ISCTE IUL REPOSITÓRIO
INSTITUTO UNIVERSITÁRIO DE LISBOA

Repositório ISCTE-IUL

Deposited in *Repositório ISCTE-IUL*:
2022-04-05

Deposited version:
Accepted Version

Peer-review status of attached file:
Peer-reviewed

Citation for published item:
Silva, R. R. & Unkelbach, C. (2021). Fluent processing leads to positive stimulus evaluations even when base rates suggest negative evaluations. Consciousness and Cognition. 96

Further information on publisher's website:
10.1016/j.concog.2021.103238

Publisher's copyright statement:
This is the peer reviewed version of the following article: Silva, R. R. & Unkelbach, C. (2021). Fluent processing leads to positive stimulus evaluations even when base rates suggest negative evaluations. Consciousness and Cognition. 96, which has been published in final form at https://dx.doi.org/10.1016/j.concog.2021.103238. This article may be used for non-commercial purposes in accordance with the Publisher's Terms and Conditions for self-archiving.

Use policy

Creative Commons CC BY 4.0
The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a link is made to the metadata record in the Repository
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Serviços de Informação e Documentação, Instituto Universitário de Lisboa (ISCTE-IUL)
Av. das Forças Armadas, Edifício II, 1649-026 Lisboa Portugal
Phone: +(351) 217 903 024 | e-mail: administrador.repositorio@iscte-iul.pt
https://repositorio.iscte-iul.pt

Running Head: Pseudocontingencies and Fluency

Fluent processing leads to positive stimulus evaluations even when base rates suggest negative evaluations

Rita R. Silva[a] and Christian Unkelbach[b]

[a] CIS_ISCTE, ISCTE-University Institute of Lisbon. Avenida das Forças Armadas, Edifício Sedas Nunes, Sala 2W08. 1649-026 Lisboa, Portugal. [b] University of Cologne. Richard-Strauß-Strasse 2. 50931 Cologne, Germany.

E-mails: rita.silva@uni-koeln.de; christian.unkelbach@uni-koeln.de.

Corresponding author: Rita R. Silva (rita.silva@iscte-iul.pt)

Accepted for publication in Consciousness and Cognition: Silva, R. R., & Unkelbach, C. (2021). Fluent processing leads to positive stimulus evaluations even when base rates suggest negative evaluations. *Consciousness and cognition*, *96*, 103238. https://doi.org/10.1016/j.concog.2021.103238

This research was supported: by the Department of Innovation, Science and Research of the state of North Rhine-Westphalia (Ministerium für Innovation, Wissenschaft und Forschung des Landes Nordrhein-Westfalen, Programm Nachwuchsforschungsgruppe, NRW 2015–2021, "Erlebte Sicherheit"); by Portuguese national funds through FCT—Fundação para a Ciência e a Tecnologia, I.P., within the project UIDB/03125/2020; by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 844296.

Abstract

Fluency is the experienced ease of ongoing mental operations, which increases the subjective positivity of stimuli attributes. This may happen because fluency is inherently positive. Alternatively, people may learn the meaning of fluency from contingencies within judgment-contexts. We test pseudocontingencies (PCs) as a mechanism through which fluency’s meaning is learned. PCs are inferred correlations between two attributes due to the observation of their jointly skewed base rates – people relate what is frequent in one attribute to what is frequent in the other. Using online seller evaluations as the dependent variable, we manipulated base rates of seller name-fluency and seller reputation, creating conditions where name-fluency aligned positively or negatively with reputation. However, participants evaluated high-fluency name sellers more positively across base-rate conditions, although we observed negative PCs between seller reputation and a fluency-neutral dimension in a follow-up study. We discuss the implications for the debate regarding fluency's positive vs. malleable nature.

<150 words >

Keywords: fluency, positivity, pseudocontingencies, pronounceability,

1. Fluent processing leads to positive stimulus evaluations even when base rates suggest negative evaluations.

Processing fluency is the experienced ease of ongoing mental processes; it may result from the encoding, the retrieval, the processing, or the production of information (Alter & Oppenheimer, 2009; Greifeneder & Unkelbach, 2013). In other words, people experience whether a stimulus is easy or difficult to perceive, to retrieve, to process, or to evaluate. This experience influences many judgments; from judgments that are closely related to the experience itself, such as familiarity (Whittlesea, 1993), or judgments that are farther removed, such as a person's intelligence (Oppenheimer, 2006; see Unkelbach & Greifeneder, 2013; for an overview of judgmental effects).

