutral vs. inconsistent) X 2 (probe word: trait, situation) within-subjects ANOVA. A main effect emerged for probe type, $F(1, 58) = 6.14, p = .016$, revealing a higher number of incorrect responses for situational probes ($M = 4.01$) than for trait probes ($M = 2.48$). Although the predicted interaction between social category and probe word was not significant ($F > 1$), a planned comparison revealed that the number of errors for situational probes was higher in the inconsistent-category condition ($M = 5.76$) than in both the neutral-category ($M = 2.54$) and consistent-category conditions ($M = 3.73$), $t(58) = 2.03, p = .023$, one tailed. Percentage of incorrect responses for situational words didn't differ significantly between consistent and neutral conditions (t

> 1). Results are presented in Figure 5. For trait probes, despite means going in the predicted direction, with more errors for the category consistent ($M = 3.72$) than for neutral ($M = 2.03$) and inconsistent category conditions ($M = 1.86$), planned comparisons revealed that differences were not statistically significant, $p > .05$ (one tailed).




Figure 5. Percentage of errors as function of probe type and category dimension

Response Times

A cutoff criterion was used for responses faster than 200 ms and longer than 2000 ms (Ratcliff, 1993). Based on this criterion, 5 participants were removed from the analysis.

Response times for correct responses were submitted to a 3 (social category: consistent vs. inconsistent vs. neutral) X 2 (probe word: trait, situation) within-subjects ANOVA. Only a main effect emerged for probe type, $F(1, 58)=10.08$, $p = .002$, showing that participants overall took more time to reject situational probes ($M = 871$ ms) than trait probes ($M = 837$ ms). No other significant effects emerged from the analysis.

Discussion

In the present experiment trait/situational ambiguous behavior descriptions were used. The same description could elicit a STI as well as a SSI. Results showed that the

percentage of incorrect responses for situational probes was higher in the inconsistent category condition than in the consistent or neutral category conditions. Thus, SSIs were more likely to occur when the category is inconsistent with the behavior. This is in line with our view that spontaneous inferences tend to be facilitated when they increase local coherence of the processed information. Our results show that, for example, upon observing the behavior “the dietician eats three plates of French fries” people are more likely to spontaneously infer the situational property “delicious” than when the same behavior is performed by a “soccer fan”. Apparently, when the actor is inconsistent with the behavior, participants are more likely to spontaneously infer something about the situation in order to comprehend what is being processed in a coherent way.

With regard to STIs, consistent with our hypothesis and with results from our previous studies, results revealed a stronger tendency for STIs in the consistent condition than in the neutral and inconsistent conditions. However, differences in this case didn't reach statistical significance.

Overall, results from this experiment showed that when ambiguous trait/situational behaviors are presented, significant differences as a function of the activated category are observed only for SSIs. Actually, Ham and Vonk (2003), using a similar behavioral material (but without the category labels), found differences in error rates only for situational probes and not for trait probes (although differences in response times were found both for STIs and SSIs). That is, participants were more likely to falsely recognize situational words after seeing experimental than control sentences, but no differences emerged for trait probes. Taken together, our results and Ham and Vonk (2003) results suggest that SSIs are more likely to occur with this type of behaviors, and are also more likely to be influenced by the type of category label provided. This aspect is interesting because it suggests that different types of behaviors tend to elicit different patterns of STIs and SSIs, which is in agreement with a more flexible view of the occurrence of spontaneous social inferences. Also, it confirms the necessity of incorporating behaviors not so highly illustrative of personality traits within STI research.

Another major consequence of our study is that it helps to clarify whether STIs are especially prevalent in the processing of social information, or whether other social inferences (e.g., situational inferences) might also occur spontaneously. Our results

show that situational inference may occur spontaneously, especially when behaviors are clearly pre-tested to allow for SSIs, and when behaviors are performed by an inconsistent actor.

Discussion of Experiments 2 – 4

In Experiments 2-4, instead of using behaviors highly illustrative of one personality trait as in previous STI research, we presented behavior descriptions that could imply a STI as well a SSI. In addition, instead of presenting the actor as an abstract entity we explicitly provided relevant information about the social category of the actor. With this material, effects of local coherence could be explored. Another innovation of these experiments was the possibility of examining not only the occurrence of STIs but also of SSIs.

According to the principle of local coherence, we predicted that both STIs and SSIs should occur in line with the basic necessity of our comprehension system of achieving local coherence. In general, results from the 3 reported experiments support our predictions. The social category associated with the actor influences both the magnitude of STIs as well as of SSIs, with the influences observed always in line with a coherent picture of what is being processed.

Specifically, the first two studies support the view that STIs are more likely to occur when the actor belongs to a consistent social category than when the actor belongs to a neutral or inconsistent social category. Results from Experiment 3, despite not significant, go in the same direction. Thus, STIs are not an unconditional process, but are clearly influenced by the type of actor that performs the behavior. These results are in line with previous findings from Wigboldus et al. (2003), but apply more complex social material.

In terms of SSIs, the first two experiments showed that SSIs tend to be inhibited when the behavior is performed by a consistent-category actor, *and* that SSIs tend to be facilitated when the behavior is performed by an inconsistent-category actor (although evidence for facilitation effects was weaker). In Experiment 3, results showed stronger evidence for the facilitation of SSIs by an inconsistent actor. The differences observed between studies are probably due to the different types of behaviors applied. Importantly, independently of the differences in terms of inhibition/facilitation effects

between studies, the observed pattern of results is always in agreement with a greater local coherence of the information processed.

The different behavior descriptions used in the different experiments may also explain the stronger effects of category dimension on STIs than on SSIs observed on Experiments 2 and 3, but not in Experiment 4. While in the first two experiments (Experiments 3 and 4) behaviors were selected to highly and unambiguously imply a trait, with situational information added a posteriori, in the Experiment 4 the same behaviors could equally imply a trait or a situational inference.

One question that is not totally clear are the different patterns for accuracy and response times obtained across studies. Previous studies that applied the recognition probe paradigm have also produced contradictory patterns. Predictions in the probe recognition paradigm are made both for error rates and for response times. However, the patterns of results are not consistent. In some cases differences are obtained only in errors and not in response times (Uleman et al., 1996, Experiment 1), in other cases differences are obtained in response times but not in errors (e.g., McKoon & Ratcliff, 1986; Wigboldus et al., 2003, and still in others differences are observed in both errors and reaction times (Ham & Vonk, 2003). It is puzzling to understand exactly why these different patterns emerge. The possibility exists of a speed-accuracy trade off (Reed, 1973). That is, people may respond faster with less accuracy, or slower with more accuracy. In line with this hypothesis, McKoon and Ratcliff (1986) found differences in error rates when a deadline procedure was applied (i.e., forcing participants to respond in a short interval of time). Apparently, forcing participants to respond faster made them sacrifice accuracy, making differences in errors more evident. However, the same pattern was not replicated by Uleman et al. (1996). In this study, motivation for participants to respond quicker was introduced by providing feedback on speed and accuracy after each trial (Experiments 2 and 3). Results showed that this manipulation reduced reaction times as expected, but at the same time *eliminated* differences in accuracy. Thus, the reasons for inconsistent patterns between response times and error rates in the probe recognition paradigm are not completely understood yet.

In sum, our results show that both STIs and SSIs are influenced by the type of actor that performed the behavior. These experiments are also important in showing that SSIs also occur spontaneously and may even be stronger than STIs, especially when

behaviors allow situational inferences to occur, and when the social category of the actor is inconstant with the behavior. Thus, our results support the view that trait inferences are not the only social inferences occurring spontaneously, and serve as an indication that different social inferences may be guided by similar rules. In this case, we specifically showed that both STIs and SSIs are influenced by the social group of the actor.

EXPERIMENT 5

Cued Recall Paradigm with Pairs of Behaviors

In the present experiment our goal was to explore the principle that the occurrence of STIs is guided by local coherence requirements by applying a different methodology. According with this principle, STIs that are made from a sentence should be largely dependent on its relation of coherence with previous text information. If previous text information supports the trait inference, then it will be more likely to occur. By contrast, if previous information from the text is not in agreement with the trait inference, it will be less likely to occur. In other words, the likelihood of making a trait inference should be function of the amount of previous information that supports the inference. This view is in line with the main theoretical perspectives within text comprehension literature (Kintsch, 1988; McKoon & Ratcliff, 1992).

Based on these formulations, we can predict that STIs are not only influenced by the activation of strong mental representations (i.e., stereotypes), as observed in our previous experiments (and also by Wigboldus et al., 2003), but also by previous behavioral information known about the actor. In order to test this prediction, we presented pairs of behaviors about the same actor. That is, instead of presenting only one behavior about each actor, as is usually the case within STI research, two trait implying behaviors were used to describe each actor. Local coherence was manipulated by varying the trait consistency between the two behaviors. In some trials, the two behaviors implied the same personality trait (consistent pairs), while in other trials the two behaviors implied opposite personality traits (inconsistent pair condition). In line with a coherence view, our general prediction is that STIs would be more likely to occur after observing consistent behavioral information about the same actor than after observing inconsistent behavioral information about the same actor.

In this experiment, we applied a new version of a cue-recall paradigm. During encoding, participants were presented with pairs of trait-implying behaviors describing different actors, under memory instructions. Each actor was described by two different behaviors. Some of the actors were described by pairs of consistent behaviors (i.e., they imply the same personality trait), while others actor were described by pairs of inconsistent behaviors (i.e., they imply opposite personality traits). After a distracting task, participants were asked to recall all the second behaviors used to describe each individual (*only* the second behaviors presented) under two different cue-recall conditions. Half of the participants received the traits implied by the second behaviors

as cues to recall, while the other half of the participants received the first behavior of each pair as a cue to recall.

As previously said, we predicted that participants would be more likely to spontaneously infer the trait from the second behavior when it is preceded by a consistent behavior than when it is preceded by an inconsistent behavior. Moreover, we make additional predictions about the encoding of incongruent behaviors. It is well established in the person memory literature that behaviors that are inconsistent with an existing expectancy are more difficult to be spontaneously encoded in trait terms (e.g., Jerónimo, 2007). In addition, there is evidence that inconsistent behaviors tend to reactivate previous received information, probably in an attempt to make sense of the received information (Sherman & Hamilton, 1994). In these studies, however, participants have explicit impression formation instructions, and have strong previous expectancies about the actor. Yet, we may hypothesize that similar effects can be observed even under the present study conditions. That is, because inconsistent behaviors are more difficult to comprehend, they will be compared with previous information. As a consequence, they will become stronger associated with the previous received behavior. Thus, we assume that the establishment of associations between behaviors may occur in the natural course of comprehension, independently of impression formation goals or strong previous expectancies. This prediction is supported by text comprehension formulations (Baker & Anderson, 1982; van der Broek, Risdén, Fletcher, & Thurlow, 1996), according to which coherence breaks result in the revision of previous information in an attempt to make sense of the new information.

Since we predict that behaviors are more likely to be spontaneously encoded in trait terms on consistent trials, our hypothesis was that on consistent trials second behaviors would be better recalled with trait cues than with first-behavior cues. By contrast, on inconsistent trials, because we predict that second behaviors are less likely to be encoded in trait terms and more likely to become associated with the previous behavior, our hypothesis was that second behaviors would be better recalled with first-behaviors cues than with trait cues.

Method

Participants and Design

Participants were 62 undergraduate students at University of California, Santa Barbara (51 women and 11 men). The critical trials of the experiment formed a 2 (Type of pairs: consistent vs. inconsistent) X 2 (cue condition: trait vs. first behavior) X 4 (Material version: version 1 vs. version 2 vs. version 3 vs. version 4). The type of pair variable was manipulated within-subjects, and the other variables were manipulated between-subjects.

Stimulus Material

We selected six pairs of behaviors that imply the same personality trait, six pairs of behaviors that imply opposite personality traits, and four pairs of neutral behaviors from norms compiled by Stroessner (1989). Because the manipulation of “type of pair” was within participants, different implied traits were chosen for the consistent and for the inconsistent behavior pairs (Trait Version I and Trait Version II, see Appendix E). The neutral pairs were the same in both “trait versions”. In addition, the order of behaviors within each pair was varied between participants. For example, if some participants saw the pair Friendly/Unfriendly, other participants saw the pair Unfriendly/Friendly. Combining the type of trait version with the order manipulation, 4 different version of material were used in total. The behaviors and the corresponding implied traits are presented in Appendix E.

Procedure

The experiment was run using the software MediaLab (Jarvis, 2002). Participants were provided with memory instructions, being informed that a series of pairs of sentences would be presented on the screen and that their task was to memorize them. After the initial instructions, the series of pairs of behaviors was randomly presented. Each sentence was presented individually in the screen for 6 seconds. The two sentences of each pair were matched with the same proper name (e.g., John) in order to differentiate the pairs of behaviors. The two sentences of each pair were presented sequentially, and the pairs of sentences were randomly presented. In the total,

16 pairs of behaviors were presented: 6 consistent pairs (the two behaviors implied the same personality trait), 6 inconsistent pairs (the two behaviors implied opposite personality traits), and 4 filler neutral pairs (formed by 2 neutral behaviors). Filler neutral pairs were included in order to dissipate the nature of the task, and also to avoid that the number of cues provided was used as a strategy by participants to count the number of behaviors that they should recall. After all pairs of behaviors were presented, participants completed a distracting anagram task for 4 minutes. After an anagram distracting task, participants were asked to remember as many of the second behaviors presented in each pair as they could. Instructions highlighted the fact that only the second behaviors of each pair should be recalled. Participants were provided 5 minutes to complete the recall task. During recall, half of the participants were provided with the traits implied by the second behaviors as recall-cues (trait-cue condition) and the other half of the participants were provided with all the first behaviors presented in each pair as recall cues (first-behavior cue condition). Participants were told that they should use the information provided as cues to behaviors recall. In both cue-conditions, the cues remained on the screen the entire time during the recall task. At the end, participants were debriefed and thanked for their participation.

Dependent Measures

As dependent measures we used the percentage of consistent and inconsistent second behaviors recalled from the total of behaviors recalled by each participant (which also included behaviors from neutral filler pairs).

Results

Behaviors recalled were categorized by an experimenter blind to the manipulations of the study. A factorial ANOVA showed no effects that included the material version variable ($F < 1$). This variable was not further considered.

Percentage of second behaviors recalled were entered in a 2 (Type of pair: consistent vs. inconsistent) X 2 (cue condition: trait vs. first behavior) mixed ANOVA. A main effect for type of pair was observed, revealing better recall of second behaviors on inconsistent trials ($M = 57.71$) than on consistent trials ($M = 39.41$), $F(1, 62) = 12.75$, $p < .001$, which suggests that in general participants recall a higher number of

second behaviors when they are inconsistent with the first behavior presented than when they are consistent with the first behavior presented.

However, and more important to our predictions, the main effect for type of pair interacted with cue condition, $F(1, 62) = 3.76, p = .057$. Planned comparisons revealed that for consistent pair trials, second behaviors were better recalled with trait-cues ($M = 43.80$) than with first-behaviors cues ($M = 35.01$), $t(62) = 1.64, p = .053$, one tailed (see Figure 6). This suggests that behaviors that are consistent with a previously presented behavior are more likely to be encoded in terms of the underlying traits. Consequently, they don't need to be compared with previously received behavioral information.




Figure 6. Percentage of errors as function of probe type and category dimension

For inconsistent pairs, second behaviors recall was better with first-behaviors cues ($M = 63.26$) than with trait-cues ($M = 52.16$), $t(62) = 2.16, p = .03$, one tailed (see Figure 6), which supports the idea that when second behaviors are inconsistent with a previous behavior they are less likely to be spontaneously encoded in terms of traits. Because of that, they are likely to be compared with previous received information, resulting in the stronger associations between the two behaviors.

Discussion

Results from this experiment showed that for consistent pairs of behaviors (i.e., both behaviors imply the same trait), recall of the second behaviors is better when the

underling traits are provided as cues than when the first-behaviors are provided as cues. In contrast, for inconsistent pairs of behaviors, recall of the second behavior is better with first behaviors cues than with trait cues. This dissociation pattern is consistent with the view that second consistent-behaviors tend to be spontaneously encoded in trait terms, while second inconsistent-behaviors are less likely to be encoded in trait terms and are stronger associated with previous received behaviors.

Thus, these data provide initial evidence that STIs are not only influenced by the activation of strong mental structures (e.g., stereotypes) but the presentation of merely *one* previous behavior is enough to influence the magnitude of subsequent spontaneous trait inferences. According to our data, STIs are more likely to occur when they are consistent with previously received information, and less likely to occur when they are inconsistent with previous information. This supports the principle that STIs are regulated by the basic need of our comprehension system for local coherence. According to this view, STIs are not only inhibited or facilitated in order to preserve existing stereotypes but work in line with the comprehension process in a broader way. In this sense, STIs should be facilitated when the trait inference is in line with a coherent narrative about the information received, and should be disrupted every time the trait inference contributed to increase the incoherence of the narrative.

Although our results are consistent with our hypothesis, there are some criticisms that can be made to this study. First, this study is not conclusive about whether STI effects occurred during encoding. In order to prove the on-line status of the observed effects, we should apply on-line inferential measures. Second, there is the danger that participants might be engaging in explicit impression formation processes. Although we have provided participants with memory instructions, the way information was presented might have instigated impression formation processes. Each behavior was presented for 6 seconds, both sentences in a pair referred to the same actor, and each pair of behaviors made reference to a name of a different actor. The social nature of the task might become apparent under these conditions. In order to clarify these aspects, in Experiment 6 we used the recognition probe paradigm with pairs of behaviors.

EXPERIMENT 6

Probe Recognition Paradigm with Pairs of Behaviors

The aim of the present experiment was to better clarify some of the questions raised by our previous study. Because in the previous study the recall task was only performed a long time after encoding, it is not clear whether differences in terms of spontaneous trait inferences really occurred during behavior encoding. Further, the fact that behavior presentation time was quite long (i.e., each behavior was presented for 6 seconds) might have elicited an explicit impression formation processes. In order to rule out these explanations for the effects on STIs when pairs of behaviors are presented, in the present experiment we applied an online measure of inferential activity, the recognition probe paradigm (McKoon & Ratcliff, 1986).

In the recognition probe paradigm, participants are presented with a trait-implicating behavior for a very brief time, and immediately afterwards the trait implied by the behavior is presented. Participants are instructed to indicate, as fast as possible, whether the trait was included in the sentence. This procedure is appropriate to rule out the two alternative explanations of our previous results. First, because the trait is presented immediately after the behavior and a fast response is required, there is not much room for deliberative retrieval processes to occur. Second, with this procedure, the more participants engage in trait inferences the more those inferences interfere with a correct answer. While in the previous experiment explicit trait inferences might be beneficial for performance, in this case engaging in explicit trait inferences work against accurate performance. Thus, it is not understandable why participants would engage in impression formation processes when doing so would be prejudicial to their performance.

In this experiment the recognition probe paradigm was applied with pairs of behaviors describing the same actor. For example, on one of the trials the second behavior was “John solved a complicated mathematics problem in his spare time”, a behavior that implies the trait inference “Intelligent” under the conditions in which STIs have been usually studied. In the consistent trials, this behavior was preceded by a behavior that implied the same trait, for example, “John can speak three different languages fluently”. In contrast, in the inconsistent trials the second behavior was preceded by a behavior that implied the opposite trait, for example, “John was visibly confused by the map in the subway station” (a behavior that implies the trait “stupid”).

Our hypothesis was the same as in the previous study. We predicted that participants would be more likely to infer the trait from the second behavior when it is preceded by a consistent behavior than when it is preceded by an inconsistent behavior. In terms of the recognition probe paradigm, this would be reflected in a greater difficulty in rejecting the trait on consistent than on inconsistent trials.

In addition, in this study we also added a control condition in which the second behavior was preceded by a neutral behavior, for example, “John ate a bacon and cheese hamburger for lunch”. This condition is important in order to clarify whether differences on STIs are due to facilitation effects on consistent trials, inhibition effects on inconsistent trials, or both.

Method

Participants and Design

Participants were 75 undergraduate students at University of California, Santa Barbara (54 women and 21 men). The critical trials of the experiment formed a 3 (Type of pair: consistent vs. inconsistent vs. control) X 2 (replication: order 1 vs. order 2) design. The first variable was manipulated within-subjects and the second variable was manipulated between-subjects.

Stimulus Material

Experimental Trials. We selected 8 behaviors illustrative of different personality traits from existing norms compiled by Stroessner (1989). On consistent pairs each behavior was preceded by one behavior that implied the same personality trait, on inconsistent pairs each behavior was preceded by one behavior that implied the opposite trait, and on neutral pairs each behavior was preceded by a neutral behavior (see Appendix F). The combination of the 8 behaviors with the different preceding behaviors (consistent, inconsistent, or neutral) resulted in 24 experimental pairs of behaviors: 8 consistent pairs, 8 inconsistent pairs, and 8 neutral pairs. The probe words presented on the experimental trials were always the traits implied by the second behavior of each pair.

The order of behaviors within each inconsistent pair was manipulated between participants, resulting in two replications of the stimulus materials. That is, if some participants were presented with the inconsistent pair “stupid behavior – intelligent behavior”, other participants were presented with the same behaviors in the reverse order “intelligent behavior – stupid behavior” (see Appendix E for the list of the stimuli used). For both replications, neutral pairs were formed by pairing the second behavior of each pair with the same neutral behavior (in the first case, “neutral behavior – intelligent behavior” and in the second case “neutral behavior – stupid behavior”). Consistent pairs were formed by pairing each second behavior with an additional behavior that implied the same trait (in the first case intelligent behavior II – intelligent behavior and in the second case stupid behavior II – stupid behavior).

Filler Trials. In order to prevent only “No” responses, the same 24 experimental pairs were tested with probe words actually included in the second sentence. In these trials, a “Yes” response was required. In addition, in order to prevent that traits elicit only “No” responses, 10 filler pairs were created in which the trait was included in the second sentence. These trait fillers were created by associating neutral behaviors with new actors (e.g., “The intelligent Ryan read the sports page” – probe “intelligent”). The traits used as probes in these trait filler sentences were the same as the experimental trait probes, and more two additional new traits (ecological and athletic). Finally, the same 10 filler pairs were presented without the trait and tested with a probe word not included in the sentence (e.g., “Ryan read the sports page”, tested with the probe “wool”). The words that served as probes were taken from the first sentence of the pair.

Procedure

We followed a similar procedure as in Experiments 1-4. The study was presented as a “reading comprehension study”. Participants were initially told that we were interested in the way people read and comprehend sentences. They were then informed that pairs of sentences will appear on the screen and that after each pair of sentences a word written in blue will be presented in the center of the screen. Their task was simply to indicate whether that word was included on the second sentence of the pair of sentences they have just read. After instructions, three training pairs were

provided so participants could get familiar with the task. After the training trials, participants were asked if the task was clear and if so the experiment started. The two initial trials were also filler trials, to avoid start-up problems. All pairs of behaviors were presented following the same sequence. The first sentence in the pair was presented for 1500 ms, followed by a white screen presented for 500 ms, and then by the second behavior also presented for 1500 ms. After that a white screen was presented for 1000 ms in order to better differentiate the different trials, and then the next trial followed.

In total, participants went through 68 trials. The 68 pairs were divided in three blocks (the first block composed by 22 pairs of behaviors, and the second and third blocks composed by 23 pairs of behaviors). Each block was composed of the same number of experimental and filler trials. Presentation of the three blocks was randomized between participants, as well as the order of the pairs of behaviors within each block. In total, participants went through 68 pairs of behaviors, 34 pairs of behaviors eliciting a “No” response, and 34 eliciting a “Yes” response.

Dependent Measures

The dependent measures were the percentage of errors and response times for correct responses.

Results

Error Rates

There were no significant effects for material replication, either in errors or in response times ($F < 1$). This variable will not be further considered.

The percentage of incorrect responses was analyzed in a 3 (type of pair: consistent vs. control vs. inconsistent) within-subjects ANOVA. The ANOVA revealed a marginal effect for type of pair, $F(2, 140) = 2.04, p = .13$. As predicted, a planned comparisons between the means revealed that the number of errors was lower in the inconsistent condition ($M = 4.57$) than in both the control ($M = 5.63$) and consistent conditions ($M = 6.88$), $t(70) = 1.93, p = .029$, one tailed. In addition, planned comparison between the consistent condition and the other two conditions showed that the number of errors was marginally superior for consistent pairs than for control and

inconsistent pairs, $t(70) = 1.58, p = .06$, one tailed. The error rate means are presented in Table 2.

Table 2

Means of Errors as a Function of Type of Pair

	Type of Pair		
	Consistent	Control	Inconsistent
<i>M</i>	6.88	5.63	4.57

Response Times

A cutoff criterion was used for responses faster than 200 ms and longer than 2000 ms (Ratcliff, 1993). Based on this criterion, 4 participants were removed from the analysis. Response times for correct responses were submitted to a 3 (type of pair: consistent vs. control vs. inconsistent) repeated measures ANOVA. No significant effects emerged from this analysis, $F(2, 140)=1.05, ns$. Thus, time to correctly reject the probes was not significantly different between different types of pair conditions. Thus, no differences were obtained in response times.

Discussion

Using an online measure, the present findings give support to the view that STIs are less likely when the behavior is preceded by a trait-inconsistent behavior than when it is preceded by a trait-consistent behavior. Participants were less likely to incorrectly indicate that the trait was included in the sentence when the sentence was preceded by a trait-inconsistent behavior than when the same behavior was preceded by a trait-consistent or neutral behavior and more likely to indicate that the trait was included in the sentence when it was preceded by a consistent behavior. The fact that participants in the consistent trials infer the trait more strongly explains why they are more likely to

make false recognitions and responding as if they have actually seen the trait in the sentence.

The present data suggests that the STI process can be inhibited or facilitated depending on the behavior that is previously presented. The pattern of STI inhibition and facilitation observed is in line with the view that the occurrence of STIs is modulated by coherence requirements. When previously received information is coherent with the trait inference, STIs are more likely to occur. When previous information is incoherent with the trait inference, STIs are less likely to occur.

Thus, overall these results support the intuitive idea that when we see someone performing an intelligent behavior and subsequently acting in a stupid way, it is less likely that we spontaneously infer the trait “stupid” than when the person acts consistently in a stupid way. In the latter case, there are no contradictions. Inferring the trait is consistent with a coherent comprehension of the actor’s narrative, and represents the fastest way of making sense of the behavior. On the other hand, if the person acts differently in two different occasions, the comprehension process is affected by that variability. In this case, the trait is not so strongly inferred spontaneously, probably because inferring the trait might turn out to be erroneous for an effective comprehension of the behavior, as well as of the person.

Moreover, our data show that we don’t need to be intentionally forming impressions about the actor in order for these effects to emerge. The modulation of the trait inference process seems to be incorporated within the comprehension process itself. Specifically, the fact that the level to which one trait concept becomes active during the comprehension of one behavior is different, depending on the preceding circumstances, indicates that the behavior is initially differently comprehended.

Discussion of Experiments 5 and 6

In Experiments 5 and 6, instead of presenting only one behavior about each actor, as is typically the case within the STI literature, we presented pairs of behaviors about the same actor. Our goal was to provide converging evidence in favor of the principle that STIs are guided by local coherence requirements. If that is the case, STIs should not only be influenced by their consistency in relation to strong existing mental structures, such as stereotypes (see our Experiments 2-4, and also Wigboldus et al.,

2003), but also by their consistency with previous behavioral information about the actor. In order to test these ideas, we presented pairs that could be convergent in terms of trait implications, or could imply opposite traits.

In Experiment 5, the following reasoning was made. On consistent behavior pairs traits are more easily inferred, while on inconsistent behavior pairs traits are less likely to be inferred. At the same time, inconsistent pairs cause a coherence break on comprehension that leads to a revision of the previous behavior. Because of this revision process, behaviors will be more strongly associated with each other on inconsistent than on consistent pairs. According to this logic, we predicted that on consistent trials implied traits would be better cues for recall of second behaviors than first behaviors, but on inconsistent trials first behaviors would be better cues than the implied traits. Our results were in line with these predictions.