Historically, processing fluency gained prominence because it serves as an explanation for the mere exposure effect (Zajonc, 1968). For example, Bornstein and D'Agostino (1994) argued that repeated exposure facilitates stimulus processing; they argued that fluent processing feels positive and because the positive feeling is misattributed to the stimulus itself, a positive stimulus evaluation follows. While this interpretation has been criticized (Zajonc, 2001) and also experimentally challenged (Newell & Shanks, 2007), it led to a wealth of research showing that manipulations that influence the ease with which people process stimuli increases participants' rating of positive stimulus attributes, such as increased liking, beauty, safety, trustworthiness, or truth, suggesting that fluency is positively connotated (see Graf et al., 2018; Reber et al., 2004).

In addition, there is psycho-physical evidence that people indeed experience processing fluency positively (e.g., Harmon-Jones & Allen, 2001; Winkielman & Cacioppo, 2001)[1]. Winkielman and colleagues (2003) summarized the available evidence of fluency on evaluative judgments as the "hedonic marking hypothesis". Another variant of this model does not assume that fluency is inherently positive, but that the experience amplifies existing evaluations (Albrecht & Carbon, 2014; see Landwehr & Eckmann, 2020, for a comparison). This variant has some support, but it needs the additional assumption that stimuli are likeable, safe, trustworthy, or true, to begin with. This is at odds with many empirical findings. For example, processing fluency increases the rated truth of statements when they are known to be false (Fazio et al., 2015), or when they are clearly labelled as false (Unkelbach & Greifeneder, 2018). Similarly, fluency increases ratings of stimuli as "old/familiar", despite them being factually new (Whittlesea, 1993).

Another line of research suggested that the influence of the fluency experience is malleable. The seminal finding for this line was presented by Mandler et al. (1987). These authors observed that repeated exposure increased participants' judgments of liking for irregular polygons, but also their judgments of lightness and blackness. Thus, processing fluency may influence any judgment dimension; it may be independent of its evaluative connotation but dependent on participants' interpretation.

For example, people typically agree with propositions if they can easily generate arguments for the proposition. Briñol et al. (2006) showed that this attitude effect reverses when participants interpret the ease of argument generation as bad rather than good. In their study, the authors influenced participants' interpretation of the fluency experience by explicit instructions. They asked participants to come up with either a small or large (i.e., easy or difficult, respectively) number of arguments in favor of a new evaluation policy and then to report their attitudes about it. When participants were told that ease in argument generation is a sign of intelligence, participants in the small number of arguments condition felt more positively about the policy. But when ease was associated with low intelligence, then participants in the small number of arguments condition reported feeling more negatively about the policy.

Using a learning task rather than explicit instruction, Unkelbach (2006, 2007) showed that after completing a feedback-learning task associating high fluency with novelty or falseness and low fluency with familiarity or truth, people on average interpreted fluently processed stimuli as being "new" rather than "old" and fluently processed statements as "false" rather than "true" (see also Olds & Westerman, 2012; Corneille et al., 2020). The general idea within this line is that people use processing fluency as a cue to inform their judgments, and the cue’s meaning and validity is context-dependent (Hertwig et al., 2008; Herzog & Hertwig, 2013).

Thus, there is evidence for the inherent positivity of processing fluency, as well as for the malleability of processing fluency's influences. The former explanation predicts unqualified main effects of fluency on positively connotated judgment attributes, while the latter predicts interactions, depending on the context, the judgment attributes, or the learning history. These lines of research are not incompatible. For the malleability position, one may still assume that fluency's positive connotation has an effect, which would imply that fluency is more readily associated with positive criteria (e.g., truth, liking), rather than negative criteria (e.g., falsity, disliking; see Mandler et al., 1987). However, the prediction of unqualified main effects does not allow the reversal, and from a logic of science perspective, it is, therefore, the stronger position, as it is easier to falsify.

1.1 The Present Research

A central question for both lines of research is how people acquire the meaning of the fluency experience. For the assumption that fluency is a positive experience, Winkielman and colleagues (2003) provided three arguments. First, fluency may be inherently pleasant: the same way people "enjoy" feeling positive emotions, people may enjoy the ease of processing. Second, fluency may indicate the successful running/processing of mental operations and/or their completion; thereby, it is a positive experience. Third, fluency signals competence and mastery, which is also a positive state.