In Experiment 6 we used an online measure and obtained data convergent with the findings from Experiment 5. In this case, we observed that participants are less likely to incorrectly indicate that the trait was included in the sentence on inconsistent than on consistent and neutral trials. The online nature of this task supports the notion that the observed effects are unintentional and occurred during behavior encoding.

These two experiments represent an initial attempt to integrate narrative information into the study of STIs. Overall, they clearly prove the importance of using more meaningful contexts in order to study the nature and occurrence of STIs. STIs occur differently depending on whether we have other behavioral information about the actor or not. Importantly, the spontaneity of STIs seems to be dependent upon a coherent picture of the received information.

SECTION III

SPONTANEOUS TRAIT INFERENCES VARY IN A CONTINUUM OF STRENGTH

EXPERIMENT 7

Trait Cued Recall

According to our third general principle concerning the mechanisms underlying spontaneous trait inferences, it is stated that trait inferences vary in a continuum of strength. Inferences can be stronger or weaker encoded. In one of the points of the continuum, the trait underlying the behavior is activated to a lower extent, and may become associated not only with the actor of the behavior, but also with other element that are present in the context. In the other point of the continuum, the trait is strongly activated and is specifically linked to the actor of the behavior.

The closer the trait inference is to the lower end of the continuum, the less likely it would be that the trait can be deliberately and consciously accessed, leaving the effects of trait inferences largely implicit. In contrast, the closer trait inferences are to the highest level of the continuum, the more likely is that the link between the actor and the trait can be consciously accessed and verbalized to others (see Ferreira et al., in press). Similar to Ferreira et al. (in press), we argue that the closer trait inferences are to the highest point of the continuum, the more likely it is that individuals will have conscious access to the output of the inference.

One of the variables that governs the continuum is the intention to make a trait inference. When an explicit intention to form an impression about the target is present, a specific link is established between the inferred trait and the actor. In this case, the output of the inference can be accessed both during encoding and also during performance in subsequent tasks. By contrast, when no intention exists to infer the trait, the process of trait inference remains unnoticed during encoding. In this case, it is also less likely that individuals gain conscious access to the behavior-trait association in posterior tasks. Since the highest level of the continuum is reached under impression formation conditions, this is a useful condition in order to contrast with trait inferences made under unintentional processing conditions.

In line with this reasoning, in the present experiment we compared trait inferences under memory (i.e., spontaneous trait inferences) and impression formation (i.e., intentional trait inferences) instructions, in order to explore the extent to which participants can have access to the output of the inference under spontaneous study conditions.

We used a paradigm that combines the most crucial features of the Winter and Uleman (1984) and Hamilton et al. (1980) paradigms. Specifically, in this experiment

participants were presented with a list of twenty four behaviors. However, contrary to what is typical in STI paradigms, instead of presenting only one behavior for each trait category, several behaviors were presented that imply the same personality trait, following the procedure of Hamilton et al. (1980). Specifically, the 24 behaviors represented four different trait categories (6 friendly behaviors, 6 intelligent behaviors, 6 musical behaviors, and 6 athletic behaviors). Notice that presenting four behaviors implying the same personality trait is especially appropriate to the purposes of the present experiment because in this case conscious access during subsequent recall to the previously inferred traits would be highly relevant for the task at hand. After a distracting task, participants were asked to recall the entire list of behaviors, either with the four underlying traits provided as cues or with no cues.

There are several relevant aspects that can be tested within this paradigm. First, according with the view of several authors (e.g., Gordon & Wyer, 1987; Hamilton et al., 1980; Srull & Wyer, 1989), because impression participants have the intention to form a coherent impression about the target, behaviors will be encoded in trait terms and will be organized according with the superordinate trait structure. Integrating the different pieces of information into a coherent impression would be an efficient organizational strategy (Gordon & Wyer, 1987; Hamilton, 1981; 1989; Hamilton, Driscoll, & Worth, 1989; Hamilton, Katz, & Leirer, 1980; Hamilton & Sherman, 1996; Hoffman, Mischel, & Mazze, 1981; Park, 1989). In line with this theoretical framework, our first hypothesis is that impression formation participants will organize the different behaviors according to the underlying traits. As a consequence, impression participants will tend to recall behaviors implying the same trait together (i.e., impression participants will exhibit trait clustering). Thus, in this experiment the level of trait-clustering exhibited in recall is taken as a measure of the extent to which behaviors were organized by trait. By contrast, we predict that because memory participants infer traits spontaneously, they will not have conscious access to the trait inference process and will not organize the different behaviors according with the underlying traits. As a consequence, they will not exhibit trait clustering during behaviors recall.

A second hypothesis concerns the number of behaviors recalled in each processing condition. Because impression formation participants encode behaviors following an organized trait structure, they will tend to recall a higher number of

behaviors, compared with the number of behaviors recalled by memory participants. Our first two hypotheses are in line with previous results obtained by Hamilton et al. (1980), showing that participants asked to form impressions exhibit higher levels of clustering and higher recall than participants asked to memorize behaviors.

In addition to the previous predictions, specific hypotheses are formulated about the effects of providing the traits as cues during retrieval. Specifically, we predict that because impression participants have conscious access to the inferred traits, they will use those traits to facilitate recall independently of the traits being explicitly provided as cues. Thus, providing or not providing the traits will have little impact on recall for impression formation participants. In contrast, because memory participants have no access to the inferred traits during encoding, we predict that they will be also less likely to use the output of the inference process in subsequent tasks, *unless* traits are explicitly provided. In the case in which traits are provided, traits can recreate the encoding context (Tulving & Thomson, 1973). As a result, the existing behavior-trait association will be activated, and recall will be facilitated. This prediction is in line with existing results from the cued recall paradigm, showing greater recall in trait-cue than in no-cue conditions (Winter & Uleman, 1984). In sum, we hypothesize that no differences between no-cues and trait-cues recall will be found in impression formation conditions, whereas in memory conditions trait cues recall will be higher than no-cues recall.

Finally, specific hypotheses were made about the effect of cue conditions on the level of clustering. For impression formation participants, similar predictions are made. Cue condition will have no effect on the level of trait clustering regardless of traits being provided or not provided, because impression participants will use the traits to organize their recall either way. In this case, however, similar predictions are made for memory participants. Because memory participants do not extract a trait organization structure during encoding, even when they are provided with traits during recall they will not be able to organize recall by trait. Thus, in the memory condition we predict that trait-cues will result in a higher number of behaviors recalled, but that will not be reflected in a greater trait clustering. These predictions were tested in the present experiment.

Method

Participants and Design

Participants were 136 undergraduate students at the University of California, Santa Barbara, whose participation earned them partial credit for a psychology course. Participants were randomly assigned to the conditions of a 2 (processing goals: impression formation vs. memory) X 2 (cue condition: trait cues vs. no cues) X 3 (order of trait cues: random order1 vs. random order2 vs. random order3) between-participants factorial design.

Stimulus Material

The 24 behaviors and the corresponding implied traits were taken from existing norms compiled by Stroessner (1989). For a list of the stimuli used, see Appendix G.

Procedure

Stimuli were presented and responses were collected using MediaLab software (Jarvis, 2002). Participants were initially provided with impression or memory instructions, according to condition. Impression participants were asked to form an impression about the target while reading the behaviors. Instructions informed them that the twenty-four behaviors were performed by the same hypothetical target (named “John”). Memory participants were asked to memorize the sentences as accurately as possible, and no reference was made about “behaviors” or “actors”. Twenty-four behavioral sentences were then randomly presented. The list of sentences was composed of behaviors illustrative of four different personality traits (6 Intelligent, 6 Friendly, 6 Musical and 6 Athletic). Each sentence was presented for 6 seconds. After participants had read the behavioral descriptions, a 5 minute anagram filler task was given in order to eliminate short-term memory effects. Following the distracting task, participants were instructed to remember as many of the behavior descriptions as they could. The time for the free recall task was 5 minutes. In the “trait cues” condition, participants were told that in order to help them in the free recall task, four words would be provided. They should use those words as memory cues while recalling the sentences. The four words were the traits implied by the behaviors (intelligent, friendly,

musical, and athletic). The cue trait words were presented for 15 seconds in a different screen. Three different orders of cue presentation were randomly created. After fifteen seconds, the screen with the cues disappeared and participants were instructed to write the behaviors. In the “no cues condition” participants recalled the behaviors without any reference to cues. After finishing the experiment participants were fully debriefed and thanked for their participation.

Dependent Measures

As dependent measures we used the number of behaviors recalled and the level of clustering in recall. As a measure of clustering, we applied the adjusted ratio of clustering (ARC) measure (Roenker, Thompson, & Brown, 1971).

Results

Behavior descriptions were categorized by a coder blind to the manipulations and goals of the study. A factorial ANOVA showed no effects of cue order for both free recall and ARC scores. These results will be not further considered.

Free Recall

In order to analyse the effects of both processing goals and cues on number of behaviors recalled, results were analysed in a 2 (processing goal: impression formation vs. memory) X 2 (cue condition: trait cues vs. no cues) ANOVA. A significant main effect was found for processing goal, $F(1, 132) = 6.22, p=.014$, revealing that participants under impression formation instructions ($M = 10.79$) recalled significantly more behaviors than participants under memory instructions ($M = 9.02$). These results are in agreement with our predictions and replicate the pattern obtained by Hamilton et al. (1980). Table 3 shows both the free recall data and the ARC scores.

A significant main effect was also revealed for cue condition, $F(1, 132) = 3.87, p=.051$, showing that when trait cues were provided the number of behaviors recalled was significantly higher ($M = 10.60$) than when no cues were provided ($M = 9.21$).

Central to our hypotheses is the interaction between cues condition and processing goal conditions. We predicted no impact of cue condition for impression formation participants. In contrast, we predicted that trait cues would facilitate recall for

memory participants, in comparison to the no cues condition. Consistent with our hypotheses, and as shown in Table 3, participants given impression formation instructions recalled approximately the same number of behaviors with trait cues (M=11.06) and without cues (M=10.51), $t(124) = .55$, *ns* (one tailed planned comparisons). Thus, providing the traits made no difference in terms of recall for impression formation participants. In contrast, participants under memory instructions conditions recalled significantly more behaviors when trait-cues were provided (M= 10.14) than when no-cues were provided (M= 7.90), $t(124) = 2.19$, $p = .03$ (one-tailed planned comparisons). These results replicate Winter and Uleman's (1984) findings and support the notion that participants do not spontaneously access the inferred traits during recall. Traits only facilitate recall when are explicitly provided.

Table 3

Free Recall and ARC Scores as a Function of Processing Goal and Cues Condition

	Free Recall		Clustering (ARC)	
	No cues	Trait Cues	No Cues	Trait Cues
Impression Formation (n=132)	10,48	10,92	.22*	.24*
Memory (n=132)	7,86	10,15	.02	.11

Level of Clustering

In order to analyse the extent to which participants imposed a trait organization on the behaviors recalled, we used the adjusted ratio of clustering score (ARC) (Roener, Thompson, & Brown, 1971). The ARC varies between -1 and 1, and measures the degree to which a participant recalls the items on a list grouped together as a function of the underlying categories. Because behaviors were initially presented in a random order, if recall is organized by trait then it is taken as evidence that participants imposed an organization on the behaviors. The ARC score cannot be calculated when all items recalled are from the same category or when only one item is recalled from

each category. Also, the ARC score is a biased measure when the number of intrusions is very high. For these reasons, clustering analysis will be reported without the ARC index for 13 participants.

A 2 (processing goal: impression formation vs. memory) X 2 (cue condition: trait-cues, no-cues) ANOVA revealed a main effect of processing goal, showing that the level of clustering is higher under impression formation instructions (ARC=.23) than under memory instructions (ARC=.08), $F(1, 119) = 3.71, p=.056$. This pattern replicates the findings obtained by Hamilton et al. (1980).

The interaction between processing goal and cue condition was not significant, $F(1, 119) < 1$. As expected, planned comparisons revealed that under impression formation conditions cues made no difference in the level of trait clustering (ARC_{nocues}=.22 and ARC_{cues}=.24), $t < 1$, (one-tailed). Under memory conditions, despite trait cues tending to increase the level of clustering observed (ARC_{nocues}=.04 and ARC_{cues}=.11), planned contrasts revealed that the difference was not significant, $t < 1$.

Further analyses were computed in order to examine whether the observed level of clustering differed from zero under both processing conditions. For impression formation participants, the mean ARC scores differed significantly from zero ($\neq 0, p < .01$), both in the cue (ARC_{cues}=.24) and in the no cue condition (ARC_{nocues}=.22). This indicates that under impression formation conditions participants recalled behaviors clustered by trait, independently of being provided or not provided with the traits as cues. The organization of behaviors by trait supports the notion that impression formation participants have conscious access to the inferred traits during encoding. By contrast, for memory participants, the level of clustering was not significantly different from zero ($\neq 0, ns$) in both cue conditions (ARC_{nocues}=.04 and ARC_{cues}=.11). Thus, not even when trait cues were explicitly provided did memory participants recall behaviors clustered by trait. Probably, the reason is that participants cannot take advantage of a trait structure that was not mentally represented during encoding in the first place. That is, when trait inferences are spontaneously made during encoding they are not used to organize behaviors (our results suggest that such organization doesn't happen even at an implicit level). Therefore, when traits are later provided, they are no longer useful for clustering because participants cannot take advantage of an organizing structure that was not previously created.

Discussion

Results from Experiment 7 support the notion that the more participants have an explicit intention to make a trait inference, the more likely is that they can consciously access those trait inferences and make deliberative use of them in subsequent tasks.

Our findings showed that participants instructed to form an impression about a target organized behaviors presented in a random order according to the superordinate trait concepts. In contrast, participants instructed to memorize behaviors showed a level of trait clustering not different from chance level. Impression participants also recalled a higher number of behaviors than memory participants. These findings are in line with previous results (Hamilton et al., 1980) and support the idea that impression formation participants encoded behaviors in trait terms and organized them in trait categories (Gordon & Wyer, 1987; Hamilton et al., 1980; Srull & Wyer, 1989). At recall, traits served as retrieval cues for the behaviors emanating from each trait. Memory participants, by contrast, did not take advantage of the underlying trait structure in order to organize behaviors and thus did not benefit from trait organization during recall. The fact that memory participants did not cluster behaviors by trait might seem puzzling if we consider that traits are inferred spontaneously under memory conditions (Winter & Uleman, 1984). If participants under memory instructions spontaneously infer traits from behaviors, why do they not use the inferred traits to organize behaviors, especially when this would be an effective recall strategy for the type of material presented?

According to our predictions, this puzzle can be easily explained. Because under memory conditions, no reference is made to social material (i.e., no reference is made to impressions, behaviors, or actors), so the trait inference process occurs spontaneously. As a consequence, participants fail to gain conscious access to the output of the inference and are unable to apprehend the trait structure that underlies behaviors.

Results from our experiment regarding the difference between cue conditions are especially relevant concerning this point. Specifically, the comparison between the level of recall and clustering in the trait-cue condition and in the no-cue condition can be seen as a non-obtrusive measure of the degree of conscious access to the inferred traits during retrieval. For impression participants, results showed that providing or not providing the traits as cues makes no difference either in the number of behaviors recalled or in the level of trait clustering. This is in agreement with the idea that

impression formation participants consciously use previous trait inferences in order to guide recall, regardless of being explicitly provided with the traits. In other words, providing the traits is apparently redundant because participants would use them as a recall device anyway.

For memory participants, on the contrary, providing the traits as cues significantly increased the number of behaviors recalled. This suggests that under normal conditions of recall, when no cues were provided, participants did not spontaneously access previously inferred traits in order to facilitate recall. In this study, participants apparently did not have the ability to consciously access spontaneously inferred traits. Only when traits were explicitly provided did they facilitate recall, probably because they reinstated the encoding context (Tulving & Thomson, 1973) or because they activated the encoded behavior-trait association.

An interesting result concerns the effect of cues on level of trait clustering under the memory condition. As predicted, and contrary to the benefit observed in terms of number of behaviors recalled, providing traits as cues had no impact on the level of trait clustering. This means that memory participants did not organize their recall by trait, even when the traits were explicitly provided to them during retrieval. According to our theoretical framework, this can be explained by the fact that memory participants did not organize behaviors by traits during encoding, in the first place. Because of that, when traits were provided at retrieval participants were no longer able to use them in order to create an organized structure. This might be especially difficult in the present study, because under memory instructions no reference was made to an actor within the sentence, which could serve as a unifying structure. The absence of impact of trait-cues on level of clustering under memory instructions indicates that effects deriving from spontaneously inferred traits remain largely implicit. Thus, despite effects of STI being visible under some implicit tasks such as in the savings in relearning paradigm (Carlston & Skowronski, 1994), participants seem to be unable to deliberately making use of those spontaneous trait inferences in order to organize recall, even when they are explicitly provided with the traits.

In sum, the fact that trait-cued recall was higher than no-cued recall under memory conditions supports the notion that participants spontaneously inferred traits during behavior encoding (see Winter & Uleman, 1984). However, the total absence of

clustering under memory instructions indicates that traits inferred spontaneously have no function (not even at an implicit level) in terms of organizing received information. The fact that providing traits do not increase the level of clustering indicates that memory participants were unable to make more deliberative use of those inferences. By contrast, impression participants tend to organize recall by traits, and providing traits has no impact on levels of clustering or in recall. This clearly indicates that participants that are forming an impression not only have conscious access to the inferred traits but deliberately make use of those traits in subsequent tasks that involve the actor. These effects are totally independent of traits being provided at retrieval.

There are, however, some questions that need further clarification. First, as some authors note (e.g., Gordon & Wyer, 1987; Wyer & Srull, 1989) clustering of behaviors may not necessarily reflect superordinate trait associations, but instead be the result of direct associations between semantically similar behaviors. However, because behaviors were chosen to be semantically different, despite implying the same personality trait, this hypothesis seems unlikely. Also, the fact that the level of clustering for memory participants was not significantly different from zero can be taken as an indication that no strong semantic relations between behaviors existed. That is, because memory participants have no explicit intentions to conceive behaviors in terms of traits, they would be more likely to use other recall strategies as, for example, thinking of behaviors in relation to one another (Gordon & Wyer, 1987). This would lead to the creation of direct associations between behaviors that share semantic features, and would be reflected in terms of clustering. The fact that memory participants exhibited no cluster indicates that no strong relations existed between behaviors that implied the same trait.

A second issue that deserves our consideration is the question about whether effects observed occur during encoding or during retrieval. For example, it is not clear whether trait clustering under impression formation conditions is the result of organizational processes that occur at the time behaviors are learned, or whether behaviors are organized by trait only during retrieval. Klein and Loftus (1990), for example, argue that trait clustering in recall can be produced merely by retrieval processes. Thus, although it is intuitively plausible that organizational processes take place as participants form impressions, and this idea is consistent with existing formulations (e.g., Gordon & Wyer, 1987; Hamilton et al., 1980; Srull & Wyer, 1989),

it is possible that organizational processes also occur during retrieval. Whichever the case, it is important to note that, for our purposes, the most crucial aspect is that under impression formation participants have conscious access to previously inferred traits. This is evidenced by the fact that trait-cues have no impact on the level of recall and clustering. Whether organization of behaviors by trait occurs during encoding, retrieval, or both is not totally clear from these data.

An additional criticism that that be made to this study is that the possibility exists that participants automatically activate impression formation goals, even under memory conditions (Chartrand & Bargh, 1996). Chartrand and Bargh (1996) make this point clearly by asserting that certain situational features may automatically trigger goals that tend to be associated with them. Applying this reasoning to our study, we may question whether reading a list of behaviors in which many of them imply the same trait would not automatically make participants think about impressions. There are some reasons, however, against this possibility. First, behaviors were presented without any reference to an actor. Thus, there were no apparently reasons why participants would think that presented sentences were all performed by the same person, and because of that engage in impression formation processes about an imaginary target. Second, another reason that favours the view that memory participants did not engaged in impression formation processes are the differences observed in our results between memory and impression formation conditions. Specifically, if memory participants were forming impressions then both the level of recall and trait clustering under no trait cue conditions should be similar to the one observed under impression conditions. That was not the case, suggesting that memory participants didn't engage in thoughtful impression formation processes. Despite these aspects, the possibility that contamination impression formation processes intervened under memory conditions cannot be completely ruled out. This is a major problem within STI research, which might be especially difficult to eliminate when different behaviors implying the same trait are presented, as in the present experiment.

Also problematic is the fact that, as some authors have pointed out, (Corbett & Doshier, 1978; D'Agostino & Beegle, 1996; McKoon & Ratcliff, 1992; Wyer & Srull, 1989), our data are not conclusive about whether the advantage of trait-cued recall over no-cued recall for memory participants is due to spontaneous trait inferences occurring

during encoding, or to processes taking place only during retrieval. Despite the fact that we used a between cue manipulation, considered to be a more appropriate manipulation (D'Agostino & Beegle, 1996), it is still possible that traits provided at retrieval facilitate behavior recall by prompting participants to generate information that in turn make it easier to recall behaviors. Thus, no definitive arguments can be made about whether STIs occurred during encoding in the present study, and data from converging paradigms should be applied in order to clarify this point.

EXPERIMENT 8

False Recognition and Reading Times

According with the principle that trait inferences vary in terms of encoding strength, it is assumed that two characteristics may differentiate between trait inferences that are weakly encoded from strongly encoded trait inferences. One of the characteristics is the ability to deliberately use trait inferences in subsequent tasks. A second differentiating characteristic is the extent to which trait inferences will work as expectations in future processing of information about the actor. While weaker inferences cannot be deliberately accessed and do not lead to expectations about the actor, stronger trait inferences can be consciously accessed in posterior tasks and guide subsequent information processing. In the previous experiment (Experiment 7) the first of these characteristics was explored. In the present experiment, our goal is to explore the conditions under which trait inferences result in the creation of expectations about the actor. Previous research has left this question largely unexplored, as well as the more general issue about which are the effects of STIs on posterior processing of information. However, these are important questions to inform us about what are the real consequences of STIs within person perception.

As previously stated, intentionality is assumed to govern the trait inference continuum. Thus, as in the previous study, we used the comparison between memory and impression conditions as an initial step to approach this question. Since impression participants are deliberately trying to understand the actor, inferred traits are likely to be mentally represented as dispositional characteristics of actors and act as expectations in future processing of information. Under memory conditions, because traits are inferred spontaneously, it is not clear whether they are represented as dispositional characteristics of actors or not. Crucial debates exist in the field about whether STIs are really characterizations of actors or mere characterizations of behaviors (Bassili, 1989a, 1989b; Carlston & Skowronki, 1994; Clayes, 1990; Todorov & Uleman, 2002), and about whether trait-actor links have an inferential or associative nature (Brown & Bassili, 2002; Carlston et al., 2005; Carlston et al., 2007; Carlston et al., 2008; Todorov & Uleman, 2004). Both of these debates speak directly to the question about the creation of expectations. If STIs are only characterizations of behaviors, and are only associated with actors, expectations are not likely to exist. However, if STIs are inferences about the personality of the actor, then they should result in the creation of expectations.

In this experiment, our goal was to explore whether STIs lead to expectancies about the actor. We applied an adaptation of the false recognition paradigm (Todorov & Uleman, 2002) and included a reading time measure. Initially, participants were presented with trait-implying sentences under memory or impression formation instructions. Each sentence was paired with the name of a different actor (e.g., “John Moore”). After a distracting task, half of the participants completed the typical false recognition measure (Todorov & Uleman, 2002). That is, participants were presented with trait-actor pairs and asked to indicate whether the trait was included in the sentence previously paired with that name. In some of the pairs, names were paired with the trait implied by the behavior previously presented with that name (*implied condition*), while in other pairs names were paired with traits implied by behaviors performed by a different actor (*mismatch condition*). If traits are spontaneously inferred during encoding and are specifically associated with the actor, more false recognitions are expected in the implied-condition than in the mismatch condition. This pattern of results has been found in previous studies (Todorov & Uleman, 2002, 2003, 2004).

The other half of the participants, instead of the false recognition task, completed a reading time measure. In this case, a list of new behaviors about the previous actors was presented. The list included three new behaviors about each one of the actors: one consistent, one inconsistent, and one neutral in relation to the trait implied by the behavior presented in the initial phase of the experiment. Participants were simply asked to read the sentences at a normal reading pace, and to press the spacebar to go to the next sentence. The time participants take to read consistent and inconsistent sentences was taken as an index of the existence of expectations. Specifically, if trait inferences made during the initial phase of the study were encoded as a property of the actor, they should act as expectations and guide subsequent processing. As a consequence, participants should take more time to read behaviors that are inconsistent with previous trait inferences, than behaviors that are consistent with previous trait inferences. Because inconsistent behaviors are in contradiction with existing expectations, participants should experience greater comprehension difficulty during their encoding.

Reading times have been generally used as a measure of comprehension difficulty (e.g., Corbett, 1984; McKoon & Ratcliff, 1980; McKoon & Ratcliff, 1992). In

addition, several authors have shown that reading times are a sensitive measure of the existence of expectations (Baker & Anderson, 1982; de Vega, 1995; Duffy, 1986; Locke & Walker, 1999; Irmen & Roßberg, 2004; O'Brien & Albrecht, 1992; Vonk, 1985). Baker and Anderson (1982), for example, showed that people take more time to read inconsistent sections of text passages. O'Brien and Albrecht (1992) used reading times to measure participant's sensitivity to violations of information about the localization of a protagonist. They found that participants take more time to read a sentence (e.g., She decided to go outside the health club) when it was inconsistent with previous information from the text (e.g., Kim stood outside the health club) than when it was consistent with previous information (e.g., Kim stood inside the health club) (see also de Vega, 1995).

Importantly, several studies within the person memory field have reported longer reading times for behaviors that are inconsistent with existing expectancies Stern and collaborators (Stern et al., 1984, Experiment 2), for example, presented participants with a list of behaviors under impression formation instructions. The first half of the list was composed of behaviors that implied the same trait in order to create an expectation. In the second half of the list some behaviors were inconsistent, others consistent, and others neutral with the expectancy created. Participants controlled the time behaviors appeared on the screen. Results showed that when all behaviors describe the same individual, participants spend more time reading inconsistent than consistent behaviors. Other studies consistently reported longer reading times for expectancy-inconsistent behaviors under impression formation conditions (see Bargh and Thein, 1985; Crocker & Vitus, 1983; Hemsley & Marmurek, 1982). Thus, it is well established in the literature that behaviors that are inconsistent with existing expectations take longer to process and capture people's attention more.

Some particularities of this study deserve some clarification. First, one difference between our paradigm and the typical false recognition paradigm was that we used names, instead of pictures. This change was implemented because presentation of photos may trigger impression formation processes, even for memory participants. This is a major problem of STI paradigms that apply faces as stimuli. People are very sensitive to stimulus faces (Johnson, Dziurawiec, Ellis, Morton, 1991; Farah, Wilson, Drain, & Tanaka, 1998; Valenza, Simion, Cassia, Umiltà, 1996). Thus, it is likely that

the simple presence of a face triggers impression formation processes independently of the instructions provided. This may be especially likely when the face is paired with a behavior, as is the case in STI paradigms. Because of that, in the present experiment we opted for the presentation of proper names to describe actors, instead of photos.

Another important feature of this study was that, in addition to “self descriptive trials” in which the behavior described was performed by the named person (e.g., “Christopher Allen visited a sick friend in the hospital”), we also included “other descriptive trials” in which the named person describes a behavior from an acquaintance (e.g., “Paul Campbell said that his friend Isaac can speak three different languages”). Previous research has shown that under these conditions the implied trait became associated with the communicator of the behavior (i.e., trait transference effect, see Carlston & Skowronski, 2005; Crawford, et al., 2007; Crawford et al., 2008; Mae et al., 1999; Skowronski et al., 1998). However, since the link between the trait and the communicator is assumed to be merely associative (Mae et al., 1999; Skowronski et al., 1998), it should not be mentally represented as a characteristic of the communicator and will not have the power of acting as an expectation. Thus, in “other-descriptive trials” we expected that the trait transference effect would be observed in the false recognition measure, similar to previous studies (Bassili & Smith, 2002; Todorov & Uleman, 2004). However, we expected no differences in reading times for consistent and inconsistent in this case. Notice that since no differences are expected to be observed in the “other-descriptive” condition, this represents an important control condition.

Another relevant aspect of this paradigm was that sentences that included the implied trait were also included in the initial list of behaviors (e.g., “Edward Williams was so responsible that he paid his taxes early”). In this case, by mentioning the implied trait an expectation about the actor is being explicitly provided. As a consequence, in the second phase of the experiment participants should take more time to read inconsistent than consistent behaviors about the actor. This is also a relevant control condition. If effects of expectations on reading times are not observed in this case, then it might be the case that the measure is not being sufficiently sensitive to detect differences in terms of expectancies.

Finally, while some of the sentences that included the trait were “self descriptive”, other sentences were “other descriptive” (e.g., “Daniel Parker said that his

friend Bob was so stubborn that he only rarely changed his mind”). In this case, we predicted that transference effects will be less likely to occur both in the reading time measure, as well as in the false recognition measure. Because the trait is explicitly mentioned, it will be easier for participants to associate the trait with the appropriate actor. Previous studies support this prediction (Crawford et al., 2007; Crawford et al., 2008; Goren & Todorov, 2009; Todorov & Uleman, 2004), by showing that making more salient the link between the trait and the appropriate actor weakens or eliminates spontaneous trait transference effects.

The main hypotheses for this experiment can be summarized in the following way. In terms of the false recognition measure, we predict the typical false recognition effect, under both impression formation and memory instructions, in agreement with previous results (Todorov & Uleman, 2002, Experiment 4). That is, participants should make a greater number of false recognitions for trials in which the implied trait is matched with the corresponding actor (i.e., implied condition) than for trials in which traits are matched with different actors (i.e., mismatch condition). Moreover, the false recognition effect should be weaker in the other-describing condition than in the self-describing condition, reflecting weaker spontaneous trait transference effects than spontaneous trait inference effects.

In terms of reading times we make different predictions depending on the specific condition. For self descriptive trials, we predicted that participants would take more time to read inconsistent than consistent behaviors under impression formation instructions. The same pattern is expected under memory conditions when traits are included in the sentences. However, under memory conditions, when traits are not included in the sentences but are only implied, there are different possible predictions. If STIs lead to the creation of expectations about the actor, then memory participants should take more time to read inconsistent than consistent behaviors. However, if STIs do not have consequences in terms of expectations about actor’s future behavior, then no differences should be observed.

In other describing trials, predictions can be made about the occurrence of trait transference effects. When impression formation instructions about the communicator are provided, no differences in terms of reading times are expected because in this case it is easily apparent that the behavior was performed by a friend, and not by the

communicator himself. Thus, a specific link should be established between the trait and the appropriate communicator. The same prediction is made for memory participants when the traits are included in the sentences. For memory participants, when traits are not included, there are two different possibilities. According to our predictions, because trait transference effects are caused by mere associative mechanism, no expectations should be formed about the actor, and no differences in reading times should be observed. However, according with the model developed by Skowronski and colleagues (Mae et al., 1999; Skowronski et al., 1998), despite their associative nature, transferred traits may have an implicit influence on impressions about the actors to whom they became associated. Thus, according to the framework of Skowronski and colleagues (Mae et al., 1999; Skowronski et al., 1998) the possibility exists that participants will take more time to read inconsistent than consistent behaviors even under memory instructions.

Method

Participants and Design

Participants were 98 undergraduate students at the University of California, Santa Barbara (54 women and 21 men) whose participation earned them partial credit for a psychology course. Participants were randomly assigned to the conditions of a 2 (processing goals: impression formation vs. memory) X 2 (measure: false recognition, reading time) X 2 (trial condition: self descriptive vs. other descriptive) X 2 (trait inclusion: trait not included vs. trait included) X 2 (type of pair: implied condition vs. mismatch condition) X 2 (material replication: replication 1 vs. replication 2) design. The first two variables were manipulated between-subjects and the last three variables were manipulated within participants.

Stimulus Material

Trait-implying behaviors were selected from existing norms compiled by Stroessner (1989). Initially, 12 behaviors were selected that implied different traits. These behaviors were presented in the initial phase of the experiment. For each of the behaviors one consistent (i.e., implied the same personality trait), one inconsistent (i.e.,

implied the opposite personality trait), and one neutral behavior were selected, to be presented in the posterior reading time task (for a complete list of the sentences, see Appendix H). All sentences were paired with male names. The names of the actors and of the acquaintances were taken from a website in the internet with information about the most common names for males in United States (for the names selected see Appendix H). Each actor was described by his first and last name (e.g., John Moore). During the initial encoding phase, 12 sentences were presented that described different actors. Six of these sentences described behaviors performed by the named person (e.g., “Christopher Allen visited a sick friend in the hospital”), while the other 6 sentences described behaviors that were being only communicated by actors about a friend’s behaviors (e.g., “Paul Campbell said that his friend Isaac can speak three different languages fluently”). For both self-describing and other-describing conditions, the trait was not included in 4 sentences while the remaining two sentences included the trait (e.g., ““Daniel Parker said that his friend Bob was so stubborn that he only rarely changed his mind”). Overall, half of the behaviors had a negative valence and half had a positive valence, equally distributed across conditions (see Appendix VIII).

In the recognition task, from the 4 trait implying sentences, half of them were represented the implied condition (i.e., the implied traits were tested with the corresponding names, in the false recognition measure), and the other half represented the mismatch condition (i.e., the implied traits were tested with the name of a different actor).

Different replications of the material were created. One corresponds to the explained previously. In other replication, 12 behaviors were selected that imply the opposite traits from the behaviors initially presented in replication 1. In this way, the congruency of the behaviors presented in the second phase was altered without changing the behaviors. In the other replications, actors’ names, self-other descriptive nature, trait inclusion, and implied-mismatch factors were varied across behaviors, in order to control material effects. In total, 8 different material replications were used.

Procedure

The experiment was run on computers using the software E-prime (Schneider et al., 2002). Participants worked individually and all instructions and responses were given through the computer. The experiment was divided in two phases.

In the first phase, participants were initially informed that they would see names of different people and that each name would be paired with a sentence. It was explained that some of the times the sentence would make reference to the person named in the sentence, while other times the person named in the beginning of the sentence would be describing a behavior of a friend. Sentences were said to have been selected from a longer list of statements that the named persons had provided in response to an interview. Depending on the condition participants were given either impression formation or memory instructions. In the impression formation condition, participants were instructed to form an impression of the personality of the actor named at the beginning of each sentence. In the memory conditions, participants were instructed to try to remember the material as well as they could. The experiment started by the presentation of three practice trials in order for participants get familiar with the task. The list of 12 experimental behaviors was then presented with each behavior being displayed on the screen for 5 seconds. A white screen was presented between sentences for 500 ms. Of the 12 sentences, 6 were “self-descriptive” and 6 were “other descriptive”. From each of these conditions, 4 sentences did not include the implied trait and 2 sentences included the implied trait. In addition, half of the behaviors were negative and half were positive in terms of valence. The list of 12 behaviors was randomly presented for each participant. After the presentation of all behaviors, participants completed an anagram filler task for 4 minutes.

In the second phase of the experiment, half of the participants completed a false recognition task and the other half completed a reading time task. In the false recognition task, participants were instructed that they would be presented with the names of the 12 actors they had seen in the first part of the study, and that each name would be accompanied by a single word. The words were the traits implied by previous sentences. Participants’ task was to decide whether the presented word was part of the sentence previously paired with the name presented. In order to do that, participants should press the button “I” on the keyboard if the exact word was in the sentence

previously paired with the name shown, and press the button “E” on the keyboard if the word was not in the sentence previously paired with the name. Participants were instructed to respond as fast and accurately as possible. Before the experimental “name-trait” pairs were presented participants were given some practice trials in order to assure that they understood the task. After that, the 12 experimental pairs were presented. Half of the names previously paired with trait-implying behaviors were paired with the trait implied by the correspondent behavior (implied condition) and the other half were paired with the trait implied by the behavior performed by another actor (mismatch condition). Actors previously paired with behaviors in which the trait was included were paired with the mentioned trait. The order of presentation of the name-trait pairs was randomized for each participant.

In the reading time task, participants were instructed that they would be presented with additional behaviors performed by the same actors named in the beginning of the experiment. Participants were asked to read the behaviors at a comfortable, normal reading pace. They were instructed to press the “spacebar” to move to the next sentence. Every time the spacebar was pressed the screen was erased and a new sentence appeared on the screen. The time between “spacebar” presses was recorded. The list of behaviors was composed of three behaviors for each one of the actors: one consistent, one inconsistent, and one neutral in relation to the trait implied by the behavior previously paired with the same actor. Sentence presentation was randomized for each participant.

Dependent Measures

In the false trait recognition task, recognition accuracy and response times for correct recognition responses were the dependent measures. In the reading time task, time between spacebar presses was the dependent measure.

Results

No significant effects were found for material replications in any of the two dependent measures ($F < 1$). This variable will not be further considered.

False Recognition Task

Accuracy. The proportion of correct recognition of presented traits was .71 ($SD = .30$) in the self-describing condition. In order to analyze spontaneous trait inferences and trait transference occurrence, trait false recognitions were entered in a 2 (processing goal: impression formation vs. memory) X 2 (trial condition: self descriptive vs. other descriptive) X 2 (type of pair: implied condition vs. mismatch condition) repeated measures ANOVA. The only emerging effect was a significant interaction between trial condition and type of pair, $F(1, 46) = 7.80, p = .007$. Planned comparisons revealed that in the self-describing condition, participants were more likely to falsely recognize the trait when it was paired with the correspondent name ($M = .53$) than when it was paired with the name of a different actor ($M = .35$), $t(46) = 2.04, p = .02$ (one-tailed planned comparisons). Thus, the typical false recognition effect was replicated in the present experiment (e.g., Todorov & Uleman, 2002). This supports the idea that personality traits were spontaneously inferred and associated with the corresponding actor. Moreover, the false recognition effect was observed both for memory and impression formation participants, given that no interactions were observed with processing goal ($F < 1$).

By contrast, planned comparisons revealed that in the other-describing condition, trait false recognitions were not significantly different in the implied and in the mismatch condition, $t(46) = 1.45, ns$ (one-tailed). Thus, no evidence for traits transference effects was obtained in this study, either in impression formation conditions or in memory conditions.

Response Times. Four participants that exhibited response times faster than 100 ms. were removed from the analysis. Remaining response times for correct responses were submitted to a 2 (processing goal: impression formation vs. memory) X 2 (trial condition: self descriptive vs. other descriptive) X 2 (type of pair: implied condition vs. mismatch condition) repeated measures ANOVA. The ANOVA revealed a significant interaction between type of processing goal and trial condition, $F(1, 18) = 4.52, p = .047$, revealing that while memory participants took more time to correctly reject traits in the self-describing condition ($M = 2402$ ms) than in the other-describing condition ($M = 1826$ ms), impression formation participants took more time to correctly reject

traits in the other-describing condition ($M = 2037$ ms) than in the self-describing condition ($M = 1696$ ms). We have no specific interpretation for this interaction. No other significant effects emerged from response time's analysis.

Reading Time Task

The time between spacebar presses was taken as a measure of comprehension difficulty. Participants that exhibited consistently very fast reading times (i.e., faster than 1000 ms), were removed from the analysis because this time was not enough to read the sentence (according with pilot reading times tests). These participants were probably not paying attention to the task. Participants who showed exceptionally long reading times were also removed from the analysis (i.e. longer than 6000). These exceptionally long reading times were probably due to momentary distractions that took participants' attention away from the screen. These criteria led to the elimination of 12 participants from the analysis.

In order to analyze whether previous trait inferences functioned as expectations, reading times on self-describing trials were submitted to a 2 (processing goal: impression formation vs. memory) X 2 (type of pair: implied vs. mismatch) X (behavior consistency: consistent, inconsistent, neutral) mixed ANOVA. No significant effects emerged from the ANOVA analysis. Thus, no differences in reading times were found as a function of behavior consistency with previous trait inferences, for any of the conditions. While no differences in reading times were expected on mismatch trials because behavior congruency is manipulated in relation to a different actor, the absence of differences on implied trials is less expected, especially under impression formation conditions. The fact that no differences were found in the time participants take to read consistent, inconsistent and neutral behaviors indicates that trait inferences made in the initial phase of the experiment were not sufficiently strong to create expectations about actors.

In addition, to explore transference effects on reading times, reading times on other-describing trials were submitted to a 2 (processing goal: impression formation vs. memory) X 2 (type of pair: implied vs. mismatch) X (behavior consistency: consistent, inconsistent, neutral) mixed ANOVA. The only effect revealed by the ANOVA analysis was a marginal main effect for type of pair, $F(1, 32) = 3.90$, $p = .056$, revealing that

participants took more time to read behaviors in the implied condition ($M = 3564$ ms) than in the mismatch condition ($M = 3286$ ms). More relevant to our hypothesis, and confirming our predictions, no differences in the time participants take to read consistent, inconsistent and neutral behaviors are observed on other-describing conditions. Thus, underlying traits are not transferred to the communicator of the behavior in the form of expectations about his own behavior.

Finally, we analyzed whether including the trait in the initial behavior led to the creation of expectations. In order to examine this issue, reading times on trials in which the personality trait was included were analyzed in a 2 (processing goal: impression formation vs. memory) X 2 (trial condition: self descriptive vs. other descriptive) X 3 (behavior consistency: consistent, inconsistent, neutral) mixed ANOVA. The ANOVA revealed a main effect for behavior consistency, $F(2, 64) = 13.64, p < .001$, revealing that participants took more time to read both consistent ($M = 3660$ ms) and inconsistent ($M = 3671$) than neutral behaviors (3028 ms).

However, this main effect interacted with trial condition, $F(2, 64) = 3.12, p = .05$. Planned comparisons revealed that in the self-describing condition participants took significantly more time to read inconsistent ($M^{\text{Inc}} = 3595$ ms) than both neutral ($M^{\text{Neu}} = 3140$ ms) and consistent behaviors ($M^{\text{Cons}} = 3394$ ms), $t(32) = 1.66, p = .053$ (one-tailed). No significant differences were observed between reading times for neutral and consistent behaviors, $t(32) = 1.77, ns$. On other describing trials, planned comparisons only revealed a significant effect showing that participants took more time to read both consistent ($M = 3947$ ms) and inconsistent behaviors ($M = 3748$ ms) than neutral behaviors ($M = 2916$ ms), $t(32) = 39.65, p < .001$ (one-tailed). No other significant effects emerged from the ANOVA analysis.

Thus, as predicted, effects of expectations emerged when traits were included in the sentences on self-describing trials. Importantly, this effect didn't interact with processing goals, $F < 1$. Thus, participants took significantly more time to process inconsistent behaviors than consistent and neutral behaviors, both under impression formation and memory conditions.

Discussion

The goal of the present experiment was to examine whether STIs work as expectations about the future behavior of actors. Using an adaptation of the false recognition paradigm (Todorov & Uleman, 2002), we found evidence for the occurrence of STIs. Specifically, participants under both impression and memory conditions made more trait false recognitions when an actor was paired with the trait implied by his previous behavior (i.e., implied condition) than when the actor was paired with a trait implied by a behavior previously performed by another actor (i.e. mismatch condition). These results are in line with previous findings (Todorov & Uleman, 2002; 2003; 2004), and suggest that participants spontaneously inferred the trait and associated it with the actor of the behavior. If the trait has been activated, but had not specifically been associated with the actor, no differences should be observed between match and mismatch conditions.

In the false recognition paradigm, no differences were obtained in response times. It should be noted, however, that the absence of differences in response times within the false recognition measure can be explained by the relatively difficulty of the task. Participants are presented with a list of 12 different actors, and the recognition task is only administered after a filler task. Given the difficulty of the task, it is not surprisingly that differences are mainly reflected in terms of the likelihood of making trait false recognitions. In line with this view, authors that use the false recognition paradigm tend to make predictions mainly in terms of false recognitions (Goren & Todorov, 2009; Todorov & Uleman, 2003, 2004) and in cases in which differences in reaction times are reported (see Todorov & Uleman, 2002), they are not consistent across experiments.

However, despite evidence that STIs were made, no subsequent differences were observed in the time participants took to read behaviors that were consistent or inconsistent in relation with the previously inferred trait. Thus, despite the evidence for STIs in the false recognition measure, no evidence was obtained in the reading time measure for subsequent processing effects of the inferred traits. Apparently, the actor-trait association created was not sufficiently strong to take the form of expectancies about the future behavior of the actor. The fact that traits spontaneously inferred do not

work as expectancies that guide future processing suggests that STIs do not become strongly mentally represented as a dispositional characteristic of the actor.

One more unexpected result in this study, however, was the absence of reading times differences between consistent and inconsistent behaviors under impression formation conditions. In this case, participants have an explicit intention to form an impression about the personality of the actor. Because of that, the inferred trait should be encoded as a characteristic of the actor and function as an expectancy about his future behavior. However, in the present paradigm, even intentional trait inferences did not have consequences in terms of future processing of consistent or inconsistent information about the actor. How can the absence of differences in reading times be explained, even when participants are intentionally inferring personality traits about the actors?

One possible interpretation is that making a trait inference from a single behavior is not sufficient to guide future processing of information about the actor, even when participants have explicit impression formation intentions. This possibility is intriguing because it suggests that even when participants are trying to extract information about the personality of the actors, a single behavioral observation will not be sufficient for expecting behavioral consistency. This might be especially true given the present study conditions (that are also the study conditions in which STIs are typically studied within STI research). In the present study conditions, behavioral information is presented about a variety of different actors, only one behavior is presented about each one of the actors, and behaviors are detached from any context or individual story. All these aspects may contribute to the absence of strong expectations about the future behavior of actors. In agreement with this explanation, Whitney, Davis, and Waring (1994) have found differences in reading times under impression formation conditions when behaviors are incorporated in more meaningful narratives about the actors. In this study, Whitney et al. presented participants with short narratives composed of 4 sentences. The first sentence from the narrative could be consistent, inconsistent, or neutral with the trait implied by the last sentence. Results showed that participants took longer times to read the last sentence when it was inconsistent with the first sentence, but only under impression formation instructions. When participants were

given factual instructions (i.e., informed that they would be given factual questions after reading the narratives) no differences in reading times were observed.

However, a crucial aspect of the present data is the fact that when the trait is included in the behavior, evidence for the existence of expectancies about the actor was obtained. This was the only condition in the study in which differences in reading times were observed. As we expected, participants took more time to read inconsistent behaviors than neutral and consistent behaviors. No differences were observed between the time they took to read consistent and neutral behaviors, however. These results indicate that inconsistent behaviors were more difficult to process than neutral behaviors and consistent behaviors. The fact that differences in reading times were observed in this case suggests that the reading times measure was sensitive to the existence of expectancies.

Finally, no differences were obtained in terms of reading times when trait consistency was manipulated in relation to the behavior of an acquaintance. Thus, as expected, when a communicator described a behavior performed by a friend the underlying trait was not used to create expectations about the future behavior of the communicator himself. However, contrary to our expectations, we also didn't find trait transference effects in the trait false recognition measure. The number of trait false recognitions in other-describing trials did not differ between implied and mismatch conditions. Thus, in this study the implied trait was not transferred to the communicator of the behavior. This result is in contradiction with previous results (e.g., Brown & Bassili, 2002; Carlston & Skowronski, 2005; Crawford et al., 2007; Crawford et al., 2008; Mae et al., 1999; Skowronski et al., 1998). One possible interpretation of the results is in terms of the type of material used in this experiment. The main difference between the present study and previous studies that reported trait transference effects was the fact that in our study we did not present photos. While in previous studies a photo of the communicator was usually shown, in our experiment the communicator was identified only by its name. Because under these conditions the perceptive salience of the communicator is lower, participants were probably less likely to establish an association between the implied trait and the communicator.

In this sense, results from this experiment might be important to help clarify the limiting conditions of the spontaneous trait transference effects. Previous research has

shown that when the salience of the actor of the behavior is increased, for example by presenting not only the photo of the communicator but also the photo of the actor (see Crawford et al., 2007; Crawford et al., 2008; Goren & Todorov, 2009; Todorov & Uleman, 2004), spontaneous trait transference effects are reduced or eliminated. Our results suggest that spontaneous trait transference reduction may be due not only to the increase of the perceptive salience of the actor of the behavior, but also to the decrease in the perceptive salience of the communicator, to whom the trait would be otherwise transferred. Consistent with this interpretation, Goren and Todorov (2009, Experiment 3) showed that spontaneous trait transference effects are eliminated when faces and behaviors are separated perceptively. In this study, behaviors were presented only 3 seconds after the irrelevant faces were shown. No transferences of the implied traits to the irrelevant faces were observed under these study conditions, despite STI effects being observed even when the actor was shown only 3 seconds after the behavior. Both Goren and Todorov's (2009) results and the findings from our present experiment suggest that transference effects are eliminated when the irrelevant actor is not present (Goren & Todorov, 2009) or sufficiently salient (Experiment 8) during the encoding of the behavior.

One possible limitation of the present experiment is that we did not include a measure to check of sentence's comprehension in the reading time task. Thus, we cannot guarantee that participants were in fact reading and processing the presented sentences. There are, however, some indications that participants were in fact processing the sentences in a meaningful way. First, before reading the behaviors instructions mentioned that participants should pay attention to the sentences because after sentence presentation some questions would be asked regarding the content of the sentences. Second, evidence that participants were paying attention to the content of the sentences comes also from the reading times. Very few participants exhibited reading times shorter than the time that was necessary to understand the sentences, and those who exhibited very short reading times were excluded from the analysis. If the goal of the participants was only get rid of the task, a higher number of very short reading times should have been observed. Finally, when the trait was included in the sentences, differences in reading times between the different types of behaviors were observed.

This clearly indicates that participants were processing behaviors in a meaningful way and were sensitive to behavior congruency.

In sum, the present experiment helps to clarify the consequences of STIs in terms of future processing of information about the actor. While evidence from the false recognition measure showed that STIs were associated with actors of the behavior, the association established was not sufficiently strong to create expectations about actor's future behavior. This experiment also revealed the surprising fact that expectations were not developed even when impression formation instructions were provided. Apparently, only when impression formations instructions were furnished within more meaningful encoding contexts were expectations developed that guide subsequent information processing (Whitney et al., 1994). Probably, the narrative structure used by Whitney et al. (1994) favors greater individuating and elaborative processing about the actor. The greater elaborative processing leads to more extreme trait inferences that became represented as a dispositional characteristic of the actor.

Discussion of Experiments 7 and 8

In the two final experiments, we explored the principle according to which trait inferences vary in a continuum of strength. Two basic characteristics of the higher level of a trait inference are the ability to make deliberative uses of the inferred trait in subsequent tasks, and the capacity of trait inferences to act as expectations about the actor. Experiment 7 explored whether trait inferences made under spontaneous conditions can be deliberately accessed and Experiment 8 explored whether STIs act as expectations.

Results from Experiment 7 confirmed that while intentional trait inferences can be deliberately accessed to guide subsequent recall of behaviors, spontaneous trait inferences remain inaccessible to subsequent conscious uses. In this study, differences in recall and in level of clustering between trait-cue and no-cue conditions were used as a non obtrusive measure of the degree of conscious trait access. Under impression formation instructions, providing the cues made no difference in terms of level of recall and trait clustering exhibited during recall. This indicates that participants spontaneously accessed the inferred traits to guide recall, independently of the traits being provided to them. By contrast, under memory conditions, participants recalled a

higher number of behaviors when traits were provided during retrieval than when no traits were provided. This suggests that participants did not spontaneously use the traits to guide recall, unless traits were explicitly provided to them. Thus, effects of previous STIs became visible only when the encoding context is reinstated.

In Experiment 8, similar to previous studies (e.g., Todorov & Uleman, 2002) evidence was found that STIs are associated with the actor of the behavior when a false trait recognition measure is used. However, the actor-trait association was not sufficiently strong for participants to expect behavior consistency from actors. No differences were observed in the time participants took to read consistent, neutral, and inconsistent behaviors in relation with the trait implied by the previous behavior. Surprisingly, no differences in reading times were also observed in our experiment under impression formation conditions. By contrast, consistent with existing evidence, when information was presented in the form of short narratives, participants took more time to read the last sentence when it was inconsistent with the trait implied by the initial sentence (Whitney et al., 1994). Taken together, the pattern of results from Experiment 8 and results from Whitney et al. (1994) suggest that STIs do not become represented as dispositional characteristics of the actor in the present study conditions. In addition, intentional trait inferences made under impression formation processing goals may or may not lead to expectations about the actor. When study conditions do not favor an individuating processing (i.e., information is presented about several different actors, only one behavior is presented about each actor, and behaviors are detached from any narrative structure or context), intentional trait inferences apparently also do not lead to expectations of actor's behavior consistency. Only when more meaningful information is provided about each actor, for example in the form of narratives, are expectations created that guide subsequent information processing.

This reasoning is in agreement with the idea that trait inferences made under memory and under impression formation instructions are not necessarily qualitatively different, but probably vary in terms of degree. That is, differences between trait inferences made under memory and impression formation conditions might be differences in terms of degree of the same underlying process. In this sense, differences between memory and impression formation instructions might be found under some study conditions. However, differences between two different impression formation

conditions might also be found. Apparently, the more the study conditions favor an individuating processing of information, the more the inferred trait is specifically linked to the behavior and represented as a dispositional characteristic of the actor.

An important point about results from Experiment 8 was that no differences were found in the reading times measure. The reading times measure may be considered an implicit memory measure (see Roediger, 1990; Schacter, 1987). Participants are asked to perform a task without any reference to the previous study episode. Thus, the influence of the previous study episode on the task remains probably unnoticed. However, despite no differences being found in terms of reading times, other studies have reported effects of STIs using different types of implicit measures. For example, Carlston and Skowronski (1994) showed effects of STIs using the savings in relearning paradigm. In this case, it is shown that participants exhibit learning facilitation for old trait-face pairs than for new trait-face pairs. Why are differences observed in one implicit measure (savings in relearning) but not with the other (i.e., reading times)?

We may hypothesize that only when the encoding context is in same way reinstated will influences of STIs become visible. In the savings in relearning task the presentation of the photo and of the implied trait may serve to reactivate the actor-trait association. By contrast, when the association between the actor and the inferred trait is not reinstated, consequences of STIs are not manifest. This is probably the case in the reading times measure. In this case, participants are only asked to read the behaviors. No reason exists for the previous actor-trait association being reactivated. We may theorize that the more the trait-actor association is reinstated during retrieval, the more effects of STIs will become evident. The more the association is obscured by the task, the less STI effects will be apparent.

In sum, results from Experiments 7 and 8 showed that when STIs are made no conscious access to the output of the inference exists in subsequent tasks and the inferred trait does not act as an expectation about the actor.

PART III

GENERAL DISCUSSION AND CONCLUSIONS

1. Main Findings

The goal of the present project was to explore the flexibility of the STI process. The question about the inevitable nature of the process has been a central one in the field. Anchored on the text comprehension literature, three principles were formulated that together sustain a flexible view of STIs occurrence. These principles state that (a) STIs are more likely when they are easily available; (b) STIs work in line with local coherence requirements; and (c) trait inferences vary along a continuum. Experiments were conducted in order to test each of these principles.

According to the first principle, STIs occur when easily inferable from the information conveyed in the text. This principle was tested in Experiment 1 with the probe recognition paradigm. Participants were presented with ambiguous behaviors that could be interpretable by two different personality traits. Evidence was obtained showing that STIs are more likely to occur when a category label is provided that disambiguates the trait interpretation of the behavior than when the behavior remains ambiguous by being presented with a neutral label, or when the category label presented favors an alternative trait interpretation. These results are consistent with our predictions and support the notion that STIs are dependent on the easiness with which the trait is inferred from the information presented.

The second principle was examined in two sets of experiments. This principle asserts that the perceiver aims to maintain text local coherence, and STIs are modeled by this basic need. In the first set of experiments that tested this principle (Experiments 2-4), information about the actor and about the situation was provided in order to create different patterns of text coherence. Evidence from the probe recognition paradigm showed that not only STIs but also SSIs occur in line with the postulated coherence principle. Specifically, in general results showed that STIs are more likely when the category label is consistent with the trait implied by the behavior. SSIs were shown to be facilitated by a category label inconsistent with the behavior and to be inhibited by a category label consistent with the trait implicated by the trait. However, the pattern of results obtained was different depending on the material provided. While in Experiments 2 and 3, stronger evidence was found for STIs effects, in Experiment 4, with different behaviors, stronger evidence was obtained for SSIs effects.

In a second set of experiments (Experiments 5 and 6), the local coherence principle was tested by providing two behaviors about the same actor that could be either trait-consistent or trait-inconsistent. Experiment 5 applied a cued recall paradigm in which participants were asked to recall the second behavior of each pair. For consistent pairs, traits implied by second behaviors were better retrieval cues than first behaviors, while for inconsistent pairs first behaviors were more effective cues than traits. Thus, behaviors that are inconsistent with previous behavioral information are less likely to be encoded in trait terms and are more likely to be associated with previous behaviors. The online status of the observed effects was examined in Experiment 6, applying the probe recognition paradigm. As expected, it was easier for participants to correctly reject the trait probe on inconsistent than on consistent trials. Both experiments reinforce the view that STIs are linked to a need for comprehension coherence. STIs are inhibited when they are incoherent with a previous behavior of the same actor, and are facilitated when they are in agreement with previously received information.

The principle according to which trait inferences vary along a continuum was explored in two final experiments (Experiments 7 and 8). Each experiment focused on one defining characteristic of the highest level of trait inferences. Experiment 7 examined whether STIs can be consciously accessed in posterior tasks and Experiment 8 explored whether STIs result in the creation of expectations about the actor. Results from Experiment 7 showed that while intentional trait inferences tend to be used to guide and organize subsequent behavior recall, spontaneous trait inferences are not used to facilitate behavior recall unless traits are explicitly provided as cues. Experiment 8 furnished evidence that STIs are associated with the actor, using the false recognition paradigm. However, it was shown that the trait-actor association established is not sufficiently strong to guide subsequent processing of information. Participants exhibited no differences in reading times for behaviors consistent and inconsistent with a previous STI.

The present experiments depart in several ways from previous existing studies. In most of them, the material applied is more complex and heterogeneous than is usually the case in STI research. The complexity of the material was reached by different means. Ambiguous behaviors were applied (Experiments 1 and Experiment 4),

information about the social category of the actor and about the situational circumstances of the behavior were provided (Experiments 2, 3 and 4), previous behavioral information about the actor was manipulated (Experiments 5 and 6), and several behaviors were used to describe the same actor, with some of them implying the same personality trait (Experiment 7). In addition, the consequences of STI were also examined in the present experiments (Experiments 7 and 8), an issue that has been left largely unexplored. We explored whether a trait that was spontaneously inferred can be consciously accessed (Experiment 7) and whether it serves as an expectation about what might be the future behavior of the actor (Experiment 8). All of these manipulations represent new ways of studying STI occurrence under more realistic social settings. As a whole, results support a flexible view of the STI process and contribute to a deeper understanding of the STI process.

Theoretically, our research represents an initial attempt to examine the general principles that guide occurrence of STI occurrence. This is especially relevant in a field that has been marked by the absence of theoretical frameworks. These principles may have an important role in orienting and stimulating future research.

2. Potential Limitations

There are potential limitations in the present experiments that should be discussed. These aspects are important in order to contrast alternative explanations for our findings, as well as to be considered in future studies. Some of these aspects have to do with limitations of the paradigms used. As outlined in the introduction, none of the paradigms developed in the literature on STIs is free from criticisms, and all of them seem to allow alternative explanations (see Uleman et al., 1996).

2.1. Paradigms Weaknesses

The probe recognition paradigm was used in five of our experiments (Experiments 1, 2, 3, 4 and 6). The benefits of this method are linked to the fact that the STI process interferes with performance in this task. In addition, the fast stimulus presentation and the fast response requirement are likely to prevent strategic and reconstructive processes. For these reasons, the probe recognition paradigm has been considered particularly adequate to prove the unintentional nature of the trait inference

process (see Overwalle et al., 1999; Uleman et al., 1996; Wigboldus et al., 2003). However, results from the probe recognition paradigm are subject to an alternative explanation. It is possible that results are not due to inferences made during encoding, but to processes that occur only after the probe word is presented (Keenan et al., 1990). According to this hypothesis, when the probe word appears, participants compare it with the meaning of the previous sentence which is still in short term memory. This *compatibility* or *context-checking* procedure (Forster, 1981) would result in a greater difficulty of rejecting the probe word because the meaning of the probe is highly compatible with the meaning of the sentence. As McKoon and Ratcliff (1986) point out, results from the probe recognition paradigm can also be explained by a combination of inferential processes and context-checking mechanisms. In this case, difficulty in rejecting the probe would be due both to spontaneous inferences made during encoding and to compatible backward processes that make the response even harder.

There are reasons to think that backwards context-checking mechanisms were not so important in the present experiments. In those studies in which the crucial comparison is between neutral sentences and trait-implying sentences, as for example in McKoon and Ratcliff (1986), the context-match between the probe and the experimental sentence is clearly higher than the context-match between the probe and the neutral sentence. In most of our experiments (Experiments 1-4) the critical comparison is between the same sentences tested with different category labels. In this case, the match would be quite similar for the different sentences, except for the category label presented. It is still possible that the label itself contributes to a stronger context-match in some conditions. However, the fact that no differences in the difficulty of rejecting the trait were observed when the different categories were paired with neutral behaviors (Experiment 2) indicates that the stereotype label per se did not determine responses, assuming the operation of context-checking mechanisms. The hypothesis that differences in responses are due to the inference activated during sentence comprehension seems to be more adequate to explain the present results.

Despite the arguments presented, results are not totally conclusive about whether a context-checking mechanism can have an influence on the pattern of results obtained. McKoon and Ratcliff (1986) proposed a speeded recognition method in order to guarantee that responses are due to automatic processes. The reasoning is that forcing

rapid responses precludes any elaborative processes that might occur during test, including context-checking processes. Future studies should use this method under the present study conditions in order to provide additional evidence for the encoding nature of the results obtained.

Although our experiments are not conclusive about whether results are due to encoding or retrieval processes, the fact that participants showed greater difficulty in rejecting the probe word in some conditions is interesting in its own right, as McKoon and Ratcliff (1986) called attention to. If the presentation of one sentence results in a greater difficulty to say “No” to a specific probe word, similar influences can be assumed to occur in other processing situations, as for example in normal reading comprehension or conversation. For example, it is expected that if a sentence makes reference to a concept that is inerrable by a previous sentence its processing will be facilitated, no matter whether the mechanisms responsible take place during encoding of the first sentence, encoding of the second sentence, or both.

In addition to the probe recognition paradigm, other paradigms used in the present project have been also subject to criticisms. For example, cue-recall paradigms have been criticized in a similar way to the probe recognition paradigm. The major problem is that it is not clear whether trait inference occurs at the time sentences are presented or only when trait cues are provided. As noted by several authors, traits can be effective cues for behavior retrieval even if no inferences were made during encoding (e.g., Corbett and Doshier, 1978; D’Agostino & Beegle, 1996; McKoon & Ratcliff, 1986; Wyer & Srull, 1989). D’Agostino and Beegle (1996) specifically argued that in cases in which cues are manipulated within subjects, trait cue efficacy can be alternatively explained by output interference effects. According with this explanation, prior trait cued recall may suppress recall of non-cued sentences, by retrieval inhibition or output interference mechanisms (e.g., Nickerson, 1984; Roediger, 1974). We applied cued-recall paradigms in two experiments (Experiments 5 and 7). In Experiment 5, consistent and inconsistent pairs of behaviors were presented, and participants were asked to recall all second behaviors with either implied traits or first behaviors given as cues. In Experiment 7 participants were provided with behaviors representative of four different personality traits about one actor and were later asked to recall the behaviors with trait cues or without cues. However, differences between cue conditions were

observed in these experiments, despite the fact that in both cases cue condition was manipulated *between* subjects. Thus, our results cannot be explainable by output interference effects (D'Agostino & Beegle, 1996). Even so, the hypothesis that retrieval mechanisms can be intervening in the pattern of results cannot be totally dismissed.

Finally, retrieval mechanisms may also intervene in the false trait recognition paradigm (Experiment 8). It is possible that participants retrieve behaviors only when name-trait pairs are presented at test. Previous results have provided evidence against retrieval accounts. It was shown that the trait false recognition effect occurs even when recall is difficult by presenting a long behavior list (Todorov & Uleman, 2002, Experiment 5), and that trait false recognition effects are observed even on trials in which behaviors cannot be recalled (Todorov & Uleman, 2002, Experiment 6). However, these arguments do not completely rule out the possibility of the intervention of retrieval mechanisms.

The difficulty of STI paradigms in avoiding the intervention of contaminating retrieval problems is related with the general problem of developing conceptual implicit tests (Hourihan & MacLeod, 2007). The difficulty relies on the fact that performance on conceptual implicit tests is based on semantic processing of items. However, semantic processing tends also to increase the memorability of information, increasing the probability of explicit retrieval processes. Ideally, test conditions should be developed in which participants perceive no benefit in explicitly retrieving previous information. The reading time measure used in Experiment 8 may fulfill this goal at least in part. In this case, participants are asked to read the sentences without any reference to the study episode. However, the development of conceptual implicit measures continues to represent a challenge to STI research.

Another potential limitation of our experiments, and also a general problem within STI research, is that it is not possible to guarantee that participants are not engaging in explicit impression formation processes, even when they are given memory instructions. We had made different attempts to rule out the intervention of impression formation processes in our experiments. In most experiments the probe recognition paradigm was used, a paradigm considered to discourage elaborative processing (Experiments 1, 2, 3, 4, and 6), and in Experiment 8 names were presented instead of photos in order to avoid the spontaneous activation of impression formation processes.

However, the intervention of elaborative processes may be problematic, especially in those cases in which different behaviors were used to describe the same actor (Experiments 5, 6, and 7). In these cases, there is the risk that information comprehension spontaneously entails impression formation processes. In Experiment 5, for example, the high recall of second behaviors on incongruent pairs may suggest the intervention of impression formation processes. Jerónimo (2007) has recently presented large evidence in favor of the relation between the inhibition of trait encoding and the establishment of inter-behavioral links under impression formation conditions. The possibility remains that impression formation processes intervened on these results.

The interplay between using more complex stimuli and at the same time assure that participants not engage in impression formation processes represents a major challenge for future research. It must be guaranteed that behaviors are comprehended, but at the same time it must be assured that participants do not use impression formation strategies as a way of comprehending behaviors. One possibility is to provide participants with processing goals that focus on more superficial aspects of the text, such as analyzing semantic features of the words, similar to Uleman and Moskowitz (1994, Experiment 2).

Until now, no paradigm within STI research has been developed that is immune to alternative explanations. The problem of developing paradigms that are able to rule out retrieval contamination problems and eliminate the intervention of elaborative processing has been central to STI research. Paradigms applied in the present experiments are also not free from objections. For that reason it might be interesting for future research to test the present findings using different paradigms.

3. Implications for STI Debates

Despite the points previously outlined, results from the present experiments largely support a flexible view of the STI process, using a variety of different methods and approaches. An important question is related to the implications of the present results for the most prominent STI debates. We will next consider some of these issues.

3.1. Automaticity of STIs

Some authors have suggested a characterization of the STI process as automatic (e.g., Todorov & Uleman, 2003; Winter et al., 1985). The research approach itself has been largely oriented by the automatic-controlled dichotomy, with different studies exploring whether STIs obey to the different automaticity criteria proposed by Bargh (1994).

Our results converge with a class of studies that indicate that STIs should not be characterized as an unconditional automatic process (e.g., Uleman & Moskowitz, 1994; Uleman et al., 1992; Wigboldus et al., 2003). Specifically, results from our experiments show that STI magnitude is influenced by the degree of relatedness between the behavior and the implied trait, by the category label that is associated with behavior, and by previous behavioral information that is provided about the actor.

Even more crucial, our theoretical approach departs from an analysis of the STI process according to the automatic-controlled dichotomy. Instead, our focus is on exploring the processing mechanisms that regulate STI occurrence. We think that this approach has a broader explanatory value and can provide us with more accurate information about the STI process.

3.2. Other Spontaneous Social Inferences

Three of our Experiments (Experiments 2-4) explored both the occurrence of STIs and spontaneous situational inferences (SSIs). Our results provided evidence that SSIs can occur spontaneously, consistent with previous studies (Lupfer et al., 1990; Lupfer et al., 1995; Ham & Vonk, 2003). Moreover, our results indicate that the occurrence of both STIs and SSIs is in line with achieving text local coherence. Both types of spontaneous inferences are more likely when they are consistent with the information provided about the actor. Finally, under some conditions SSIs were shown to be more likely than STIs. This happens when the behavior is ambiguous in terms of trait/situational implications, and the behavior is inconsistent with the social category of the actor.

In sum, situational inferences can occur spontaneously, are guided by consistency with known information about the actor, and in some cases are stronger than STIs. These findings are in agreement with the hypothesis that the same

mechanisms that underlie STIs occurrence may also be responsible for the occurrence of other spontaneous social inferences (for a similar point see Uleman et al., 1996). Given existing evidence for the spontaneous of other social inferences (e.g., Gernsbacher et al., 1992; Hassin et al., 2005; McKoon & Ratcliff, 1986; Winter et al., 1985), the question that naturally arises is whether STIs are different from other elaborative inferences involved in comprehension. Should the centrality of trait concepts in person perception (e.g., Fiske & Cox, 1979) be taken as an indication that STIs are operated by specific mechanisms? Or are STIs driven by the same rules that guide inference generation under the normal processes of comprehension? These are important and controversial questions that must be further considered by future research.

3.3. Debates Concerning Actor vs. Behavior and Association vs. Inference

Two crucial debates in the literature are whether STIs are characterizations of actors or mere behavior characterizations, and whether STIs have an associative or inferential nature. The present work can add some contributions into these discussions.

According to the third principle proposed, spontaneous trait inferences vary in a continuum of strength. In the lowest level of the continuum, the inferred trait can be associated with any element from the encoding context, but will be more strongly associated with the most distinctive and relevant elements (for obvious reasons, usually the *actor* of the behavior). Thus, at this level, spontaneous trait inferences are not specifically *about the actor* (since they can also be associated with other elements, as “bananas”, see Brown & Bassili, 2002), and are naturally explained by the operation of associative mechanisms.

In line with this reasoning, present results (Experiment 8) confirm that an association is established between the actor and the trait, but this association is not sufficiently strong to create expectations about the actor. The trait-actor association has also no role (not even an implicit role) in structuring or organizing the different received behavioral information about the actor (Experiment 7). Thus, at least under the present study conditions, it seems that the trait is not mentally represented as a characteristic of the actor (i.e., is not strongly associated with the actor).

At the strongest level of the continuum, trait inferences are specifically linked with the actor. In this case, traits serve as expectancies and may be deliberately accessed in subsequent tasks. The trait-actor association has now an explicit *inferential* status. At this level, the link between the trait and the actor become clearly differentiated from trait transferences to other elements of the context. Usually, the strongest level of trait inferences occurs when impression formation goals are present.

An important question is to know whether weaker inferences may occur with impression formation goals, and whether stronger trait inferences may result from spontaneous processes. According to our proposal the answer is yes. Results from Experiment 8 showed, for example, that under some impression formation conditions expectancies are not created (i.e., when only one behavior is provided about each actor and several different actors are presented), despite association between the actor and the trait being established. In a similar way, we may think that conditions in which the association between the actor and the trait is made more salient, strong trait inferences may be established, even under spontaneous inference conditions.

This view has some important consequences. First, the continuum approach opens the possibility that differences between trait inference processes between memory and impression formation might be in some cases merely a question of degree of trait-actor association strength. Second, it dilutes the distinction between associative and attributional processes proposed by some authors (Carlston & Skowronski, 2005; Crawford et al., 2007; Crawford et al., 2008), seen as responsible for STTs and STIs effects respectively. According to a continuum view, at a lower level of the continuum, the components of the processes responsible for STTs and STIs are the same (i.e., associative), only varying in strength. As a consequence of this view, we would predict that similar differences in trait-actor association strength could be observed between two different actors (one “main actor” and one “secondary actor”), and not only between actors and non-actors. Thus, a way of testing whether observed differences between STI and STT effects are merely quantitative would be by presenting a behavior performed by two different actors, having one of them a more prominent role in the action (i.e., the main actor). In this case, we would expect stronger evidence for STIs for the most central actor in the action and weaker evidence for the less relevant actor, similar to differences between actors and communicators (e.g., Skowronski et al., 1998).

However, because in this case the two individuals are actors of the behavior, results cannot be explained by qualitative processing differences, but only by quantitative differences of the same underlying process. In other words, if the same differences are observed between actors and non-actors, and between main actors and secondary actors, it is likely that the same processes operate in both conditions.

4. A closer look into the three principles underlying STI

In the present project, three principles were proposed about the mechanism underlying STI occurrence. These principles are theoretically and empirically grounded on the minimalist framework (McKoon & Ratcliff, 1992) and on the gradual view of inference generation (McKoon & Ratcliff, 1986, 1989a, 1990a). This approach represents an initial attempt to describe the processes that underlie STIs. However, we believe that these principles will be revised and specified as more data are obtained.

Certainly, there are some questions that need further clarification. For example, the mechanisms that make a trait easily inferable were not addressed. Minimalist authors have faced similar questions (e.g., McKoon et al., 1996). However, this “easily inferable” notion is easily incorporated by the general principles of global memory models (e.g., Gillund and Shiffrin 1984; Hintzman 1986, 1988; Murdock 1982). According to global memory principles, text information activates related concepts in memory in a fast and parallel way. The more the concept is activated, the more likely it is that it will be integrated into the representation of the event. The possibility that STIs may be influenced by semantic associations was also raised by Uleman and collaborators. They state that “with more lifelike nontextual stimuli, perhaps some contexts support particular trait inferences more than others because they contain strong associates of them” (Uleman et al., 1996, p. 392). However, these ideas were not further developed in subsequent studies.

Our second principle is quite straightforward. According to a local coherence view, concepts are not spontaneously activated if they make the text more difficult to understand. This is considered a basic need of the comprehension system (e.g., Graesser et al., 1994; Kintsch, 1988; 1998; Lorch & O'Brien, 1995; McNamara & Kintsch, 1996; Tapiero & Kintsch, 2007). The principle of local coherence can also be seen as being intrinsically related with the first principle. For example, according to the construction

integration model (Kintsch, 1988; 1998), the most semantically coherent concepts are naturally included in the representation of the text, through repeated cycles of parallel activation. That is, the easiest inferences are those that are more likely to contribute to a coherent picture of the processed information. And those inferences that bring coherence to the text are the ones that are easiest to infer. This would make sense to an efficient cognitive system.

It is important to notice that, as it was manipulated in the present experiments, talking about local coherence is equivalent to talk about consistency or inconsistency of the inference with existing expectations about the actor that were either activated by a category label (Experiments 2-4) or by a previous behavior (Experiment 5 and 6). However, as it is usually defined, the concept of local coherence is broader than that. Specifically, it is assumed that the main function of inferential processes is a “coherence maintaining function” (Van den Broek et al., 1995, p. 357), not only of the information about the actor, but also of situational, goal, and also causal information. Therefore, although we have only manipulated coherence *about the actor* in the present experiments, it would be important for future studies to explore how these other aspects may interplay to guide the inferential activity in general, and STIs in particular. A way to exemplify that coherence is not reduced to consistency of the information about the actor would consist in the presentation of two behaviors with consistent trait information. However, depending on the condition, the two behaviors could form a more or less coherent representation of the text. We would predict stronger STIs from the second behaviors, in the case in which the two behaviors are more coherent (i.e., are more likely to belong together). Thus, differences in STIs magnitude are predicted even under conditions in which there is no inconsistency between information provided about the actor.

Another way of demonstrating how coherence is determinant for inferential processes would be by showing that STIs are undermined when the basic structure of the text is not achieved in the first place. For example, if we read the sentence: “John climbed the stairs to the lower floor in order to receive the scientific award” we will probably not infer the trait intelligent. Specifically, the local incoherence presented in the text (“climbed the stairs to the lower floor”) will probably interfere with subsequent inferences. In this case, no incoherence exists concerning the information provided

about the actor himself (the inference intelligent is not inconsistent with any information provided about the actor). However, because the structure of the text is incoherent, the perceiver would be especially worried in clarifying the incoherency. This study represent a way of showing the primacy of local coherence in text comprehension, and to demonstrate that local coherence is in some sense more “important” to the perceiver than engaging in automatic trait inferences. Again, this would be in agreement with a flexible and conditional view of the STI process.

Another important aspect related with the coherence issue is related to the fact that we assumed that STIs are determined by local coherence, but not by text global coherence (i.e., integration of received information with previous information no longer in short term memory). There are reasons for that. While the occurrence of inferences that are necessary for local coherence is well documented; whether inferences occur automatically in order to assure global coherence is a question that remains open (Graesser et al., 1994; McKoon & Ratcliff, 1992; O’Brien & Myers, 1999; Suh & Trabasso, 1993). According to a minimalist approach, for instance, the reader is mainly concerned with achieving local coherence. However, although not directly tested in the present studies, we may hypothesize that global coherence may also contribute to the modulation of STIs. Results provided by Albrecht and O’ Brien (1993) are consistent with this hypothesis. In this study, a characteristic of an actor (e.g., Mary was a strict vegetarian) was described. After a number of sentences, a new sentence was presented that was locally coherent with previous text, but that could be either consistent, inconsistent, or neutral with the initial information, Results showed that participants took longer to read inconsistent sentences (e.g., Mary ordered a cheeseburger and fries). Thus, global coherence breaks can also potentially affect inferences that are automatically generated during comprehension (see also Myers, O’Brien, Albrecht, & Mason, 1994; O’Brien et al., 1998). This question can be an interesting venue for future studies.

Finally, the third principle postulated that spontaneous trait inferences vary along a continuum. Data from the present studies support the view that STIs are in some cases only weakly or minimally associated with the actor (Experiment 8) and results from other experiments show that STI magnitude can vary depending on the type of category label, type of behavior, and previous information associated with the actor

(Experiments 1-6). All these data are inconsistent with an all or none view of the STI process, and thus they support a continuum view of the process. However, additional studies in which more levels of the variables are manipulated are certainly necessary to prove the continuum nature of STIs.

5. Future venues

5.1. Effects of Trait Accessibility on STIs.

In Experiments 5 and 6, consistent or inconsistent pairs of behaviors were provided to participants. We interpreted the results from these experiments as an indication that the trait is less likely to be spontaneously inferred when it is incoherent with previous information (i.e., inconsistent pairs), and more likely to be inferred when it is coherent with previous information (i.e., consistent pairs) about the actor.

However, we may question whether similar effects can be observed due to the operation of more general trait accessibility effects, working independently of coherence issues. A good way of testing this hypothesis would be by presenting pairs of behaviors of different actors. If the two behaviors belong to different individuals, comprehension of one behavior is independent of the comprehension of the other behavior. Because coherence issues do not come into play under these conditions, the principles postulated in the present proposal do not allow specific predictions to be made about what would be the pattern of results under these conditions.

Notice, however, that if differences in STIs from the second behavior are observed even when the two behaviors refer to different actors, and assuming no reasons to expect coherence between the two behaviors, then it would mean that STIs may be affected by general accessibility priming effects. The possibility of STIs being influenced by trait accessibility is not contradictory with previous accounts for the flexibility of STIs (Wigboldus et al., 2003). Wigboldus et al. (2003) explained stereotype effects on STIs by the temporary accessibility of consistent traits and temporary inhibition of inconsistent traits provoked by stereotype activation. We might suppose that similar accessibility effects may be produced by different means, as for example by presenting previous behaviors that might serve as self-generated trait primes (Moskowitz & Roman, 1992; Stapel et al., 1996) or by directly priming personality

traits, either supraliminally (e.g., Higgins et al., 1977) or subliminally (e.g., Bargh & Pietromonaco, 1982). If future studies show that STI may be influenced by priming manipulations, it would suggest a picture of the STIs as highly contextual malleable.

Importantly, however, these effects would not be incompatible with a minimalist hypothesis of inference generation (McKoon & Ratcliff, 1992). According to a minimalist account, received information activates memory contents in a parallel and fast way. Concepts more active will be the more likely to be available for inference generation (see McKoon, et al., 1996). Thus, a minimalist proposal could easily integrate the occurrence of more general accessibility effects. These issues might represent important venues for future studies.

5.2. Contributions of STI research for Text Comprehension

In addition to the points open to future research outlined throughout the general discussion, a further interesting question would be to examine potential consequences of the STI research for the text comprehension literature. The approach that we followed in the present proposal was to import ideas from the text comprehension literature to the STIs field. However, we think that the intersection between the two fields can also work the other way around.

A potentially interesting question for research would be to explore whether spontaneous transference effects occur with narrative structures. While transference effects have been broadly examined within STI research, they have not been studied at all within a text comprehension framework. However, it would be important to explore whether inferences that are generated during the natural course of narrative comprehension get associated with other, irrelevant elements of the narrative. This kind of evidence would be highly informative about the mechanisms involved in normal reading. Importantly, this issue can have important implications to the current tension between constructionist and minimalist approaches. While transference effects are incompatible with a constructionist view of inference generation (Graesser et al., 1994), they can be easily integrated by a minimalist framework (McKoon & Ratcliff, 1992). According to a constructionist model, the goal of the reader is to construct a coherent and meaningful representation of the text (Graesser et al., 1994). Thus, peripheral associations between generated inferences and other irrelevant elements of the context

are not contemplated by such a view. By contrast, according with a minimalist view, inference generation is seen as being governed by the same mechanisms that rule knowledge activation in more general terms (e.g., McKoon et al., 1996). Information from the text activates memory contents in a parallel way, and activated concepts become available for inference generation. This view integrates a gradual view of inference generation, with highly activated concepts being strongly encoded and other inferences being only minimally encoded. Such a theoretical framework would easily accommodate the notion that activated concepts may establish residual associations with peripheral elements of the text.

Another possible contribution of the STI literature to the study of text comprehension processes might be the integration between narrative and non-narrative materials, such as pictures. This has been frequently done in STI research (e.g., Carlston & Skowronski, 1994; Todorov & Uleman, 2002), but not in the text comprehension field. It would be interesting to examine whether presenting pictures would have influence on the types of inferences that are automatically generated online during text comprehension. This would also represent a way of integrating the text and discourse comprehension literatures with the literature on face perception. These two domains have been quite separated, despite the natural interplay between discourse processing and face perception.

5.6. Conclusions

In conclusion, the present work adds to 25 years of research on STIs. Our major contribution consisted in demonstrating different facets of the flexibility of the STI process. Understanding when people go from the observation of actors' external behaviors to inferences about the defining traits of those actors has been a central question on social cognition research. As social perceivers, we are able to observe how people *act*, but we are unable to see with our eyes how people really *are*. This information can only be inferred. How and when we move beyond observed behavior to inferences about actors' traits is certainly an intriguing question. Our work suggests that the process by which we infer traits from behaviors is quite flexible, and presents a great variation depending on the specific encoding conditions.

REFERENCES

- Aarts, H., Dijksterhuis, A., Dik, G. (2008). Goal contagion: Inferring goals from others' actions - and what it leads to. In J. Y. Shah & W. Gardner (Eds.), *Handbook of motivation science* (pp. 265-280). New York: Guilford.
- Aarts, H., & Hassin, R. R. (2005). Automatic goal inference and contagion: On pursuing goals one perceives in other people's behavior. In J. P. Forgas, K. D. Williams, & S. M. Laham (Eds.), *Social motivation: Conscious and unconscious processes* (pp. 153-167). New York: Cambridge University Press.
- Abelson, R. P. (1981). Psychological status of the script concept. *American Psychologist*, *36*(7), 715-729.
- Abelson, R. P., & Lalljee, M. (1988). Knowledge structures and causal explanation. In D. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality*. Brighton, England: Harvester Press.
- Abraham, S.C.S. (1988). Seeing the connections in lay causal comprehension. In: D.J. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality* (pp. 175-203). Brighton, England: Harvester Press.
- Ahn, W., & Bailenson, J. (1996). Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *Cognitive Psychology*, *31*(1), 82-123.
- Ajzen, I. (1971). Attribution of dispositions to an actor: Effects of perceived decision freedom and behavioral utilities. *Journal of Personality and Social Psychology*, *18*, 144-156.
- Ajzen, I., Dalto, C. A., & Blyth, D. P. (1979). Consistency and bias in the attribution of attitudes. *Journal of Personality and Social Psychology*, *37*(10), 1871-1876.
- Alba, J. W., & Hasher, L. (1983). Is memory schematic. *Psychological Bulletin*, *93*(2), 203-231.
- Albrecht, J. E., & O'Brien, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(5), 1061-1070.
- Ames, D.R. (2005). Everyday solutions to the problem of other minds: which tools are used when? In B. F. Malle, & S.D. Hodges SD (Eds.), *Other minds: how humans bridge the divide between self and other* (158-173). New York: Guilford Press.
- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*(3), 261-295.
- Anderson, J. R. (1990). Inferences about word meanings". In A. Graesser & G. Bower (Eds.), *Inferences and Text Comprehension*. New York: Academic Press.

- Arkin, R. M., & Duval, S. (1975). Focus of Attention and Casual Attributions of Actors and Observers. *Journal of Experimental Social Psychology*, 11(5), 427-438.
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 41, 258-290.
- Asch, S. (1952). *Social Psychology*. NY: Prentice Hall.
- Baker, L., & Anderson, R. I. (1982). Effects of inconsistent information on text processing: Evidence for comprehension monitoring. *Reading Research Quarterly*, 281-294.
- Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, 10(3), 340-357.
- Balota, D. A., Flores d'Arcais, G. B. & Rayner K. (1990). (Eds.), *Comprehension processes in reading*. Hillsdale, NJ: Erlbaum.
- Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology*, 65, 272-272.
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3-51). New York: Guilford Press.
- Bargh, J. A. (1990). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition*, Vol. 2 (pp. 93-130). New York: Guilford.
- Bargh, J. A. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *The American Journal of Psychology*, 181-199.
- Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer, Jr. & T. K. Srull (Eds.), *Handbook of social cognition: Basic processes* (2nd ed., pp. 1-40). New York: Guilford Press.
- Bargh, J. A. (1997). The automaticity of everyday life. In R S. Wyer Jr. (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1-61). Mahwah, NJ: Erlbaum.
- Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process models in social psychology* (pp. 361-382). New-York: Guilford.

- Bargh, J. (2005). Bypassing the will: Towards demystifying the nonconscious control of social behavior. In R. Hassin, J. Uleman, & J. Bargh (Eds.), *The New Unconscious* (pp. 37-60). New York: Oxford University Press.
- Bargh, J. A. (2006). What have we been priming all these years? On the development, mechanisms, and ecology of nonconscious social behavior. *European Journal of Social Psychology*, *36*(2), 147-168.
- Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic goals and motives. In P. M. Gollwitzer, & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457-481). New York: Guilford Press.
- Bargh, J. A., & Chartrand, T. L. (1999). The Unbearable Automaticity of Being. *American Psychologist*, *54*(7), 462-479.
- Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin*, *126*(6), 925-945.
- Bargh, J. A., & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, *43*(3), 437-449.
- Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology*, *49*(1), 129-146.
- Bargh, J., & Williams, E. (2006). The automaticity of social life. *Current Directions in Psychological Science*, *15*, 1-4.
- Bartlett, F. C. (1932). *Remembering*. Cambridge, MA: Cambridge University Press.
- Bassili, J. N. (1989a). Trait encoding in behavior identification and dispositional inference. *Personality and Social Psychology Bulletin*, *15*(3), 285-296.
- Bassili, J. N. (1989b). Traits as action categories versus traits as person attributes in social cognition. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp.61-89). Hillsdale, NJ: Erlbaum.
- Bassili, J. N. (1993) Procedural efficiency and the spontaneity of trait inference. *Personality and Social Psychology Bulletin*, *19*, 199-204.
- Bassili, J. N., & Smith, M. C. (1986). On the spontaneity of trait attribution: Converging evidence for the role of cognitive strategy. *Journal of Personality and Social Psychology*, *50*(2), 239-245.

- Beck, A. T. (1995). Cognitive therapy: Past, present and future. In M. J. Mahoney (Ed.), *Cognitive and constructive psychotherapies: Theory, research and practice* (pp. 29–40). New York: Springer Publishing Company.
- Bem, DJ (1965). An experimental analysis of self-persuasion. *Journal of Experimental Social Psychology, 1*, 199-218.
- Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review, 74*(3), 183-200.
- Bernardo, J. M., & Smith, A. F. M. (1994). *Bayesian Theory*. Chichester: Wiley.
- Berscheid, E., Graziano, W., Monson, T., & Dermer, M. (1976). Outcome dependency: Attention, attribution, and attraction. *Journal of Personality and Social Psychology, 34*(5), 978-988.
- Bestgen, Y., & Vonk, W. (1995). The role of temporal segmentation markers in discourse processing. *Discourse Processes, 19*, 385-406.
- Bever, T.G., & McElree, B. (1988). Empty categories access their antecedents during comprehension, *Linguistic Inquiry, 19*, 34-43.
- Blair, I. V., Banaji, M. R. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology, 70*, 1142-1163.
- Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in text comprehension and memory. *Cognitive Psychology, 11*, 177-220.
- Bradbury, T. N., & Fincham, F. D. (1990). Attributions in marriage: Review and critique. *Psychological Bulletin, 107*(1), 3-33.
- Bradley, G. W. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question. *Journal of Personality and Social Psychology, 36*(1), 56-71.
- Braine, M. D. S. (1978). On the Relation between the Natural Logic of Reasoning and Standard Logic. *Psychological Review, 85*(1), 1-21.
- Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior, 11*, 717-726.
- Brewin, C. R. (1986). Internal attribution and self-esteem in depression: A theoretical note. *Cognitive Therapy and Research, 10*(4), 469-475.
- Britton BK, Graesser AC, (Eds.) (1996). *Models of Understanding Text*. Mahwah, NJ: Erlbaum.

- Brown, R. D., & Bassili, J. N. (2002). Spontaneous trait associations and the case of the superstitious banana. *Journal of Experimental Social Psychology*, 38(1), 87-92.
- Bruner, J. S. (1957). Going beyond the information given. In H. Gruber, G. Terrell, & M. Wertheimer (Eds.), *Contemporary approaches to cognition* (pp. 258-290). Cambridge, MA: Harvard University Press.
- Burt, H. E. (1941). An experimental study of early childhood memory: Final report. *Journal of Genetic Psychology*, 58, 435-439.
- Butler, L. T., & Berry, D. C. (2001). Implicit memory: Intention and awareness revisited. *Trends in Cognitive Sciences*, 5(5), 192-197.
- Caldwell, T. L., & Newman, L. S. (2005). The timeline of threat processing in repressors: More evidence for early vigilance and late avoidance. *Personality and Individual Differences*, 38(8), 1957-1967.
- Carlston, D. E., & Skowronski, J. J. (1994). Savings in the Relearning of Trait Information as Evidence for Spontaneous Inference Generation. *Journal of Personality & Social Psychology*, 66(5), 840-856.
- Carlston, D. E., & Skowronski, J. J. (2005). Linking Versus Thinking: Evidence for the Different Associative and Attributional Bases of Spontaneous Trait Transference and Spontaneous Trait Inference. *Journal of Personality & Social Psychology* 89(6), 884-898.
- Carlston, D. E., Skowronski, J. J., & Sparks, C. (1995). Savings in relearning: II. On the formation of behavior-based trait associations and inferences. *Journal of Personality and Social Psychology*, 69(3), 420-436.
- Carreiras, M., Garnham, A., Oakhill, J., & Cain, K. (1996). The use of stereotypical gender information in constructing a mental model: Evidence from English and Spanish. *Quarterly Journal of Experimental Psychology*, 49A, 639-663.
- Chaiken, S., & Trope, Y. (Eds.) (1999). *Dual-process theories in social psychology*. New York: Guilford Press.
- Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, 74(3), 271-280.
- Charniak, E. (1977). Framed PAINTING: The Representation of a Common Sense Knowledge Fragment. *Cognitive Science: A Multidisciplinary Journal*, 1(4), 355-394.
- Charniak, E. (1978). On the use of framed knowledge in language comprehension. *Artificial Intelligence*, 11(3), 225-265.

- Chartrand, T. L., & Bargh, J. A. (1996). Automatic activation of impression formation and memorization goals: Nonconscious goal priming reproduces effects of explicit task instructions. *Journal of Personality and Social Psychology*, *71*, 464-478.
- Claeys, W. (1990). On the spontaneity of behaviour categorization and its implications for personality measurement. *European Journal of Personality*, *4*(3), 173-186.
- Clark, H. H., & Lucy, P. (1975). Understanding what is meant from what is said: A study in conversationally conveyed requests. *Journal of Verbal Learning and Verbal Behavior*, *14*(1), 56-72.
- Clark, H. H., & Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory & Cognition*, *7*(1), 35-41.
- Clary, E. G., & Tesser, A. (1983). Reactions to unexpected events: The naive scientist and interpretive activity. *Personality and Social Psychology Bulletin*, *9*(4), 609-620.
- Claus, B., & Kelter, S. (2006). Comprehending narratives containing flashbacks: Evidence for temporally organized representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(5), 1031-1044.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407-428.
- Corbett, A. T. (1984). Comprehension Prenominal adjectives and the disambiguation of anaphoric nouns. *Journal of Verbal Learning and Verbal Behavior*, *23*, 683-695
- Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: accessing potential antecedents. *Memory & Cognition*, *11*(3), 283-294.
- Corbett, A. T., & Doshier, B. A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior*, *17*, 479-491.
- Corneille, O., Leyens, J. P., Yzerbyt, V. Y., & Walther, E. (1999). Judgeability concerns: The interplay of information, applicability, and accountability in the overattribution bias. *Journal of Personality and Social Psychology*, *76*(3), 377-387.
- Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*, 268-294.
- Crawford, M. T., Skowronski, J. J., & Stiff, C. (2007). Limiting the spread of spontaneous trait transference. *Journal of Experimental Social Psychology*, *43*(3), 466-472.

- Crawford, M. T., Skowronski, J. J., Stiff, C., & Leonards, U. (2008). Seeing, but not thinking: Limiting the spread of spontaneous trait transference II. *Journal of Experimental Social Psychology, 44*(3), 840-847.
- Crawford, M. T., Skowronski, J. J., Stiff, C., & Scherer, C. R. (2007). Interfering with inferential, but not associative, processes underlying spontaneous trait inference. *Personality and Social Psychology Bulletin, 33*(5), 677-690.
- Crocker, J., Hannah, D. B., & Weber, R. (1983). Person memory and causal attributions. *Journal of Personality and Social Psychology, 44*(1), 55-66.
- Csibra, G., Gergely, G., Biro, S., Koos, O., & Brockbank, M. (1999). Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition, 72*(3), 237-267.
- Cullingford, R. (1978). Script application: computer understanding of newspaper stories. *Research Report 116*. Department of Computer Science, Yale University.
- D'Agostino, P. R., & Beegle, W. (1996). A reevaluation of the evidence for spontaneous trait inferences. *Journal of Experimental Social Psychology, 32*(2), 153-164.
- Deaton, J. A., & Gernsbacher, M. A. (1997). Causal conjunctions and implicit causality cue mapping in sentence comprehension. *Journal of Memory and Language, 36*, 293-310.
- DeJong, G. (1979). Prediction and substantiation: A new approach to natural language processing. *Cognitive Science: A Multidisciplinary Journal, 3*(3), 251-271.
- Dell, G. S., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior, 22*, 121-132.
- Dennett, D. C. (1984). *Elbow room: The varieties of free will worth having*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1991). *Consciousness Explained*. New York: Penguin Books.
- De Vega, M. (1995). Backward updating of mental models during continuous reading of narratives. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*, 373-373.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*(1), 5-18.
- Dewhurst, S. A., Holmes, S. J., Swannell, E. R., & Barry, C. (2008). Beyond the text: Illusions of recollection caused by script-based inferences. *European Journal of Cognitive Psychology, 20*(2), 367-386.

- Diener, C. I., & Dweck, C. S. (1978). An analysis of learned helplessness: Continuous changes in performance, strategy, and achievement cognitions following failure. *Journal of Personality and Social Psychology*, 36(5), 451-462.
- Dijksterhuis, A., & Van Knippenberg, A. D. (1996). The knife that cuts both ways: Facilitated and inhibited access to traits as a result of stereotype activation. *Journal of Experimental Social Psychology*, 32(3), 271-288.
- Dijksterhuis, A., & Van Knippenberg, A. (1998). The relation between perception and behavior, or how to win a game of trivial pursuit. *Journal of Personality and Social Psychology*, 74(4), 865-877.
- Dooling, D. J., & Lachman, R. (1971). Effects of comprehension on retention of prose. *Journal of Experimental Psychology*, 88(2), 216-222.
- Dopkins S. (1996). Representation of superordinate goal inferences in memory. *Discourse Processes*, 21, 85-104.
- Dosher, B. A., & Corbett, A. T. (1982). Instrument inferences and verb schemata. *Memory & Cognition*, 10(6), 531-539.
- Duffy, S. A. (1986). Role of expectations in sentence integration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12(2), 208-219.
- Duffy, S., & Rayner, K. (1990). Eye movements and anaphor resolution: Effects of antecedent typicality and distance. *Language and Speech*, 33(2), 103-119.
- Ebbinghaus, H. (1885/1964). *Memory: A contribution to experimental psychology*. New York: Dover Publications.
- Ehrlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. *Journal of Verbal Learning and Verbal Behavior*, 21(3), 296-306.
- Ellis, R. J., & Holmes, J. G. (1982). Focus of attention and self-evaluation in social interaction. *Journal of Personality and Social Psychology*, 43(1), 67-77.
- Emmott, C. (1997). *Narrative comprehension: A discourse perspective*. Oxford: Clarendon Press.
- Ender, P. B., & Bohart, A. C. (1974). Attributions of success and failure. *Psychological Reports*, 35(1), 275-278.
- Enzle, ME, & Shopflocher, D. (1978). Instigation of attribution processes by attributional questions. *Personality and Social Psychology Bulletin*, 4, 595-599.
- Epstein, R. (1997). Folk wisdom: Was your grandmother right? *Psychology Today*, 76, 46-50.

- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, *102*(2), 211-244.
- Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, *87*(3), 215-251.
- Erickson, D. J., & Krull, D. S. (1999). Distinguishing judgments about what from judgments about why: Effects of behavior extremity on correspondent inferences and causal attributions. *Basic and Applied Social Psychology*, *21*(1), 1-11.
- Eysenck, M.W. Keane, M.T. (2005) *Cognitive Psychology: A Student's Handbook (5th Edition)*. London: Psychology Press.
- Farah, M. J. (2008). Neuroethics and the problem of other minds: Implications of neuroscience for the moral status of brain-damaged patients and nonhuman animals. *Neuroethics*, *1*(1), 9-18.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). "What is special about face perception?". *Psychological Review*, *105*, 482-498.
- Fazio, R. H, Jackson, J., Dunton, B., & Williams, C. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013-1027.
- Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology*, *70*, 1164-1184.
- Fein, S., Hilton, J. L., & Miller, D. T. (1990). Suspicion of ulterior motivation and the correspondence bias. *Journal of Personality and Social Psychology*, *58*(5), 753-764.
- Ferreira, F., Bailey, K. G. D., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, *11*, 11-15.
- Ferreira, M., Garcia-Marques, L., Ramos, T., Hamilton, D., & Jerónimo, R (*Manuscript in preparation*). On The Relation between Spontaneous Trait Inferences and Intentional Inferences: An Activation-Monitoring Hypothesis.
- Fiske, S. T. (1989). Examining the role of intent: Toward understanding its role in stereotyping and prejudice. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 253-283). New York: Guilford.
- Fiske, S. T., & Cox, M. G. (1979). Person concepts: The effect of target familiarity and descriptive purpose on the process of describing others. *Journal of Personality*, *47*(1), 136-161.

- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences, 11*(2), 77-83.
- Fiske, S. T., & Taylor, S. E. (1991). *Social cognition*. New York: McGraw-Hill.
- Fleming, J. H., & Darley, J. M. (1989). Perceiving choice and constraint: The effects of contextual and behavioral cues on attitude attribution. *Journal of Personality and Social Psychology, 56*(1), 27-40.
- Ford, T. E., & Kruglanski, A. W. (1995). Effects of epistemic motivations on the use of accessible constructs in social judgment. *Personality and Social Psychology Bulletin, 21*(9), 950-962.
- Ford, J. A., & Milosky, L. M. (2008). Inference generation during discourse and its relationship to social competence: An online investigation of abilities of children with and without language impairment. *Journal of Speech, Language, and Hearing Research, 57*, 367-386.
- Forgas, J. P. (1998). On being happy and mistaken: Mood effects on the fundamental attribution error. *Journal of Personality and Social Psychology, 75*(2), 318-331.
- Forster, K. I. (1981). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *The Quarterly Journal of Experimental Psychology Section A, 33*(4), 465-495.
- Francolini, C. M., & Egeth, H. E. (1980). On the nonautomaticity of "automatic" activation: evidence of selective seeing. *Perception & Psychophysics, 27*(4), 331-342.
- Galambos, J. A., Abelson, R. P., & Black J. B. (1986). *Knowledge Structures*. NJ: Lawrence Erlbaum Associates.
- Garcia-Marques, L. (1988). O modelo Anova: Análise, críticas e alternativas. *Análise Psicológica, 3-4*(VI), 321-346.
- Garcia-Marques, L., & Garcia-Marques, T. (2003). O legado de Harold H. Kelley: O estudo da interdependência e atribuição de causas. *Análise Psicológica, 21*(3), 259-266.
- Garcia-Marques, L., & Hamilton, D. L. (1996). Resolving the apparent discrepancy between the incongruency effect and the expectancy-based illusory correlation effect: The TRAP model. *Journal of Personality and Social Psychology, 71*, 845-860.
- Gawronski, B. (2003). Implicational schemata and the correspondence bias: On the diagnostic value of situationally constrained behavior. *Journal of Personality and Social Psychology, 84*(6), 1154-1171.

- Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European Review of Social Psychology, 15*, 183-217.
- Gergely, G., Nadasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition, 56*(2), 165-193.
- Gernsbacher, M.A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.
- Gernsbacher, M. A. (1994). Activating knowledge of fictional characters' emotional states. In C. A. Weaver, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 141-155). Hillsdale, NJ: Erlbaum.
- Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*(2), 245-262.
- Gernsbacher, M. A., Goldsmith, H. H., & Robertson, R. R. W. (1992). Do readers mentally represent characters' emotional states? *Cognition & Emotion, 6*(2), 89-111.
- Gernsbacher, M. A., Hallada, B. M., & Robertson, R. R. W. (1998). How Automatically Do Readers Infer Fictional Characters Emotional States? *Scientific Studies of Reading, 2*(3), 271-300.
- Gernsbacher, M. A., & Robertson, R. R. W. (1992). Knowledge activation versus sentence mapping when representing fictional characters' emotional states. *Language and Cognitive Processes, 7*(3), 353-371.
- Gibbs, R. (1979). Contextual effects in understanding indirect requests. *Discourse Processes, 2*, 1-10.
- Gibbs, R. (1983). Do people always process the literal meanings of indirect requests? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 524-533.
- Gibbs, R. (1984). Literal meaning and psychological theory. *Cognitive Science, 8*, 275-304.
- Gibbs, R. (1986). Comprehension and memory for nonliteral utterances: The problem of sarcastic indirect requests. *Acta Psychologica, 53*, 41-57.
- Gibbs, R. (1994). *The poetics of mind: Figurative thought, language, and understanding*. New York: Cambridge University Press.

- Gilbert, D. T. & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology*, 60, 509-517.
- Gilbert, D. T., & Jones, E. E. (1986). Perceiver-induced constraint: Interpretations of self-generated reality. *Journal of Personality and Social Psychology*, 50(2), 269-280.
- Gilbert, D. T., Krull, D. S., & Pelham, B. W. (1988). Of thoughts unspoken: Social inference and the self-regulation of behavior. *Journal of Personality and Social Psychology*, 55(5), 685-694.
- Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin*, 117(1), 21-38.
- Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, 54(5), 733-740.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91(1), 1-67.
- Gilovich T, Griffin, D.W., & Kahneman D. (Eds.) (2002). *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge Univ. Press.
- Glenberg, A.M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Language*, 26, 69-83.
- Goedert, K. M., & Spellman, B. A. (2005). Nonnormative discounting: There is more to cue interaction effects than controlling for alternative causes. *Learning & Behavior*, 33(2), 197-210.
- Goldam, S. R., Graesser, A. C., & Van der Broek, P. (1999). *Narrative comprehension, Causality, and Coherence: Essays in Honor of Tom Trabasso*. NJ: Lawrence Erlbaum Associates.
- Gollan, T., & Witte, E. H. (2008). "It was right to do it, because...": Understanding Justifications of Actions as Prescriptive Attributions. *Social Psychology*, 39(3), 189-196.
- Gordon, S. E., & Wyer, R. (1987). Person memory: Category-set-size effects on the recall of a person's behaviors. *Journal of Personality and Social Psychology*, 53(4), 648-662.
- Goren, A., & Todorov, A. (2009). Two Faces Are Better Than One: Eliminating False Trait Associations With Faces. *Social Cognition*, 27(2), 222-248.

- Graesser, A.C. & Bower, G.H., (Eds.) (1990). *Inferences and Text Comprehension*. New York: Academic Press.
- Graesser, A. C., Gernsbacher, M. A., & Goldman, S. R. (2003). *Handbook of discourse processes*. Mahwah, NJ: Erlbaum.
- Graesser, A. C., Gordon, S. E., & Sawyer, J. D. (1979). Recognition memory for typical and atypical actions in scripted activities: Tests of a script pointer+ tag hypothesis. *Journal of Verbal Learning and Verbal Behavior*, 18(3), 319-332.
- Graesser, A. C., Millis, K. K., & Zwaan, R. A. (1997). Discourse comprehension. *Annual Review of Psychology*, 48(1), 163-189.
- Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3), 371-395.
- Graesser, A. C., Swamer, S. S., Baggett, W. B., & Sell, M. A. (1996). New models of deep comprehension. In B. K. Britton & A. C. Graesser (Eds.), *Models of understanding text* (pp. 1-32). Hillsdale, NJ: Erlbaum.
- Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory*, 6(5), 503-515.
- Granger, R., Eiselt, K., & Holbrook, J. (1986). Parsing with parallelism: a spreading-activation model of inference processing during text understanding. In J. L. Kolodner & C. K. Riesbeck (Eds), *Experience, Memory, and Reasoning* (pp.227-246). NJ: L. Erlbaum.
- Greene, S.B., Gerrig, R.J., McKoon, G., & Ratcliff, R. (1994). Unheralded pronouns and the management of common ground. *Journal of Memory and Language*, 33, 511-526.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4-26.
- Gygax, P., Oakhill, J., & Garnham, A. (2003). The representation of characters emotional responses: Do readers infer specific emotions? *Cognition & Emotion*, 17(3), 413-428.
- Haenggi, D., Kintsch, W., & Gernsbacher, M. A. (1995). Spatial situation models and text comprehension. *Discourse Processes*, 19, 173-199.
- Hakala C. M., & O'Brien E. J. (1995). Strategies for resolving coherence breaks in reading. *Discourse Processes*. 20, 167-86

- Halberstadt, J. B., & Niedenthal, P. M. (2001). Effects of emotion concepts on perceptual memory for emotional expressions. *Journal of Personality and Social Psychology*, 81(4), 587–598.
- Ham, J., & Van den Bos, K. (2008). Not fair for me! The influence of personal relevance on social justice inferences. *Journal of Experimental Social Psychology*, 44(3), 699-705.
- Ham, J., & Vonk, R. (2003). Smart and easy: Co-occurring activation of spontaneous trait inferences and spontaneous situational inferences. *Journal of Experimental Social Psychology*, 39(5), 434-447.
- Hamilton, D. L. (1981). Stereotyping and intergroup behavior: Some thoughts on the cognitive approach. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 333-354). Hillsdale, NJ: Erlbaum.
- Hamilton, D.L. (1988). Causal attributions viewed from an information-processing perspective. In D. Bar-Tal & A.W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 369-385). Cambridge, England: Cambridge University Press.
- Hamilton, D.L. (1989). Understanding impression formation: What has memory research contributed? In P.R. Solomon, G.R. Goethals, C.M. Kelley, & B.R. Stephens (Eds), *Memory: Interdisciplinary approaches* (pp. 221-242). New York: Springer-Verlag,
- Hamilton, D. L. (1998). Dispositional and attributional inferences in person perception. In J. M. Darley & J. Cooper (Eds.), *Attribution and social interaction* (pp. 99-114). Washington, DC: American Psychological Association.
- Hamilton, D. L. (Ed.). (2005). *Social cognition: Key Readings*. New York: Psychology Press.
- Hamilton, D. L., Driscoll, D. M., & Worth, L. T. (1989). Cognitive organization of impressions: Effects of incongruency in complex representations. *Journal of Personality and Social Psychology*, 57(6), 925-939.
- Hamilton D. L., Gifford R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12, 392-407.
- Hamilton, D. L., Katz, L. B., & Leirer, V. O. (1980). Cognitive representation of personality impressions: Organizational processes in first impression formation. *Journal of Personality and Social Psychology*, 39(6), 1050-1063.
- Hamilton, D. L., & Rose, T. L. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology*, 39(5), 832-845.

- Hamilton, D. L., & Sherman, S. J. (1996). Perceiving persons and groups. *Psychological Review*, *103*(2), 336-355.
- Hansen, R. D., & Hall, C. A. (1985). Discounting and Augmenting Facilitative and Inhibitory Forces: The Winner Takes Almost All. *Journal of Personality and Social Psychology*, *49*(6), 1482-1493.
- Harvey, J. H., & McGlynn, R. P. (1982). Matching words to phenomena: The case of the fundamental attribution error. *Journal of Personality and Social Psychology*, *43*(2), 345-346.
- Harvey, J. H., Town, J. P., & Yarkin, K. L. (1981). How fundamental is" the fundamental attribution error. *Journal of Personality and Social Psychology*, *40*(2), 346-349.
- Harvey, J. H., & Weary, G. (1984). Current issues in attribution theory and research. *Annual Review of Psychology*, *35*(1), 427-459.
- Harvey, J. H., Yarkin, K. L., Lightner, J. M., & Town, J. P. (1980). Unsolicited interpretation and recall of interpersonal events. *Journal of Personality and Social Psychology*, *38*(4), 551-568.
- Hassin, R. R. (2005). Non-conscious control and implicit working memory. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The New Unconscious* (pp. 196-224). New York: Oxford University Press.
- Hassin, R. R., Aarts, H., & Ferguson, M. J. (2005). Automatic goal inferences. *Journal of Experimental Social Psychology*, *41*(2), 129-140.
- Hassin, R. R., Bargh, J. A., & Uleman, J. S. (2002). Spontaneous causal inferences. *Journal of Experimental Social Psychology*, *38*, 515-522.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, *46*(1), 44-56.
- Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, *37*(1), 25-38.
- Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*. *13*(5), 512-521.
- Heider, F. (1944). Social perception and phenomenal causality. *Psychological Review*, *51*(6), 358-374.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.

- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 243-259.
- Hemsley, G. D., & Marmurek, H. H. C. (1982). Person Memory the Processing of Consistent and Inconsistent Person Information. *Personality and Social Psychology Bulletin*, 8(3), 433-438.
- Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology*, 51(6), 1106-1115.
- Hewstone, M. (1989). *Causal attribution: From cognitive processes to collective beliefs*. Cambridge, MA: Basil Blackwell.
- Higgins, T., E. (2000). Social cognition: Learning about what matters in the social world. *European Journal of Social Psychology*, 30(1), 3-39.
- Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology*, 38(1), 369-425.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13(2), 141-154.
- Hilton, D. J., Smith, R. H., & Kim, S. H. (1995). Processes of causal explanation and dispositional attribution. *Journal of Personality and Social Psychology*, 68, 377-387.
- Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review*, 93(4), 411-428.
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95(4), 528-551.
- Hoffman, C., Mischel, W., & Mazze, K. (1981). The role of purpose in the organization of information about behavior: Trait-based versus goal-based categories in person cognition. *Journal of Personality and Social Psychology*, 40(2), 211-225.
- Houriham, K. L., & MacLeod, C. M. (2007). Capturing conceptual implicit memory: the time it takes to produce an association. *Memory & Cognition*, 35(6), 1187-1196.
- Ichheiser, G. (1949). Misunderstandings in human relations: A study in false social perception. *American Journal of Sociology*, 55, 1-70.
- Idson, C. L., & Mischel, W. (2001). The personality of familiar and significant people: The lay perceiver as a social-cognitive theorist. *Journal of Personality and Social Psychology*, 80(4), 585-596.

- Irmen, L., & Roßberg, N. (2004). Gender markedness of language: The impact of grammatical and nonlinguistic information on the mental representation of person information. *Journal of Language and Social Psychology, 23*, 272-307.
- Jacoby, L. L. (1983). Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior, 22*, 485-508.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*(5), 513-541.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110*(3), 306-340.
- Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology General, 122*, 139-139.
- Jacoby, L. L., Yonelinas, P., & Jennings, J. M. (1997). The relation between conscious and unconscious (automatic) influences: a declaration of independence. In J. D. Cohen, & J. W. Schooler (Eds.), *Scientific approaches to consciousness* (pp. 13-48). Mahwah, NJ: Lawrence Erlbaum Associates Inc.
- James, W. (1948). *Psychology*. New York: World. (Original work published 1890).
- Jarvis, B. (2002). MediaLab Research Software (Version 5.3)[Computer software]. New York: Empirisoft.
- Jaspars, J., Fincham, F. D., & Hewstone, M. (1983). *Attribution theory and research*. London: Academic Press.
- Jerónimo, R. (2007). *The secret life of incongruency: From trait inference to trait inhibition in impression formation*. Doctoral Dissertation. Lisbon University Institute, Lisbon.
- Johnson, M. K. (1997). Source monitoring and memory distortion. *Philosophical Transactions: Biological Sciences, 352*(1362), 1733-1745.
- Johnson, M. H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition, 40*, 1-19.
- Johnson, M. K., Bransford, J. D., & Solomon, S. K. (1973). Memory for tacit implications of sentences. *Journal of Experimental Psychology, 98*(1), 203-205.
- Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3-28.

- Johnson, J. T., Jemmott, J., & Pettigrew, T. F. (1984). Causal attribution and dispositional inference: Evidence of inconsistent judgments. *Journal of Experimental Social Psychology, 20*(6), 567-585.
- Johnson-Laird, P. N. (1983). *Mental Models*. Cambridge, MA: Harvard University Press.
- Johnson-Laird, P. N. (1987). The mental representation of the meaning of words. *Cognition, 25*(1-2), 189-211.
- Jones, E. E. (1979). The rocky road from acts to dispositions. *The American Psychologist, 34*(2), 107-117.
- Jones, E. E. (1990). *Interpersonal perception*. New York: Freeman.
- Jones, E. E., & Berglas, S. (1976). A recency effect in attitude attribution. *Journal of Personality, 44*(3), 433-448.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219-266). New York: Academic Press.
- Jones, E. E., Davis, K. E., & Gergen, K. J. (1961). Role playing variations and their informational value for person perception. *Journal of Abnormal and Social Psychology, 63*, 302-310.
- Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology, 3*(1), 1-24.
- Jones, E. E., & McGillis, D. (1976). Correspondent inferences and the attribution cube: A comparative reappraisal. In J.H. Harvey, W.J. Ickes, & R.F. Kidd (Eds.). *New directions in attribution research* (Vol. 1, pp.389-420). Hillsdale, NJ: Erlbaum.
- Jones, E. E., & Nisbett, R. E. (1971). *The actor and the observer: Divergent perceptions of the causes of behavior*. New York: General Learning Press.
- Jones, E. E., Rock, L., Shaver, K. G., Goethals, G. R., & Ward, L. M. (1968). Pattern of performance and ability attribution: An unexpected primacy effect. *Journal of Personality and Social Psychology, 10*(4), 317-340.
- Jones, E. E., Worchel, S., Goethals, G. R., & Grumet, J. F. (1971). Prior expectancy and behavioral extremity as determinants of attitude attribution. *Journal of Experimental Social Psychology, 7*(1), 59-80.
- Kanazawa, S. (1992). Outcome or expectancy? Antecedent of spontaneous causal attribution. *Personality and Social Psychology Bulletin, 18*(6), 659-668.
- Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.

- Kahneman, D., & Chajczyk, D. (1983). Tests of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 9(4), 497-509.
- Kahneman D., & Henik A. (1981). Perceptual organization and attention. In M. Kubovy, & J.R. Pomerantz (Eds), *Perceptual Organization* (pp. 181–211). Hillsdale, NJ: Erlbaum.
- Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman, D. R. Davies, & J. Beatty (Eds.), *Varieties of attention* (pp.29-61). New York: Academic Press.
- Kassin, S. M. (1979). Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology*, 37(11), 1966-1981.
- Keenan, J. M., Bailett, S. D., & Brown, P. (1984). The effects of causal argument overlap on comprehension and memory. *Journal of Verbal Learning and Verbal Behavior*, 23, 115-126.
- Keenan, J. M., Potts, G. R., Golding, J. M., & Jennings, T. M. (1990). Which elaborative inferences are drawn during reading? A question of methodologies. In D. A. Balotta, G. B. Flores dlArcais, & K. Rayner (Eds.), *Comprehension processes in reading* (pp.377-402). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska Symposium on Motivation* (Vol. 15, pp. 192-240). Lincoln: University of Nebraska Press.
- Kelley, H. H. (1971). *Attribution in social interaction*. Morristown, New Jersey: General Learning Press.
- Kelley, H. H. (1972) Attribution in social interaction, In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 1-26). Morristown, NJ: General Learning Press.
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, 28(2), 107-128.
- Kelley, H. H., & Michela, J. L. (1980). Attribution theory and research. *Annual Review of Psychology*, 31(1), 457-501.
- Kiesler, C. A., Nisbett, R. E., & Zanna, M. P. (1969). On inferring one's beliefs from one's behavior. *Journal of Personality and Social Psychology*, 11(4), 321–327.

- Kintsch, W. (1988). The role of knowledge in discourse comprehension construction-integration model. *Psychological Review*, *95*, 163-182.
- Kintsch, W. (1998) *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.
- Kintsch, W. (2000). Metaphor comprehension: A computational theory. *Psychonomic Bulletin and Review*, *7*(2), 257-266.
- Kintsch, W., & Mross, E. F. (1985). Context effects in word identification. *Journal of Memory and Language*, *24*(3), 336-349.
- Klein, S. B., & Loftus, J. (1990). Rethinking the role of organization in person memory: An independent trace storage model. *Journal of Personality and Social Psychology*, *59*(3), 400-410.
- Klin, C. M., Guzmán, A. E., Weingartner, K. M., & Ralano, A. S. (2006). When anaphor resolution fails: Partial encoding of anaphoric inferences. *Journal of Memory and Language*, *54*(1), 131-143.
- Klin, C. M., & Myers, J. L. (1993). Reinstatement of causal information during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(3), 554-560.
- Knobe, J., & Malle, B. F. (2002). Self and other in the explanation of behavior: 30 years later. *Psychologica Belgica*, *42*, 113-130.
- Krull, D. S. (1993). Does the grist change the mill? The effect of the perceiver's inferential goal on the process of social inference. *Personality and Social Psychology Bulletin*, *19*(3), 340-348.
- Krull, D. S. (2001). On partitioning the fundamental attribution error: Dispositionalism and the correspondence bias. In G. B. Moskowitz (Ed.), *Cognitive social psychology* (pp. 211–227). Mahwah, NJ: Erlbaum.
- Krull, D. S., & Erickson, D. J. (1995). Judging situations: On the effortful process of taking dispositional information into account. *Social cognition*, *13*(4), 417-438.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480-498.
- Kunda, Z. (1999). *Social cognition: making sense of people*. Cambridge, MA: MIT Press.
- LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, *6*(2), 293-323.

- Lampinen, J. M., Faries, J. M., Neuschatz, J. S., & Toggia, M. P. (2000). Recollections of things schematic: The influence of scripts on recollective experience. *Applied Cognitive Psychology, 14*(6), 543-554.
- Lau, R. R., & Russell, D. (1980). Attributions in the sports pages. *Journal of Personality and Social Psychology, 39*(1), 29-38.
- Lay, C. H., Burron, B. F., & Jackson, D. N. (1973). Base rates and informational value in impression formation. *Journal of Personality and Social Psychology, 28*(3), 390-395.
- Lay, C., Ziegler, M., Hershfield, L., & Miller, D. (1974). The perception of situational consistency in behaviour: Assessing the actor-observer bias. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement, 6*(4), 376-384.
- Lea, R. B. (1995). Lexical decision evidence for spontaneous trait inference. *Unpublished raw data.*
- Leddo J. & Abelson R.P. (1986). The Nature of explanations. In J.A., Galambos, R.P. Abelson, & Black, J.B. (Eds). *Knowledge structures* (pp.103-122). Hillsdale, New Jersey: LEA.
- Leddo, J., Abelson, R. P., & Gross, P. H. (1984). Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology, 47*(5), 933-943.
- Lepore, L., Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology, 72*, 275-287.
- Lewin, K. (1931). The conflict between Aristotelian and Galilean modes of thought in contemporary psychology. *Journal of Genetic Psychology, 5*, 141-177.
- Lewin, K. (1951). *Field theory in social science*. New York: Harper & Brothers.
- Leyens, J., Yzerbyt, V., & Corneille, O. (1996). The role of applicability in the emergence of the overattribution bias. *Journal of Personality and Social Psychology, 70*, 219-229.
- Locke, V., & Walker, I. (1999). Stereotyping, processing goals and social identity: Inveterate and fugacious characteristics of stereotypes. In D. Abrams & M. A. Hogg (Eds.), *Social identity and social cognition* (pp. 164-182). Oxford: Blackwell.
- Logan, G. D. (1979). On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance, 5*(2), 189-207.

- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology*, *12*, 523-553.
- Logan, G. D. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology*, *39*, 367-386.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295-327.
- Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin*, *13*(3), 411-429.
- Long, D. L., & Golding, J. M. (1993). Superordinate Goal Inferences: Are They Automatically Generated during Comprehension? *Discourse Processes*, *16*, 55-73.
- Long, D. L., Golding, J. M., & Graesser, A. C. (1992). A test of the on-line status of goal-related inferences. *Journal of Memory and Language*, *31*(5), 634-647.
- Long, D., Golding, J., Graesser, A. C., & Clark, L. F. (1990). Goal, event, and state inferences: An investigation of inference generation during story comprehension. In A. C. Graesser & G. H. Bower (Eds.), *The Psychology of learning and motivation* (Vol. 25, pp. 89-102). New York: Academic Press.
- Lorch, R.F. & O'Brien, J.D. (1995) (Eds.). *Sources of Coherence in Reading*. Mahwah, NJ: Erlbaum.
- Lupfer, M. B., Clark, L. F., & Hutcherson, H. W. (1990). Impact of Context on Spontaneous Trait and Situational Attributions. *Journal of Personality & Social Psychology*, *58*(2), 1239-1249.
- Lutz, M. F., & Radvansky, G. A. (1997). The fate of completed goal information in narrative comprehension. *Journal of Memory and Language*, *36*(2), 293-310.
- MacLeod, C. M. (1991). Half Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin*, *109*(2), 163-203.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Calvini, G. (1999). Seeing more than we can know: Visual attention and category activation. *Journal of Experimental Social Psychology*, *35*(6), 590-602.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology*, *67*, 808-817.

- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., Thorn, T. M. J., & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology, 33*, 471-489.
- Mae, L., Carlston, D. E., & Skowronski, J. J. (1999). Spontaneous Trait Transference to Familiar Communicators: Is a Little Knowledge a Dangerous Thing? *Journal of Personality & Social Psychology, 77*(2), 233-246.
- Magliano, J. P., Baggett, W. 8., Johnson, 8. K., & Graesser, A C. (1993). The time course of generating causal antecedent and causal consequence inferences. *Discourse Processes, 16*, 35-53.
- Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review, 3*(1), 23-48.
- Malle, B. F. (2004). *How the mind explains behavior: Folk explanations, meaning, and social interaction*. Cambridge: MIT Press.
- Malle, B. F. (2006). The actor-observer asymmetry in attribution: A (surprising) meta-analysis. *Psychological Bulletin, 132*(6), 895-919.
- Malle, B. F. (2008). Fritz Heider's legacy: Celebrated insights, many of them misunderstood. *Social Psychology, 39*(3), 163-173.
- Malle, B. F., & Knobe, J. (1997). Which behaviors do people explain? A basic actor-observer asymmetry. *Journal of Personality and Social Psychology, 72*, 288-304.
- Mandler, J. M. (1986). On the comprehension of temporal order. *Language and Cognitive Processes, 1*, 309-320.
- Massad, P. M., & Hulsey, T. L. (2006). Causal attributions in posttraumatic stress disorder: Implications for clinical research and practice. *Psychotherapy: Theory, Research, Practice, Training, 43*(2), 201-215.
- McArthur, L. Z. (1972). The how and what of why: Some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology, 22*, 171-193.
- McArthur, L. Z. (1976). The lesser influence of consensus than distinctiveness information on causal attributions: A test of the person-thing hypothesis. *Journal of Personality and Social Psychology, 33*(6), 733-742.
- McArthur, L. Z., & Post, D. L. (1977). Figural emphasis and person perception. *Journal of Experimental Social Psychology, 13*(6), 520-535.
- McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86*(4), 287-330.

- McDermott, K. B. (1996). The persistence of false memories in list recall. *Journal of Memory and Language*, 35(2), 212-230.
- McEleney, A., & Byrne, R. M. J. (2006). Spontaneous causal and counterfactual thoughts. *Thinking and Reasoning*, 12(2), 235-255.
- McGuire, W. J. (1997). Going beyond the banalities of bubba-psychology: A perspectivist social psychology. In C. McGarty, & S.A. Haslam (Eds.), *The message of social psychology: Perspectives on mind in society* (pp. 221-237). Malden, MA: Blackwell.
- McKoon, G., Gerrig, R.J. & Greene, S.B. (1996). Pronoun resolution without pronouns: Some consequences of memory based text processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 919-932.
- McKoon, G., & Ratcliff, R. (1980). The comprehension processes and memory structures involved in anaphoric reference. *Journal of Verbal Learning and Verbal Behavior*, 19(6), 668-682.
- McKoon, G., & Ratcliff, R. (1981). The comprehension processes and memory structures involved in instrumental inference. *Journal of Verbal Learning and Verbal Behavior*, 20, 671-682.
- McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12(1), 82-91.
- McKoon, G. & Ratcliff, R. (1988). Contextually relevant aspects of meaning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 331-343.
- McKoon, G., & Ratcliff, R. (1989a). Semantic associations and elaborative inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 326-338.
- McKoon, G., & Ratcliff, R. (1989b). Assessing the occurrence of elaborative inference with recognition: Compatibility checking vs. compound cue theory. *Journal of Memory and Language*, 28, 547-563.
- McKoon, G., & Ratcliff, R. (1990a). Dimensions of inference. In A. Graesser and G. Bower (Eds.), *The Psychology of Learning and Motivation: Vol. 25. Inferences and text comprehension* (pp. 313-328). New York: Academic Press.
- McKoon, G. & Ratcliff, R. (1990b). Textual inferences: Models and measures. In D.A. Balota, G.B. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension Processes in Reading* (pp. 403-421). Hillsdale, NJ: Erlbaum.
- McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, 99(3), 440-466.

- McKoon G, Ratcliff R. (1995). The minimalist hypothesis: directions for research. In CA Weaver, S Mannes, CR Fletcher (Eds.). *Discourse Comprehension: Essays in Honor of Walter Kintsch* (pp. 97-116). Hillsdale, NJ: Erlbaum.
- McNamara, D. S., & Kintsch, W. (1996). Learning from texts: Effects of prior knowledge and text coherence. *Discourse processes*, 22, 247-288.
- Miller, A. G., Ashton, W., & Mishal, M. (1990). Beliefs concerning the features of constrained behavior: A basis for the fundamental attribution error. *Journal of Personality and Social Psychology*, 59(4), 635-650.
- Miller, A. G., & Rorer, L. G. (1982). Toward an understanding of the fundamental attribution error: Essay diagnosticity in the attitude attribution paradigm. *Journal of Research in Personality*, 16(1), 41-59.
- Miller, D. T., & Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction. *Psychological Bulletin*, 82(2), 213-225.
- Miller, D. T., & Turnbull, W. (1986). Expectancies and interpersonal processes. *Annual Review of Psychology*, 37(1), 233-256.
- Miller, F. D., Smith, E. R., & Uleman, J. (1981). Measurement and interpretation of situational and dispositional attributions. *Journal of Experimental Social Psychology*, 17(1), 80-95.
- Miller, M. B., & Gazzaniga, M. S. (1998). Creating false memories for visual scenes. *Neuropsychologia*, 36(6), 513-520.
- Millis, K. K., & Just, M. A. (1994). The influence of connectives on sentence comprehension. *Journal of Memory and Language*, 33(1), 128-147.
- Minar, E. H. (1998). Wittgenstein on the metaphysics of the self: the dialectic of solipsism in philosophical investigations. *Pacific Philosophical Quarterly*, 79(4), 329-354.
- Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, 132(2), 297-326.
- Morris, M. W., Smith, E. E., & Turner, K. (1998). Parsimony in intuitive explanations for behavior: Reconciling the discounting principle and preference for conjunctive explanations. *Basic and Applied Social Psychology*, 20(1), 71-85.
- Morrow, D.G., Bower, G.H., & Greenspan, S.L. (1989). Updating situation models during narrative comprehension. *Journal of Memory and Language*, 28, 292-312.

- Morrow, D. G., Greenspan, S. L., & Bower, G. H. (1987). Accessibility and situation models in narrative comprehension. *Journal of Memory and Language*, 26, 165-187.
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76(2), 165-178.
- Moskowitz, G. B. (1993a). Person organization with a memory set: are spontaneous trait inferences personality characterizations or behaviour labels? *European Journal of Personality*, 7(3), 195-208.
- Moskowitz, G. B. (1993b). Individual differences in social categorization: The influence of personal need for structure on spontaneous trait inferences. *Journal of Personality and Social Psychology*, 65, 132-132.
- Moskowitz, G. B., & Grant, H. (Eds.). (2009). *The psychology of goals*. New York: Guilford Press.
- Moskowitz, G. B., & Roman, R. J. (1992). Spontaneous Trait Inferences as Self-Generated Primes: Implications for Conscious Social Judgment. *Journal of Personality & Social Psychology*, 62(5), 728-738.
- Moskowitz, G. B., Salomon, A. R., & Taylor, C. M. (2000). Preconsciously controlling stereotyping: Implicitly activated egalitarian goals prevent the activation of stereotypes. *Social Cognition*, 18, 151-177.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89(6), 316-338.
- Murphy, F. C., Wilde, G., Ogden, N., Barnard, P. J., & Calder, A. J. (2009). Assessing the automaticity of moral processing: Efficient coding of moral information during narrative comprehension. *The Quarterly Journal of Experimental Psychology*, 62(1), 41-49.
- Myers, J. L., & O'Brien, E. J. (1998). Accessing the discourse representation during reading. *Discourse Processes*, 26, 131-158.
- Myers, J. L., O'Brien, E. J., Albrecht, J. E., & Mason, R. A. (1994). Maintaining global coherence during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 876-876.
- Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language*, 26(4), 453-465.
- Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, 86(3), 214-255.

- Neale, J. M., & Friend, R. M. (1972). Attributional Determinants of Reactions to Performance in Academic Situations. *Perceptual and Motor Skills*, 34(1), 35-40.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, 106(3), 226-254.
- Nelson, K., & Gruendel, J. (1979). At morning it's lunchtime: A scriptal view of children's dialogues. *Discourse Processes*, 2, 73-94.
- Neuschatz, J. S., Lampinen, J. M., Preston, E. L., Hawkins, E. R., & Toglia, M. P. (2002). The effect of memory schemata on memory and the phenomenological experience of naturalistic situations. *Applied Cognitive Psychology*, 16(6), 687-708.
- Newman, L. S. (1991). Why are traits inferred spontaneously? A developmental approach. *Social Cognition*, 9(3), 221-253.
- Newman, L. S. (1993). How individualists interpret behavior: Idiocentrism and spontaneous trait inference. *Social Cognition*, 11(2), 243-269.
- Newman, L. S., & Uleman, J. S. (1990). Assimilation and contrast effects in spontaneous trait inference. *Personality and Social Psychology Bulletin*, 16(2), 224-240.
- Newman, L.S., & Uleman, J.S. (1993). When are you what you did? Behavior identification and dispositional inference in person memory, attribution, and social judgment. *Personality and Social Psychology Bulletin*, 19, 513-525.
- Nickerson, R. S. (1984). Retrieval inhibition from par-set cuing: a persisting enigma in memory research. *Memory & Cognition*, 12(6), 531-552.
- Nisbett, R. E., Caputo, C., Legant, P., & Marecek, J. (1973). Behavior as seen by the actor and as seen by the observer. *Journal of Personality and Social Psychology*, 27(2), 154-164.
- Nisbett, R. E. Wilson. T.D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231-259.
- Norvig, P., & Russell, S. J. (2003). *Artificial intelligence: a modern approach*: Prentice Hall.
- O'Brien, E., & Albrecht, J. (1992). Comprehension strategies in the development of a mental model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 777-784.

- O'Brien, E. J., Duffy, S. A., & Myers, J. L. (1986). Anaphoric inference during reading. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *12*, 346–352.
- O'Brien, E. J., & Myers, J. L. (1999). Text comprehension: A view from the bottom up. In Goldman, S. R., Graesser, A. C., and Trabasso, T. (Eds.), *Narrative Comprehension, Causality, and Coherence: Essays in Honor of Tom Trabasso*, (pp. 35–55). NJ: Erlbaum.
- O'Brien, E. J., Rizzella, M. L., Albrecht, J. E., & Halleran, J. G. (1998). Updating a situation model: A memory-based text processing view. *Journal of Experimental Psychology: Learning Memory and Cognition*, *24*(5), 1200-1210.
- Ogden, W.C., Martin, D.W., & Paap, K.R. (1980) Processing demands of encoding: What does secondary task performance reflect? *Journal of Experimental Psychology*, *6*, 355-367.
- Ohtsuka, K., & Brewer, W. F. (1992). Discourse organization in the comprehension of temporal order in narrative texts. *Discourse Processes*, *15*, 317-336.
- Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology*, *32*(4), 605-616.
- Paap, K. R., & Ogden, W. C. (1981). Letter encoding is an obligatory but capacity-demanding operation. *Journal of Experimental Psychology: Human Perception and Performance*, *7*(3), 518-527.
- Park, B. (1989). Trait attributes as on-line organizers in person impressions. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 39–59). Hillsdale, NJ: Erlbaum.
- Perfetti, C., & McCutchen, D. (1987). Schooled language competence: Linguistic abilities in reading and writing. In S. Rosenberg (Ed.), *Advances in applied psycholinguistics: Vol. 2. Reading, writing, and language learning* (pp. 105-141). New York: Cambridge University press.
- Peterson, C., & Seligman, M. E. P. (1984). Causal explanations as a risk factor for depression: Theory and evidence. *Psychological Review*, *91*(3), 347-374.
- Piaget, J. (1926). *The language and thought of the child*. New York: Harcourt Brace.
- Pittman, T. S., & Pittman, N. L. (1980). Deprivation of control and the attribution process. *Journal of Personality and Social Psychology*, *39*(3), 377-389.
- Posner, M.I. and Snyder, C.R.R. (1975). Attention and cognitive control. In Solso, R. (Eds.), *Information Processing and Cognition: The Loyola symposium* (pp 55-86). Hillsdale, NJ: Erlbaum.

- Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language*, 27(4), 399-415.
- Premack, D., & Premack, A. J. (1997). Infants attribute value to the goal-directed actions of self-propelled objects. *Journal of Cognitive Neuroscience*, 9(6), 848-856.
- Proust, M. (1930). *The sweet cheat gone (The Fugitive)*. Retrieved July, 10, 2009, from <http://feedbooks.com/book/1452>.
- Purdue, C. W., & Gurtman, M. B. (1990). Evidence for automatic ageism. *Journal of Experimental Social Psychology*, 26, 199-216.
- Pyszczynski, T. A., & Greenberg, J. (1981). Role of disconfirmed expectancies in the instigation of attributional processing. *Journal of Personality and Social Psychology*, 40(1), 31-38.
- Quattrone, G. A. (1982). Overattribution and unit formation: When behavior engulfs the person. *Journal of Personality and Social Psychology*, 42(4), 593-607.
- Quine, W. V. (1970). *Philosophy of Logic*. Englewood Cliffs, NJ: Prentice-Hall.
- Rajaram, S., & Roediger, H. L., III (1997). Remembering and knowing as states of consciousness during retrieval. In J. D. Cohen & J.W. Scholler (Eds.), *Scientific approaches to consciousness* (pp. 213-240). Mahwah, NJ: Erlbaum.
- Ram, A., & Hunter, L. (1992). The use of explicit goals for knowledge to guide inference and learning. *Applied Intelligence*, 2(1), 47-73.
- Ramos, T. (2006). Quando uma ação permite diferentes interpretações: Pré-teste de comportamentos ambíguos: Pré-teste de comportamentos ambíguos. *Laboratório de Psicologia*, 4 (1), 45-63.
- Ratcliff, R., & McKoon, G. (1978). Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behaviour*, 17(4), 403-417.
- Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, 104, 319-343.
- Ratcliff, R., & McKoon, G. (2008). Passive parallel automatic minimalist processing. In C. Engel & W. Singer, (Eds.), *Better than Conscious? Decision Making, the Human Mind, and Implications for Institutions*. Cambridge, MA: MIT Press.
- Read, S. J. (1987). Constructing causal scenarios: A knowledge structure approach to causal reasoning. *Journal of Personality and Social Psychology*, 52(2), 288-302.

- Read, S. J., Jones, D. K., & Miller, L. C. (1990). Traits as goal-based categories: The importance of goals in the coherence of dispositional categories. *Journal of Personality and Social Psychology*, 58(6), 1048-1061.
- Read, S. J., & Miller, L. C. (1989). Inter-personalism: Toward a goal-based theory of persons in relationships. In L. Pervin (Ed.), *Goal concepts in personality and social psychology* (pp. 413-472). Hillsdale, NJ: Erlbaum.
- Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy": Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin*, 19(5), 526-540.
- Read, S. J., & Miller, L. C. (2005). Explanatory coherence and goal-based knowledge structures in making dispositional inferences. In B. Malle & S. Hodges (Eds.), *Other Minds: How humans bridge the divide between self and others* (pp.124-142). New York: Guilford Press.
- Reed, A. V. (1973). Speed-accuracy trade-off in recognition memory. *Science*, 181, 574-576.
- Reeder, G. D. (2009). Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychological Inquiry*, 20, 1-18.
- Reeder, G. D., Kumar, S., Hesson-McInnis, M. S., & Trafimow, D. (2002). Inferences about the morality of an aggressor: The role of perceived motive. *Journal of Personality and Social Psychology*, 83(4), 789-803.
- Reeder, G. D., & Spores, J. M. (1983). The attribution of morality. *Journal of Personality and Social Psychology*, 44, 736-745.
- Reeder, G. D., Vonk, R., Ronk, M. J., Ham, J., & Lawrence, M. (2004). Dispositional attribution: Multiple inferences about motive-related traits. *Journal of Personality and Social Psychology*, 86(4), 530-544.
- Reisenzein, R., & Rudolph, U. (2008). The discovery of commonsense psychology. *Social Psychology*, 39(3), 125-133.
- Richardson-Klavehn, A., & Bjork, R. A. (1988). Measures of memory. *Annual Review of Psychology*, 39(1), 475-543.
- Rickheit, G, Schnotz, W & Strohner, H (1985). The concept of inference in discourse comprehension, In Rickheit, G & Strohner, H (Eds.), *Inferences in text processing* (pp.183-204). Amsterdam: North-Holland.
- Rickheit, G., & Strohner, H. (1985) (Eds.). *Inferences in Text Processing*. Amsterdam: North-Holland.

- Rieger, C. (1975). Conceptual memory. In R. C. Schank (Ed.), *Conceptual information processing* (pp.157-288). Amsterdam: North-Holland.
- Rinck, M., & Weber, U. (2003). Who when where: An experimental test of the event-indexing model. *Memory & Cognition*, *31*(8), 1284-1292.
- Rinck, M., Williams, E, Bower, G. H., & Becket, E. S. (1996). Spatial situation models and narrative understanding: Some generalizations and extensions. *Discourse Processes*, *21*, 23-55.
- Robins, R. W., Spranca, M. D., & Mendelsohn, G. A. (1996). The actor-observer effect revisited: Effects of individual differences and repeated social interactions on actor and observer attributions. *Journal of Personality and Social Psychology*, *71*(2), 375-389.
- Roediger, H. L. (1974). Inhibiting effects of recall. *Memory & Cognition*, *2*(2), 261-269.
- Roediger, H. L. (1990). Implicit memory. *American Psychologist*, *45*, 1043-1056.
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 803-814.
- Roenker, D. L., Thompson, C. P., & Brown, S. C. (1971). Comparison of measures for the estimation of clustering in free recall. *Psychological Bulletin*, *76*(1), 45-48.
- Roese, N. J. (1997). Counterfactual thinking. *Psychological Bulletin*, *121*, 133-148.
- Roese, N. J., & Olson, J. M. (2007). Counterfactual thinking: The intersection of affect and function. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 29, pp. 1-59). San Diego, CA: Academic Press.
- Roese N. J., Sanna L. J., & Galinsky, A. D. (2005). The mechanics of imagination: Automaticity and control in counterfactual thinking. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The New Unconscious* (pp.138-170). Oxford University Press; New York.
- Ross, L. (1977). The intuitive psychologist and his shortcomings. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173-220). San Diego, CA: Academic Press.
- Ross, L. D., Amabile, T. M., & Steinmetz, J. L. (1977). Social roles, social control, and biases in social-perception processes. *Journal of Personality and Social Psychology*, *35*, 485-494.
- Ross, M. & Sicoly, F. (1979). Egocentric biases in availability and attribution. *Journal of Personality and Social Psychology*, *37*, 322-336.

- Ruble, T. L. (1973). Effects of actor and observer roles on attribution of causality in situations of success and failure. *Journal of Social Psychology, 90*, 41-44.
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review, 89*(1), 60-94.
- Russell, B. (1914). *Our knowledge of the external world*. Chicago and London: Open Court Publishing.
- Salomon, G. (1984). Television is "easy" and print is "tough": The differential investment of mental effort in learning as a function of perceptions and attributions. *Journal of Educational Psychology, 76*(4), 647-658.
- Sanford, A. J., & Garrod, S. C. (1981). *Understanding written language*. New York: Wiley.
- Sanford, A. J., & Graesser, A. C. (2006). Shallow processing and underspecification. *Discourse Processes, 42*(2), 99-108.
- Sanna, L. J., & Turley, K. J. (1996). Antecedents to spontaneous counterfactual thinking: Effects of expectancy violation and outcome valence. *Personality and Social Psychology Bulletin, 22*(9), 906-919.
- Sanna, L. J., & Turley-Ames, K. J. (2000). Counterfactual intensity. *European Journal of Social Psychology, 30*(2), 273-296.
- Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory & Cognition, 13*, 501-518.
- Schank, R. C. (1975). Using Knowledge to Understand. In R. Schank and B. L. Nash-Webber (Eds.), *Theoretical Issues in Natural Language Processing* (pp.131-135). Cambridge, MA.
- Schank, R. C. (1982). *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge: Cambridge University Press.
- Schank, R. C. (1987). What is AI, anyway? *AI Magazine, 8*(4), 59-65.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. NJ: Lawrence Erlbaum Associates.
- Schank, R. C., & Abelson, R. P. (1995). Knowledge and memory: The real story. In R. S. Wyer Jr. (Ed.), *Advances in social cognition* (Vol. 8, pp. 1-86). Hillsdale, NJ: Erlbaum.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-Prime User's Guide*. Pittsburgh: Psychology Software Tools Inc.

- Schwarz, N. (2001). Feelings as information: Implications for affective influences on information processing. In L. L. Martin, & G. L. Clore (Eds.), *Theories of mood and cognition: A user's handbook* (pp. 159-176). Mahwah, NJ: Erlbaum.
- Schwarz, N., & Bless, H. (1991). Happy and mindless, but sad and smart? The impact of affective states on analytic reasoning. In J. Forgas (Ed.), *Emotion and social judgments* (pp. 55-71). Oxford: Pergamon.
- Schwarz, N. & Bless, H. (1992). Constructing reality and its alternatives: An inclusion/exclusion model of assimilation and contrast effects in social judgment. In L. Martin & A. Tesser (Eds.), *The construction of social judgment* (pp. 217-245). Hillsdale, NJ: Lawrence Erlbaum.
- Searle, J. R. (1979). *Expression and meaning: Studies in the theory of speech acts*. Cambridge: Cambridge University Press.
- Sedikides, C. (1990). Effects of fortuitously activated constructs versus activated communication goals on person impressions. *Journal of Personality and Social Psychology*, 58(3), 397-408.
- Seligman, M. E., Abramson, L. Y., Semmel, A., & Von Baeyer, C. (1979). Depressive attributional style. *Journal of Abnormal Psychology*, 88(3), 242-247.
- Shallice, T. (1988) *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Sharkey, N. E. (1986). A model of knowledge-based expectations in text comprehension. In J.A. Galambos, R.P. Abelson, & J.B. Black (Eds.), *Knowledge structures* (pp. 49-70). Hillsdale, NJ: Erlbaum.
- Sharkey, N. E., & Mitchell, D. C. (1985). Word recognition in a functional context: The use of scripts in reading. *Journal of Memory and Language*, 24(2), 253-270.
- Shiffrin, R. M. & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- Shimamura, A. P. (1986). Priming effects in amnesia: Evidence for a dissociable memory function. *Quarterly Journal of Experimental Psychology*, 38A, 619-644.
- Simon, H. (1957). *Models of man: Social and rational*. New York: Wiley.
- Singer, M. (1979). Processes of inference during sentence encoding. *Memory & Cognition*, 7(3), 192-200.
- Singer, M. (1993). Global inferences of text situations. *Discourse Processes*, 16, 161-168.

- Singer, M., & Ferreira, F. (1983). Inferring consequences in story comprehension. *Journal of Verbal Learning and Verbal Behavior*, 22(4), 437-448.
- Singer, M., Halldorson, M., Lear, J. c., & Andrusiak, P. (1992). Validation of causal bridging inferences in discourse understanding. *Journal of Memory and Language*, 4, 507-524.
- Singer, M. & Leon, J. (2007) Psychological studies of higher language processes: Behavioral and empirical approaches. In F. Schmalhofer, & C. A. Pefetti (Eds.), *Higher level language processes in the brain: Inference and comprehension processes*. Mahwah, NJ: Erlbaum.
- Skowronski, J. J., Carlston, D. E., Mae, L., & Crawford, M. T. (1998). Spontaneous trait transference: Communicators take on the qualities they describe in others. *Journal of Personality and Social Psychology*, 74, 837-848.
- Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4(2), 108-131.
- Smith, E. R., & Miller, F. D. (1983). Mediation among attributional inferences and comprehension processes: Initial findings and a general method. *Journal of Personality and Social Psychology*, 44(3), 492-505.
- Snyder, M. L., & Frankel, A. (1976). Observer bias: A stringent test of behavior engulfing the field. *Journal of Personality and Social Psychology*, 34(5), 857-864.
- Snyder, M. L., & Jones, E. E. (1974). Attitude attribution when behavior is constrained. *Journal of Experimental Social Psychology*, 10(6), 585-600.
- Speer, N. K., & Zacks, J. M. (2005). Temporal changes as event boundaries: Processing and memory consequences of narrative time shifts. *Journal of Memory and Language*, 53(1), 125-140.
- Spencer, S. J., Fein, S., Wolfe, C. T., Fong, C., & Dunn, M. (1998). Automatic activation of stereotypes: The role of self-image threat. *Personality and Social Psychology Bulletin*, 24, 1139-1152.
- Sperry, R. W. (1993). The Impact and Promise of the Cognitive Revolution. *American Psychologist*, 48(8), 878-885.
- Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, 37(10), 1660-1672.
- Srull, T. K., & Wyer, R. S. (1989). Person memory and judgment. *Psychological Review*, 96(1), 58-83.

- Stapel, D. A., Koomen, W., & Van der Pligt, J. (1996). The referents of trait inferences: The impact of trait concepts versus actor-trait links on subsequent judgments. *Journal of Personality and Social Psychology, 70*, 437-450.
- Stapel, D. A., Koomen, W., & Van Der Pligt, J. (1997). Categories of category accessibility: The impact of trait concept versus exemplar priming on person judgments. *Journal of Experimental Social Psychology, 33*(1), 47-76.
- Stern, L. D., Marrs, S., Millar, M. G., & Cole, E. (1984). Processing time and the recall of inconsistent and consistent behaviors of individuals and groups. *Journal of Personality and Social Psychology, 47*(2), 253-262.
- Storms, M. D. (1973). Videotape and the attribution process: Reversing actors' and observers' points of view. *Journal of Personality and Social Psychology, 27*(2), 165-175.
- Strack, F., & Hannover, B. (1996). Awareness of influence as a precondition for implementing correctional goals. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 579-596). New York: Guilford.
- Stratton, P. (2003). Causal attributions during therapy I: Responsibility and blame. *Journal of Family Therapy, 25*(2), 136-160.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General, 18*, 643-662.
- Stroessner, S. (1989). Behavior and ratings pool (Compilation). Unpublished manuscript
- Suh, S., & Trabasso, T. (1993). Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language, 24*, 595-611.
- Swann, W. B., Stephenson, B., & Pittman, T. S. (1981). Curiosity and control: On the determinants of the search for social knowledge. *Journal of Personality and Social Psychology, 40*(4), 635-642.
- Swinney, D. A., & Osterhout, L. (1990). Inference generation during auditory language comprehension. In A. C. Graesser, & G. H. Bower (Ed.), *Inferences and Text Comprehension* (pp. 17-33). San Diego, CA : Academic Press.
- Tapiero, S. (2007). *Situation models and levels of coherence: toward a definition of comprehension*. New York: Lawrence Erlbaum Associates
- Taylor, S. E., & Fiske, S. T. (1975). Point of view and perceptions of causality. *Journal of Personality and Social Psychology, 32*(3), 439-445.

- Taylor, S. E., & Koivumaki, J. H. (1976). The perception of self and others: acquaintanceship, affect, and actor-observer differences. *Journal of Personality and Social Psychology*, 33(4), 403-408.
- Tetlock, P. E. (1985). Accountability: A social check on the fundamental attribution error. *Social Psychology Quarterly*, 48, 227-236.
- Thomson, D. M., & Tulving, E. (1970). Associative encoding and retrieval weak and strong cues. *Journal of Experimental Psychology*, 86(2), 255-262.
- Thompson, E. P., Roman, R. J., Moskowitz, G. B., Chaiken, S., & Bargh, J. A. (1994). Systematic processing and the debasing of covert primacy effects in impression formation: Unshackling the motivated perceiver from the constraints of accessibility. *Journal of Personality and Social Psychology*, 66, 474-489.
- Titchener, E. B. (1923). Relearning after forty-six years. *American Journal of Psychology*, 34, 468-469.
- Todorov, A., & Uleman, J. S. (2002). Spontaneous trait inferences are bound to actors' faces: Evidence from a false recognition paradigm. *Journal of Personality & Social Psychology*, 83(5), 1051-1065.
- Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology*, 39(6), 549-562.
- Todorov, A., & Uleman, J. S. (2004). The person reference process in spontaneous trait inferences. *Journal of Personality and Social Psychology*, 87(4), 482-493.
- Trabasso, T., & Magliano, J. E. (1996). Conscious understanding during comprehension. *Discourse Processes*, 21, 255-287.
- Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, 24(5), 595-611.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, 93, 239-257.
- Trope, Y., Cohen, O., & Maoz, Y. (1988). The perceptual and inferential effects of situational inducements on dispositional attribution. *Journal of Personality and Social Psychology*, 55(2), 165-177.
- Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology*, 79(3), 344-354.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology*, 25, 1-12.

- Tulving, E., & Osler, S. (1968). Effectiveness of retrieval cues in memory for words. *Journal of Experimental Psychology*, 77(4), 593-601.
- Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80(5), 352-373.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
- Tzelgov, J., Henik, A., & Berger, J. (1992). Controlling Stroop effects by manipulating expectations for color words. *Memory & Cognition*, 20(6), 727-735.
- Tzelgov, J., Henik, A., & Leiser, D. (1990). Controlling Stroop interference: Evidence from a bilingual task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(5), 760-771.
- Uleman, J. S. (1999). Spontaneous versus intentional inferences in impression formation. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 141-160). New York: Guilford.
- Uleman, J. S., & Blader, S. (2001). A process dissociation analysis of using spontaneous trait inferences. *Unpublished data*.
- Uleman, J. S., Blader, S., & Todorov, A. (2005). Implicit impressions. In R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The New Unconscious* (pp.362-392). New York: Oxford University Press.
- Uleman, J. S., Hon, A., Roman, R. J., & Moskowitz, G. B. (1996). On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin*, 22(4), 377-394.
- Uleman, J. S., & Moskowitz, G. B. (1994). Unintended Effects of Goals on Unintended Inferences. *Journal of Personality & Social Psychology*, 66(3), 490-501.
- Uleman, J. S., Moskowitz, G. B., Roman, R. J., & Rhee, E. (1993). Tacit, manifest, and intentional reference: How spontaneous trait inferences refer to persons. *Social Cognition*, 11(3), 321-351.
- Uleman, J. S., Newman, L. S., & Moskowitz, G. B. (1996). People as flexible interpreters: Evidence and issues from spontaneous trait inference. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 28, pp. 211-279). San Diego, CA: Academic Press.
- Uleman, J. S., Newman, L. S., & Winter, L. (1992). Can personality traits be inferred automatically? Spontaneous inferences require cognitive capacity at encoding. *Consciousness and Cognition*, 1, 77-90.

- Uleman, J. S., Saribay, S. A., & Gonzalez, C. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology*, *59*, 329-360.
- Uleman, J. S., Winborne, W. C., Winter, L., & Shechter, D. (1986). Personality differences in spontaneous trait inferences at encoding. *Journal of Personality and Social Psychology*, *51*, 396-404.
- Uller, C. (2004). Disposition to recognize goals in infant chimpanzees. *Animal Cognition*, *7*(3), 154-161.
- Uller, C., & Nichols, S. (2000). Goal attribution in chimpanzees. *Cognition*, *76*(2), 27-34.
- Valenza, E., Simion, F., Cassia, V. M., & Umiltà, C. (1996). Face preference at birth. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(4), 892-902.
- Van den Broek, P., Ridsen, K., Fletcher, C.R., & Thurlow, R. (1996). A 'landscape' view of reading: fluctuating patterns of activation and the construction of a stable memory representation. In B.K. Britton & A.C. Graesser (Eds.), *Models of understanding text* (pp. 165-187). Hillsdale, NJ: Erlbaum.
- Van den Broek P, Ridsen K, & Husebye-Hartmann E. (1995). The role of readers standards for coherence in the generation of inferences during reading. In R.F. Lorch, & O'Brien (Eds.), *Sources of Coherence in Reading* (pp.353-374). Hillsdale, NJ: Erlbaum
- Van der Cruyssen, L., Van Duynslaeger, M., Cortoos, A., & Van Overwalle, F. (in press). ERP time course and brain areas of spontaneous and intentional goal inferences. *Social Neuroscience*.
- Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.
- Van Overwalle, F., Drenth, T., & Marsman, G. (1999). Spontaneous trait inferences: Are they linked to the actor or to the action? *Personality and Social Psychology Bulletin*, *25*(4), 450-462.
- Van Overwalle, F., & Timmermans, B. (2005). Discounting and the role of the relation between causes. *European Journal of Social Psychology*, *35*(2), 199-223.
- Van Overwalle, F., & Van Rooy, D. (2001). How one cause discounts or augments another: A connectionist account of causal competition. *Personality and Social Psychology Bulletin*, *27*(12), 1613.
- Vonk, W. (1985). The immediacy of inferences in the understanding of pronouns. In G. Rickheit & H. Stronher (Eds.), *Inferences in text processing* (pp.205-218). New York: Elsevier Science Publishers.

- Walker, C. H., & Yekovich, F. R. (1984). Script-based inferences: Effects of text and knowledge variables on recognition memory. *Journal of Verbal Learning and Verbal Behavior*, 23(3), 357-370.
- Warrington, E. K., & Weiskrantz, L. (1968). New method of testing long-term retention with special reference to amnesic patients. *Nature*, 217, 972-974.
- Warrington, E. K., & Weiskrantz, L. (1970). Amnesic syndrome: Consolidation or retrieval? *Nature*, 228, 629-630.
- Warrington, E. K., & Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychologia*, 12, 419-428.
- Watson, D. (1982). The actor and the observer: How are their perceptions of causality divergent. *Psychological Bulletin*, 92(3), 682-700.
- Webster, D. M. (1993). Motivated augmentation and reduction of the overattribution bias. *Journal of Personality and Social Psychology*, 65, 261-261.
- Webster, D. M., & Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67, 1049-1049.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, 101(1), 34-52.
- Wegner, D. M., & Bargh, J. (1998). Control and automaticity in social life. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of Social Psychology* (4th ed., pp. 446-496). New York: McGraw-Hill.
- Wegner, D. M., & Erber, R. (1992). The hyperaccessibility of suppressed thoughts. *Journal of Personality and Social Psychology*, 63, 903-912.
- Weisz, C., & Jones, E. E. (1993). Expectancy disconfirmation and dispositional inference: Latent strength of target-based and category-based expectancies. *Personality and Social Psychology Bulletin*, 19(5), 563-573.
- Whitney, P., Davis, P. A., & Waring, D. A. (1994). Task effects on trait inference: Distinguishing categorization from characterization. *Social Cognition*, 12(1), 19-35.
- Whitney, P., Waring, D. A., & Zingmark, B. (1992). Task effects on the spontaneous activation of trait concepts. *Social Cognition*, 10(4), 377-396.
- Whitney, P., & Williams-Whitney, D. (1990). Toward a contextualist view of elaborative inferences. In A. Graesser, & G. H. Bower (Eds.), *The Psychology of Learning and Motivation: Vol.25. Inference and Text Comprehension* (pp. 279-293). New York: Academic Press.

- Wickens, C. D. (1980). The structure of attentional resources. In R. Nickerson (Ed.), *Attention and performance VIII*. (pp. 239-257). Hillsdale, NJ: Erlbaum.
- Wickens, C. D. (2008). Multiple resources and mental workload. *Human Factors*, *50*(3), 449-455.
- Wigboldus, D. I. H. J., Dijksterhuis, A., & Van Knippenberg, A. (2003). When Stereotypes Get in the Way: Stereotypes Obstruct Stereotype-Inconsistent Trait Inferences. *Journal of Personality & Social Psychology*, *84*(3), 470-484.
- Wigboldus, D. I. H. J., Sherman, J. W., Franzese, H. L., & van Knippenberg, A. (2004). Capacity and Comprehension: Spontaneous stereotyping under cognitive load. *Social Cognition*, *22*(3), 292-309.
- Wilson, T. D., & Nisbett, R. E. (1978). The accuracy of verbal reports about the effects of stimuli on evaluations and behavior. *Social Psychology*, *41*(2), 118-131.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology*, *47*(2), 237-252.
- Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology*, *49*(4), 904-917.
- Wittgenstein, L. (1953) *Philosophical Investigations*. Translated by GEM Anscombe. Oxford: Blackwell.
- Wong, P. T. P., & Weiner, B. (1981). When people ask “why” questions, and the heuristics of attributional search. *Journal of Personality and Social Psychology*, *40*(4), 650-663.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1-34.
- Woolfolk, R. L., Doris, J. M., & Darley, J. M. (2006). Identification, situational constraint, and social cognition: Studies in the attribution of moral responsibility. *Cognition*, *100*(2), 283-301.
- Wright, E. F., & Wells, G. L. (1988). Is the attitude-attribution paradigm suitable for investigating the dispositional bias? *Personality and Social Psychology Bulletin*, *14*(1), 183-190.
- Wyer, R. S. (2004). *Social comprehension and judgment: The role of situation models, narratives, and implicit theories*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Wyer, R. S., Jr., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.

- Wyer, R. S., & Srull, T. K. (1994). *Handbook of Social Cognition: Applications*. NJ: Lawrence Erlbaum Associates.
- Yarkin-Levin, K. (1983). Anticipated interaction, attribution, and social interaction. *Social Psychology Quarterly*, 46(4), 302-311.
- Zbrodoff, N. J., & Logan, G. D. (1986). On the autonomy of mental processes: A case study of arithmetic. *Journal of Experimental Psychology: General*, 115(2), 118-130.
- Zelli, A., Cervone, D., & Huesmann, L. R. (1996). Behavioral experience and social inference: individual differences in aggressive experience and spontaneous versus deliberate trait inference. *Social Cognition*, 14(2), 165-190.
- Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science*, 6, 292-297.
- Zwaan, R. A., Magliano, J. P., & Graesser, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 386-397.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162-185.
- Zwaan, R.A., Radvansky, G.A., Hilliard, A.E., & Curiel, J. M. (1998). Constructing multidimensional situation models during reading. *Scientific Studies of Reading*, 2, 199-220.

APPENDICES

APPENDIX A

Stimuli Material and Instructions of Experiment 1

Stimuli Material used for the Experiment 1

Ambiguous Behaviors and Implied Traits

Ambiguous Behaviors	Traits	
His contacts with people are rather limited (Higgins, Rholes, & Jones, 1977)	Shy	Independent
Only rarely did he change his mind even when it it's better if he had (Higgins, Rholes, & Jones, 1977)	Persistent	Stubborn
In order to try to improve his life he avoids loaning money to anyone (adapted from Sedikides, 1990)	Miser	Thrifty
In parties his humour is quick to address the faults that people have or the mistakes that they have made (Sedikides, 1990)	Funny	Rude
Started attempting to keep up to date with cultural knowledge (Sedikides, 1990)	Cultivated	Effortfull
Openly smiled to a colleague with whom he argued the day before (Ramos, 2006)	False	Kind
Didn't thank to the person who gave him a present (Ramos, 2006)	Distracted	Unfriendly

Ambiguous Sentences and Category Labels

Match (Mismatch) (Neutral) Conditions	Trait
The businessman (librarian) (person) contacts with people are rather limited	Independent
The librarian (businessman) (person) contacts with people are rather limited	Shy
Only rarely did the old men (manager) (person) change his mind even when it's better if he had	Stubborn
Only rarely did the manager (old men) (person) change his mind even when it's better if he had	Persistent
The professor (model) (person) started attempting to keep up to date with cultural knowledge	Sultivated
The model (professor) (person) started attempting to keep up to date with cultural knowledge	Srtificial
In order to try to improve his life the construction worker (accountant laborer) (person) avoids loaning money to anyone	Thrifty
In order to try to improve his life the accountant laborer (construction worker) (person) avoids loaning money to anyone	Miser
The skinhead (child) (person) didn't thank to the person who gave him a present	Unfriendly
The child (skinhead) (person) didn't thank to the person who gave him a present	Distracted
Often in parties the plumber's (actor's) (person's) humour is quick to point out the faults of other people	Rude
Often in parties the actor's (plumber's) (person's) humour is quick to point out the faults of other people	Funny
The politician (infant teacher) (person) openly smiled to a colleague with whom he argued the day before	False
The infant teacher (politician) (person) openly smiled to a colleague with whom he argued the day before	Kind

Instructions Experiment 1

(Similar Instructions were used in Experiment 2, 3, 4, and 6)

Instructions I

Study on Comprehension Speed

This experiment is concerned with the way people read and comprehend sentences.

In a few moments you will see a series of sentences on the computer screen and after each sentence you will be asked to respond to a word written in blue that will appear on the screen. Upon the presentation of the word (written in blue) press the button (I) on the keyboard if that exact word had been presented in the sentence you had just read, and press the button (E) on the keyboard if the word had not been presented in the previous sentence.

Please respond as quickly and as accurately as possible.

Click the “Spacebar” to continue.

Instructions II

Remember, your task is simply to indicate whether the word that appears on the screen was presented in the preceding sentence.

In order to help you make these task as fast as you can, please keep your index fingers on the E (“No”) and I (“Yes”) keys throughout the task. This will help you respond more quickly, as you won’t have to move your hand to make your decision.

If you have any questions please ask the experimenter. Otherwise press the “Spacebar” to continue.

Instructions III

In order to familiarize you with the task we will provide you some practice trials. Your task is to indicate, as fast and accurately as possible, whether the probe word that appears after each sentence was part of that sentence. Press the “Spacebar” to continue.

Instructions IV

(After Practice Trials)

Now that you are familiar with the task, the study will begin. Remember to be fast and accurate and to pay attention to the screen at all times. Keep one index finger on the “yes” (I) key and the other index finger on the “No” (E) key.

When you are ready to begin, press the “Spacebar” with your thumb and the program will start immediately.

Thank you for your participation.

APPENDIX B

Stimuli Material of Experiment 2

Stimuli used for the Experiment 2

Sentences and Probe Words

Dispositional Sentence	Situational Continuation	Trait Probe	Situational Probe
Devolveu o dinheiro do troco em excesso	depois de hesitar algum tempo	Honesto	Remorsos
Pisou os pés do seu par	depois de um longo dia de trabalho	Desastrado	Cansaço
Deu bastante dinheiro para caridade	sabendo que as doações são dedutíveis nos impostos	Generoso	Calculismo
Acertou em todas as respostas de um teste de cultura	pois já tinha feito um teste semelhante	Inteligente	Memória
Chegou três horas atrasado a uma importante reunião.	depois de ter tido um acidente de automóvel	Irresponsável	Desastre
Perguntou de onde vêm as estrelas	sendo esta uma pergunta do trabalho de casa	Curioso	Obrigação

Sentences and Category Labels

Consistent/ Inconsistent Label

O padre/toxicod dependente devolveu o dinheiro do troco em excesso

O velhote/dançarino pisou os pés do seu par

A psicóloga/economista deu bastante dinheiro para caridade

O professor/homem do lixo acertou em todas as respostas de um teste de cultura

O cigano/gestor chegou três horas atrasado a uma importante reunião

O melhor aluno da escola/o aluno mais famoso da escola perguntou de onde vêm as estrelas

APPENDIX C

Stimuli Material of Experiment 3

Stimuli used for the Experiment 3

Sentences – Trait/Situation Order Inverted

Sentences

Consistent/ Inconsistent/ Neutral Label

Depois de hesitar algum tempo o padre/toxicodependente/ele devolveu o dinheiro do troco em excesso

Depois de um dia de trabalho o velhote/dançarino/ ele pisou os pés do seu par de dança

Sabendo que as doações são dedutíveis nos impostos a psicóloga/ o economista/ele deu bastante dinheiro para caridade.

Como já tinha feito um teste semelhante professor/homem do lixo/ele acertou em todas as respostas de um teste de cultura

Depois de ter tido um acidente de carro o cigano/ gestor/ ele chegou três horas atrasado a uma importante reunião

Sendo uma das perguntas do seu trabalho de casa o melhor aluno da escola/ o aluno mais famoso da escola/ ele perguntou de onde vêm as estrelas

APPENDIX D

Stimuli Material of Experiment 4

Stimuli used for the Experiment 4

Sentences and Probe Words

Behavior Description	Trait	Situation
Doesn't feel like moving the table	Lazy	Heavy
Leaves the cinema smiling	Happy	Funny
Tells the waitress the food tastes good	Complementary	Delicious
Eats three plates of French fries	Gluttonous	Delicious
Stays at work till 12 p.m. every night	Diligent	Hard
Drops the pot of potatoes	Clumsy	Hot
Cries while watching the movie	Emotional	Touching
Wins the quiz	Smart	Easy
Doesn't like to go to class	Lazy	Boring
Cannot read the letter	Illiterate	Unclear

Sentences and Social Category Labels

Sentences

Consistent/ Inconsistent/ Neutral

The delinquent/ weightlifter/ person doesn't feel like moving the table

The kid/ depressed person/ person leaves the cinema smiling

The gentleman/ connoisseur/ person tells the waitress the food tastes good

The soccer fan/ dietician/ person eats three plates of French fries

The sales manager/ Mexican/ person stays at work till 12 p.m. every night

The new waitress/ top chef/ person drops the pot of potatoes

The teenage girl/ teenage boy/ person cries while watching the movie

The physician/ high school dropout/ person wins the quiz

The party kid/ the PhD student/ person doesn't like to go to class

The aboriginal/ teacher/ person cannot read the letter

APPENDIX E

Stimuli Material and Instructions of Experiment 5

Stimuli used for the Experiment 5

Material – Trait Version I

Inconsistent Pairs

	Order 1	Order 2
John	Gave as much as he could anonymously to a charity (Friendly)	Insulted a secretary in his office without provocation (Unfriendly)
	Insulted a secretary in his office without provocation (Unfriendly)	Gave as much as he could anonymously to a charity (Friendly)
Peter	Made new friends easily (Extroverted)	Did not maintain eye contact during the conversation (Introverted)
	Did not maintain eye contact during the conversation (Introverted)	Made new friends easily (Extroverted)
Brian	Told his friend he was interested in backpacking across U.S. (Adventurous)	Reluctantly followed the guide and other people into the cave (Scared)
	Reluctantly followed the guide and other people into the cave (Scared)	Told his friend he was interested in backpacking across U.S. (Adventurous)
Daniel	Make sure all the bills were paid at the end of the month (Responsible)	Purposely did not finish his work, so someone else had to (Irresponsible)
	Purposely did not finish his work, so someone else had to (Irresponsible)	Make sure all the bills were paid at the end of the month (Responsible)
Michael	Returned a wallet he found to the owner (Honest)	Cheated in a game of Monopoly (Dishonest)
	Cheated in a game of Monopoly (Dishonest)	Returned a wallet he found to the owner (Honest)
Eric	Watches the evening news regularly to keep up on current events (Curious)	Sleeps most of his class lectures and does not take notes (Uninterested)
	Sleeps most of his class lectures and does not take notes (Uninterested)	Watches the evening news regularly to keep up on current events (Curious)

Consistent Pairs		
	Order 1	Order 2
James	Solved a complicated mathematics problem in his spare time (Intelligent)	Impressed the history lecturer with his questions (Intelligent)
	Impressed the history lecturer with his questions (Intelligent)	Solved a complicated mathematics problem in his spare time (Intelligent)
Robert	Whenever possible he reuses the materials (Ecological)	Whenever possible he reuses the materials (Ecological)
	Doesn't use products that can damage the ozone layer (Ecological)	Doesn't use products that can damage the ozone layer (Ecological)
Donald	Kept working even though the job was boring (Determined)	Runs five miles a day to keep in shape (Determined)
	Runs five miles a day to keep in shape (Determined)	Kept working even though the job was boring (Determined)
Philip	Couldn't keep a secret (Disloyal)	Told a lie about a good friend (Disloyal)
	Told a lie about a good friend (Disloyal)	Couldn't keep a secret (Disloyal)
Anthony	He is reluctant to give his opinion about any issue (Feeble)	Was afraid to tell the waitress about the spoiled meat (Feeble)
	Was afraid to tell the waitress about the spoiled meat (Feeble)	He is not capable of giving his opinion about any issue ((Feeble)
Chris	Accidentally deleted all hi word-processing files on his computer (Absent-minded)	Put on different pairs of shoes and didn't notice (Absent-minded)
	Put his sandals on the wrong feet (Absent-minded)	Accidentally deleted all hi word processing files on his computer (Absent-minded)

Material – Trait Version II

Consistent Pairs		
	Order 1	Order 2
John	Gave as much as he could anonymously to a charity (Friendly)	Offered some directions to some tourists who looked lost (Friendly)
	Offered some directions to some tourists who looked lost (Friendly)	Gave as much as he could anonymously to a charity (Friendly)
Peter	Made new friends easily (Extroverted)	Was the captain of the local softball team (Extroverted)
	Was the captain of the local softball team (Extroverted)	Made new friends easily (Extroverted)
Brian	Told his friend he was interested in backpacking across U.S. (Adventurous)	Tried to go surfing although the waves were enormous (Adventurous)
	Tried to go surfing although the waves were enormous (Adventurous)	Told his friend he was interested in backpacking across U.S. (Adventurous)
Daniel	Purposely did not finish his work, so someone else had to (Irresponsible)	Forgot to take the important medication (Irresponsible)
	Forgot to take the important medication (Irresponsible)	Purposely did not finish his work, so someone else had to (Irresponsible)
Michael	Cheated in a game of Monopoly (Dishonest)	Claimed credit for someone else's idea (Dishonest)
	Claimed credit for someone else's idea (Dishonest)	Cheated in a game of Monopoly (Dishonest)
Eric	Sleeps most of his class lectures and does not take notes (Uninterested)	Never votes in elections (Uninterested)
	Never votes in elections (Uninterested)	Sleeps most of his class lectures and does not take notes (Uninterested)

Consistent Pairs		
	Order 1	Order 2
James	Solved a complicated mathematics problem in his spare time (Intelligent)	Did not realize the group of people were being sarcastic (Stupid)
	Did not realize the group of people were being sarcastic (Stupid)	Solved a complicated mathematics problem in his spare time (Intelligent)
Robert	Whenever possible he reuses the materials (Ecological)	Threw an empty bottle out of his window (Non-Ecological)
	Threw an empty bottle out of his window (Non-Ecological)	Whenever possible he reuses the materials (Ecological)
Donald	During the Easter avoids eating meat (religious)	Said he has doubts about the existence of God (non religious)
	Said he has doubts about the existence of God (non religious)	During the Easter avoids eating meat (religious)
Philip	Couldn't keep a secret (Disloyal)	Supported the team even with their five year losing streak (Loyal)
	Supported the team even with their five year losing streak (Loyal)	Couldn't keep a secret (Disloyal)
Anthony	Kept working even though the job was boring (Determined)	Slept until noon and missed the exam (Lazy)
	Slept until noon and missed the exam (Lazy)	Kept working even though the job was boring (Determined)
Chris	Always forget his wedding anniversary day (Disorganized)	Keep the apartment orderly and clean (Organized)
	Keep the apartment orderly and clean (Organized)	Always forget his wedding anniversary day (Disorganized)

Neutral Pairs (Used in both Trait Versions)

	Order 1	Order 2
William	Watched a late movie on television	Ate a bacon and cheese hamburger for lunch
	Ate a bacon and cheese hamburger for lunch	Watched a late movie on television
Thomas	Took a short nap in the middle of the afternoon	Heated a cup of coffee in the microwave
	Heated a cup of coffee in the microwave	Took a short nap in the middle of the afternoon
Frank	Took his car to the carwash to get it cleaned	Made some orange juice for breakfast
	Made some orange juice for breakfast	Took his car to the carwash to get it cleaned
Gabriel	Took a leisurely walk around the block	Bought stamps at the post office
	Bought stamps at the post office	Took a leisurely walk around the block

Instructions Experiment 5

Instructions I

Memory Experiment

One of the most interesting aspects of human cognition is how we can remember so many and diverse bits of information, often with minimum information available and with little effort. And often those memories are quite accurate. Although this happens all the time in our everyday experiences, we don't fully understand how it happens.

This experiment is concerned with the way in which we remember different information. In a few moments you will be shown a series of pairs of sentences. Each pair of sentences makes reference to a different individual.

Press the "Spacebar" to continue.

Instructions II

Please read the series of sentences carefully, studying each one until the next one appears on the screen. Try to remember each sentence as accurately as you can. At the end of the session, we will ask you some questions concerning that information.

Please read the sentences carefully and try to retain them in your memory so that you can reproduce them later.

Please press the "Spacebar" to start the presentation of the behaviors.

Recall with Traits

Recall I

Earlier in this experiment you read pairs of sentences. Each pair described a different individual. We would like you to recall as many of the second sentences presented in each pair.

We realize that it would be extremely difficult for you to remember all the sentences, but please try to recall as many as you can, and as accurately as possible.

Remember, your task is to recall only the phrasing of the **second sentences** used to describe each individual. Recall as many of the second sentences as you can.

Recall II

In order to help you in this task we will provide some words that may help you recall the sentences. Those words will be presented in the next screen. Try to use the given words as memory cues while recalling the sentences.

Please continue in your attempt to recall as many sentences as you can. But remember, just recall the **SECOND SENTENCES** used to describe each individual, not the sentences presented first.

Recall III

Use these words as memory cues while you are trying to recall the sentences. Note that the number of cues provided does not necessarily correspond to the number of sentences you have to recall.

You will have five minutes for this task.

Recall with First Behaviors

Recall I

Earlier in this experiment you read pairs of sentences. Each pair described a different individual. We would like you to recall as many of the **SECOND SENTENCES** presented in each pair.

We realize that it would be extremely difficult for you to remember all the sentences, but please try to recall as many as you can, and as accurately as possible.

Remember, your task is to recall only the phrasing of the **second sentences** used to describe each individual. Recall as many of the second sentences as you can.

Recall II

In order to help you in this task we will provide cues of first sentences of each pair. Those sentences will be presented in the next screen. Try to use the given sentences as memory cues while recalling the second sentences.

Please continue in your attempt to recall as many sentences as you can. But remember, just recall the **SECOND SENTENCES** used to describe each individual, since we are giving you the sentence presented first.

Recall III

The set of example sentences presented in this task correspond to the first phrasings to describe each individual. Please write the sentences that were originally paired with the initial phrasing. The number of cues provided does not necessarily correspond to the number of sentences you have to recall. You will have five minutes to complete this task.

Filler Task used in Experiment 5

(The same filler task was used in Experiments 7 and 8)

This experiment is concerned with the way we can attend to multiple aspects of social environment. The following task aims to explore the way we perceive and extract information. In the next task, the letters of some English words have been scrambled. Rearrange them so that they form correct words. You will have 4/5 minutes to complete the task. Click the “Spacebar” to start.

Anagram <i>(Randomly Presented)</i>	Correct Answer	Anagram <i>(Randomly presented)</i>	Correct Answer
FEAL	leaf	NPE	pen
ORFK	fork	TYUNIRVEIS	university
LOIN	lion	EFEFCO	coffee
SOBS	boss	CUMSI	music
SOAK	oaks	WOLLYE	yellow
ICAD	acid	IEOOKC	cookie
DOOF	food	ESHUO	house
OTPAOT	potato	YNOUCTR	country
BOALLONO	balloon	NEAOC	ocean
ODLPNHI	dolphin	NALD	land
RIAN	rain	YBO	boy
BLBU	bulb	YICT	city
AERBFAKTS	breakfast	ORWDL	world
RCA	car	CABK	back
CEHRYR	cherry	EIMT	time
EEY	eye	WOT	two
BXO	box	EMHO	home
ETA	tea	DOVIE	video
ANLCDE	candle	EKEW	week
EPI	pie	UCDOENTM	document
KESD	desk	ERIF	fire
ADEI	idea	LGOD	gold
EMAIG	image	TRNAG	grant
TEAS	seat	PTEMBERES	september

APPENDIX F

Stimuli Material of Experiment 6

Stimuli used for the Experiment 6

Pairs of Sentences and Implied Traits

Material Replication I		
Behavior 1	Behavior 2	Traits
Gave as much as he could anonymously to a charity	Gave up his seat on the subway to an elderly man	Friendly-Friendly
Insulted a secretary in his office without provocation	Gave up his seat on the subway to an elderly man	Unfriendly-Friendly
Went shopping at the market during the afternoon	Gave up his seat on the subway to an elderly man	Neutral-Friendly
Can speak three different languages fluently	Solved a complicated mathematics problem in his spare time	Intelligent-Intelligent
Was visibly confused by the map in the subway station	Solved a complicated mathematics problem in his spare time	Stupid-Intelligent
Ate a bacon and cheese hamburger for lunch	Solved a complicated mathematics problem in his spare time	Neutral-Intelligent
Was willing to try the new and exotic food	Told his friend he was interested in backpacking across U.S.	Adventurous-Adventurous
Reluctantly followed the guide and other people into the cave	Told his friend he was interested in backpacking across U.S.	Scared- Adventurous
Took a short nap in the middle of the afternoon	Told his friend he was interested in backpacking across U.S.	Neutral-Adventurous
Cleaned up the picnic area before leaving	Made sure all the bills were paid at the end of the month	Responsible-Responsible
Purposely did not finish his work, so someone else had to	Made sure all the bills were paid at the end of the month	Irresponsible-Responsible
Heated a cup of milk in the microwave	Made sure all the bills were paid at the end of the month	Neutral-Responsible
Refused to cheat on an exam	Returned a wallet he found to the owner	Honest-Honest
Cheated in a game of Monopoly	Returned a wallet he found to the owner	Dishonest-Honest
Took his car to the carwash to get it cleaned	Returned a wallet he found to the owner	Neutral-Honest
Was interested in exploring the new art sculpture	Watches the evening news regularly to keep up on current events	Curious-Curious
Sleeps through most of his class lectures and does not take notes	Watches the evening news regularly to keep up on current events	Uncurious-Curious

Material Replication I (Cont.)		
Behavior 1	Behavior 2	Traits
Made some orange juice for breakfast	Watches the evening news regularly to keep up on current events	Neutral-Curious
Told several amusing stories	Made new friends easily	Extroverted-Extroverted
Did not maintain eye contact during the conversation	Made new friends easily	Introverted-Extroverted
Took a leisurely walk around the block	Made new friends easily	Neutral-Extroverted
Told a lie about a good friend	Couldn't keep a secret	Disloyal-Disloyal
Defended an unjustly criticized person	Couldn't keep a secret	Loyal-Disloyal
Bought stamps at the post office	Couldn't keep a secret	Neutral-Disloyal

Material Replication II		
Behavior 1	Behavior 2	Traits
Shouted at the waitress when she bought the wrong order	Refused to go on a coffee break with a friend	Unfriendly-Unfriendly
Gave up his seat on the subway to an elderly man	Refused to go on a coffee break with a friend	Friendly- Unfriendly
Went shopping at the market during the afternoon	Refused to go on a coffee break with a friend	Neutral-Unfriendly
Did not understand the joke	Was visibly confused by the map in the subway station	Stupid-Stupid
Solved a complicated mathematics problem in his spare time	Was visibly confused by the map in the subway station	Intelligent- Stupid
Ate a bacon and cheese hamburger for lunch	Was visibly confused by the map in the subway station	Neutral-Stupid
Was made nervous by the glass elevator and wouldn't go on it	Reluctantly followed the guide and other people into the cave	Scared-Scared
Told his friend he was interested in backpacking across U.S.	Reluctantly followed the guide and other people into the cave	Adventurous -Scared

Material Replication II (Cont.)

Behavior 1	Behavior 2	Traits
Took a short nap in the middle of the afternoon	Reluctantly followed the guide and other people into the cave	Neutral-Scared
Went out drinking the night before an important exam	Purposely did not finish his work, so someone else had to	Irresponsible-Irresponsible
Made sure all the bills were paid at the end of the month	Purposely did not finish his work, so someone else had to	Responsible-Irresponsible
Heated a cup of milk in the microwave	Purposely did not finish his work, so someone else had to	Neutral-Irresponsible
Cheated in a game of Monopoly	Claimed credit for someone else's idea	Dishonest-Dishonest
Returned a wallet he found to the owner	Claimed credit for someone else's idea	Honest- Dishonest
Took his car to the carwash to get it cleaned	Claimed credit for someone else's idea	Neutral-Dishonest
Never votes in any elections	Sleeps though most of his class lectures and does not take notes	Uninterested-Uninterested
Watches the evening news regularly to keep up with current events	Sleeps though most of his class lectures and does not take notes	Curious-Uninterested
Made some orange juice for breakfast	Sleeps though most of his class lectures and does not take notes	Neutral- Uninterested
Was reluctant to ask the bus driver for instructions	Did not maintain eye contact during the conversation	Introverted-Introverted
Made new friends easily	Did not maintain eye contact during the conversation	Extroverted-Introverted
Took a leisurely walk around the block	Did not maintain eye contact during the conversation	Neutral-Introverted
Supported the team even with their five year losing streak	Defended an unjustly criticized person	Loyal-Loyal
Couldn't keep a secret	Defended an unjustly criticized person	Disloyal-Loyal
Bought stamps at the post office	Defended an unjustly criticized person	Neutral-Disloyal

APPENDIX G

Stimuli Material and Instructions of Experiment 7

Stimuli used for the Experiment 7

Sentences and Trait Dimensions

Friendly Behaviors

Shared the umbrella with a stranger during the rain
Visited a sick friend in the hospital
Gave up his seat on the subway to an elderly man
Helped a neighbor move boxes up the stairway of the apartment house
Invited an unpopular co-worker to go to lunch with him
Showed a foreign exchange student around campus

Intelligent Behaviors

Designed a new computer system for his company
Had several job offers upon graduating from college
Impressed the speaker afterward by asking perceptive questions
Can speak three different languages fluently
Discussed some new economic theories with a colleague
Solved a complicated mathematics problem in his spare time

Musical Behaviors

Has a large collection of CDs
Has a really great sense of rhythm
Sings along with songs on the radio
Traveled a long way to see a band concert
Plays his guitar every single day
Likes listening to music while studying

Athletic Behaviors

Jogs every morning before going work
Goes regularly to sports stores
Plays basketball in a city recreation league
Bought a new tennis racket
Does sit-ups every morning
Loves to go surfing in the beach

Instructions Experiment 7

Impression Formation

Instruction I

One of the most interesting aspects of social interaction is how we can form clear impressions of the persons with whom we interact, often with minimum information available and with little effort. And often those impressions are quite accurate. Although this happens all the time in our everyday experiences, we don't fully understand how it happens.

This experiment is concerned with the way in which we form an impression of a person on the basis of his or her actions. In a few moments you will be shown a series of sentences, each sentence containing a single description of a behaviour performed by a person named John.

Press "Spacebar" to continue.

Instruction II

Please read these behaviors carefully, studying each one until the next one appears on the screen. Try to form an overall impression of what John is like while reading the behaviors. At the end of the session, we will ask you some questions concerning the impression that you have formed of John.

Please read John's behaviors carefully, try to form an impression of his personality, and try to imagine the kind of person John is.

Please press the "Spacebar" to start the presentation of the behaviors.

Memory

Instruction I

It is one of the most fascinating aspects of human cognition how we can, often with minimum available information and without effort, to memorize so many and diverse information. And often those memories reveal themselves as being quite accurate.

This experiment is concerned with the way in which we memorize different sentences. In a few moments you will be shown a series of slides, each slide containing a single sentence.

Press “Spacebar” to continue.

Instruction II

Please read these sentences carefully, studying each one until the next slide appears on the screen. Try to remember the exact wording of each single sentence as accurately as you can. At the end of the session, we will ask you some questions pertaining to the information contained in these sentences.

Please read carefully the sentences and try to retain them in your memory so that you can reproduce them later.

When you are finished, please press “Spacebar” for starting sentences presentation.

Free Recall

(Impression Formation)

Earlier in this experiment you read a series of sentences, each of which described a behavior performed by a person named John. We would like you to recall as many of these behaviors as possible. We realize that it would be impossible for you to remember the behaviors but please try to recall as many as you can as precisely as possible.

Your task is to recall the behaviors you had read earlier. Please do not provide your general impressions of John. Recall as many of the sentences as you can in any order.

Free Recall

(Memory)

Earlier in this experiment you read a series of sentences. We would like you to recall as many of these sentences as possible. We realize that it would be impossible for you to remember the sentences word for word but please try to recall as many as you can as precisely as possible.

Your task is to recall the sentences you had read earlier. Recall as many of the sentences as you can in any order.

(Instruction added to Trait-Cue Conditions)

In order to help you in this task we will provide four words that may help you recall the sentences. Those words will be presented in the next screen. Try to use those words as memory cues while you recall the sentences. You will have a few minutes to complete the task. Please, continue in your attempt to recall as many sentences as you can.

APPENDIX H

Stimuli Material and Instructions of Experiment 8

Stimuli used for the Experiment 8

Material Replication I

Name	Behavior	Trait	Trial	Condition	Valence
Edward Williams	Was so responsible that he paid his taxes early	Responsible	Self	Trait	Pos.
John Moore	Was so introverted that he was reluctant to ask the bus driver for instructions	Introverted	Self	Trait	Neg.
Thomas Taylor	Expressed his truthful opinion on the new changes at work	Honest	Self	Mismatch	Pos.
Christopher Allen	Visited a sick friend in the hospital	Friendly	Self	Implied	Pos.
Jeff Hill	Proposed playing tackle football instead of flag football	Aggressive	Self	Mismatch	Neg.
Kevin Garcia	Lied on the entry form to be eligible for the contest	Disloyal	Self	Implied	Neg.
Steven Lewis	Said that his friend Glenn was so curious that was interested in exploring the new art sculpture	Curious	Other	Trait	Pos.
Daniel Parker	Said that his friend Bob was so stubborn that he only rarely changed his mind	Stubborn	Other	Trait	Neg.
Mark Brown	Said that his friend Andy was interested in backpacking across the U.S.	Adventurous	Other	Mismatch	Pos.
Paul Campbell	Said that his friend Isaac can speak three different languages fluently	Intelligent	Other	Implied	Pos.
James Rogers	Said that his friend Rafael had pizza delivered even though he lived only a block away	Lazy	Other	Mismatch	Neg.
Carl Smith	Said that his friend Lester accidentally deleted all files on his computer	Distracted	Other	Implied	Neg.

Material Replication II

Name	Behavior	Trait	Trial	Condition	Valence
Edward Williams	Was so irresponsible that he went drinking the night before an important exam	Irresponsible	Self	Trait	Pos.
John Moore	Was so extroverted that he told several amusing stories and jokes	Extroverted	Self	Trait	Neg.
Thomas Taylor	Put the wallet he found in his own pocket	Dishonest	Self	Mismatch	Pos.
Christopher Allen	Refused to go with a friend to a coffee break	Unfriendly	Self	Implied	Pos.
Jeff Hill	Noticed a lost old man and helped him find his way	Caring	Self	Mismatch	Neg.
Kevin Garcia	Supported the team even during their five year losing streak.	Loyal	Self	Implied	Neg.
Steven Lewis	Said that his friend Glenn was so uninterested in politics that he never votes in any elections	Uninterested	Other	Trait	Pos.
Daniel Parker	Said that his friend Bob was so flexible that he considered the opposing point of view	Flexible	Other	Trait	Neg.
Mark Brown	Said that his friend Andy wouldn't go on it on the glass elevator	Scared	Other	Mismatch	Pos.
Paul Campbell	Said that his friend Isaac was visibly confused by the map in the subway station	Stupid	Other	Implied	Pos.
James Rogers	Said that his friend Rafael rewrote the letter until it was perfectly phrased	Determined	Other	Mismatch	Neg.
Carl Smith	Said that his friend Lester told that his acquaintance kept the apartment orderly and clean	Organized	Other	Implied	Neg.

Reading Time Measure
Consistent, Inconsistent, and Neutral Behaviors

Consistent (Replication I) / Inconsistent (Replication II)	Inconsistent (Replication I) / Consistent (Replication II)	Neutral Behaviors
Cleaned up the picnic area before leaving (Responsible)	Purposely did not finish his work so someone else had to (Irresponsible)	Heated a cup of milk in the microwave
Did not maintain eye contact during the conversation (Introverted)	Made new friends easily (Extroverted)	Took his car to the carwash to get it cleaned
Refused to cheat on an exam (Honest)	Claimed credit for someone else's idea (Dishonest)	Made some orange juice for breakfast
Shared the umbrella with a stranger during the rain (Friendly)	Loudly criticized the clothes of the other people at the party (Unfriendly)	Took a leisurely walk around the block
Shouted at the waitress when she brought the wrong order (Aggressive)	Sent out Christmas cards to all his acquaintances (Caring)	Bought stamps at the post office
Cheated on his income tax report (Disloyal)	Defended an unjustly criticized friend (Loyal)	Went shopping at the market during the afternoon
Watches the evening news regularly to keep up on current events (Curious)	Sleeps during most of his class lectures and does not take notes (Uninterested)	Ate a cheese hamburger for lunch
Refused to accept needed help on a project (Stubborn)	Always apologizes when he feels is wrong (Flexible)	Ate a piece of fruit after dinner.
Was willing to try the new and exotic food (Adventurous)	Decided not to surf because of the threat of sharks (Scared)	Bought a magazine in his way home
Solved a complicated mathematics problem in his spare time (Intelligent)	Did badly on the easy exam (Stupid)	Took the key of his pocket to open the door
Fell asleep at work while the boss was out (Lazy)	Runs five miles a day to keep in shape (Determined)	Watched a late movie on television
Put on different pairs of shoes and didn't notice (distracted)	Threw his clothes in their usual place on the chair (Organized)	Drank a coffee during the morning

Instructions Experiment 8

Instruction I

In the first part of this experiment, you will see names of different people. Each name will be paired with a description of a behavior. Sometimes the sentence describes a behavior performed by the person named in the sentence (self-descriptive condition). Other times the person named in the sentence describes a behavior performed by an acquaintance (other-descriptive condition).

These behaviors were selected from a longer list of statements that the named persons have provided in response to questions from an interview.

Press “Spacebar to continue”.

Instruction II

This experiment is concerned with the way in which we form an impression of a person on the basis of his or her actions. When the information is presented, read the behaviors carefully and form an impression of the personality of each person described. Please try to form an impression of the person described by each sentence since that information will be useful later in the experiment. In order to familiarize you with the task, we will present you with two practice trials. Press the “Spacebar” in order to see the practice trials. **(Impression Formation)**

This experiment is concerned with the way in which we remember verbal descriptions of actions. When the information is presented, read the sentences carefully and try to retain them in your memory so that you can reproduce them later. Please try to remember the names and the accompanying statements, since your memory for these pairings will be assessed later in the experiment. In order to familiarize you with the task, we will present you with two practice trials. Press the “Spacebar” in order to see the practice trials. **(Memory)**

Instruction III

Now that you are familiar with the task, the study will begin.

In a few moments you will see a series of screens. Each screen will consist of a name and a behavior description. Sometimes the behavior described was performed by the named person (self description condition) while other times the named person is describing a behavior from an acquaintance (other descriptive condition)

Each screen will be presented for 5 seconds.

Pay attention and read the information carefully. Your memory for this information will be tested later. (**Memory**)

Form an impression about each person described. You will be asked about your impressions of these persons later. (**Impression**)

Please press the “Spacebar” to start the presentation of the behaviors.

False Recognition Measure

Instructions I

We will now present you with the names you saw in the first part of the study, and each one will be accompanied by a single word. Your task is to decide whether that word was in the sentence paired with the name presented.

After you are presented with the name and word, press the button “I” on the keyboard if that exact word was in the sentence paired with the name shown, and press the button “E” on the keyboard if the word was not in the sentence paired with the name shown.

Please respond as quickly and as accurately as possible.

In order to familiarize you with the task, we will present you some practice trials. Press the “Spacebar” in order to see the practice trials.

Instructions II

(After Practice)

Now that you are familiar with the task, we may start the experimental trials. Your task is to decide whether you had seen the word in the sentence paired with the name presented. Press the “I” key if you believe that you had seen the word in the study phase or the “E” key if you believe you had not seen the word.

To help you complete the task as fast as you can, please keep your index fingers on the E (“No”) and I (“Yes”) keys throughout the task. This will help you respond more quickly, as you won’t have to move your hand to after you make your decision.

Click the “Spacebar” to continue with the experiment.

Reading Time Measure

In this phase of the research we will continue the first phase of the study. Specifically, we will present you with a list of additional behaviors. These behaviors were performed by the same persons named in the first earlier set of sentences, so you will see the same names again. You just have to read the behaviors at a comfortable, normal reading pace. After reading each behavior simply press “Enter” in the keyboard to go to the next sentence. Press the spacebar to start presentation of the sentences.