For the assumption of fluency's malleability, there are two classes of explanations that allow variable interpretations of the fluency experience. The first class assumes that participants have lay theories about the fluency experience’s meaning (e.g., Schwarz, 2006, 2015). For example, people may *believe* that instances that come fluently to mind are frequent or that fluent argument generation indicates intelligence. Thereby, beliefs about the meaning of the experience shape the resulting judgment (see Unkelbach & Greifeneder, 2013; Wänke, 2013). The second class assumes that the interpretation is learned: for example, people learn ecologically that frequent events come easily to mind (Hertwig et al., 2009; Herzog & Hertwig, 2013; Unkelbach, 2006).

The assumptions for an unconditional positive connotation of fluency are almost a priori true. For both explanations of malleability, naïve theories and ecological learning, the origins of the interpretations afford an explanatory challenge. Assuming that the environment rarely provides direct learning opportunities or statements about the meaning of subjective experiences, how do people acquire the meaning of fluency? Here, we suggest *pseudocontingencies* (PC; Fiedler & Freytag, 2004; Fiedler et al., 2009) as a possible explanation of how people learn the interpretation of processing fluency. For example, to explain positive fluency effects on stimulus evaluations, the PC explanation only requires that positive information is more prevalent than negative information and that easy-to-process information is more prevalent than difficult-to-process information (see Unkelbach et al., 2019, 2020; for reviews). We delineate this explanation below and test it in one experiment in comparison with the hedonic marking hypothesis (Winkielman et al., 2003).

1.2 A Pseudocontingencies Paradigm

Pseudocontingencies are people's perception of a correlation between two attributes, based on the alignment of two skewed base-rates. (e.g., Fiedler & Freytag, 2004; Fiedler et al., 2009). If the base-rates are skewed in the same direction, people infer a positive correlation between the two attributes. If the base-rates are skewed in opposite directions, people infer a negative correlation. Let us illustrate this abstract notion with an example. Assume that most seller ratings in an online market are positive, and most sellers are from the US; according to PCs, people should infer that US sellers are good; and conversely, that sellers from other countries are less good.

This inference influences phenomena from low-level processes such as operant learning (e.g., Kutzner et al., 2008), to classic social-cognitive phenomena such as stereotyping (e.g., Fiedler et al., 2007), to deliberate, outcome-dependent decisions (e.g., Vogel & Kutzner, 2017). Kutzner and colleagues (2011) also showed by simulations that PC inferences from base-rates and factual correlations within data sets align most of the time, highlighting the adaptive nature of PCs as a simplified strategy to detect correlations in the environment.

There are three central points regarding PCs. First, PCs do not require perception and encoding of the joint occurrences of the attributes but require only encoding of each attribute's base-rates. The only pre-requisite seems to be that participants encode base-rate alignment (e.g., both US sellers and positive ratings must be frequent). Second, the skewed base-rates allow for factual correlations to be positive, negative, and zero. In this sense, PCs represent a potentially erroneous inference (but see Kutzner et al., 2011). Third, although

PCs are logically not warranted, they may represent a case of adaptive cognition: the stronger the skew in the base-rates, the more constrained the factual contingencies. Because it is substantially easier to track base-rates instead of contingencies (see Fiedler et al., 2013), pseudocontingencies may constitute a simplified cognitive strategy to detect environmental contingencies.

Factual contingencies between fluency and positivity may also exist (i.e., positive information is in general processed and retrieved faster; Unkelbach et al., 2008, 2010). However, the PC case would afford an elegant explanation for the learning of fluency's meaning, as they are less susceptible to the vagaries of contingency learning tasks. In the following, we delineate how we tested the PC paradigm as an explanation in comparison to fluency's potential inherent positivity.

1.3 The Present Experiment

To test whether PCs may explain how individuals learn the interpretation of processing fluency, we designed an experiment in which participants read the usernames of online sellers and their average seller reputation. We varied two attributes of the sellers: the fluency of username pronounceability (i.e., fluent vs. disfluent) and the online sellers' reputation (i.e., "good" vs. "bad"). From a PC perspective, participants should infer a positive correlation between the two attributes when the skews of their distributions are in the same direction. Thus, when both fluent usernames and good reputations are frequent, participants should infer that fluency correlates with positive reputations. Following the same logic, when both fluent usernames and bad reputations are frequent, participants should infer a negative correlation between fluency and reputation.

The DV is then participants' predictions of new sellers' reputations based on seller name fluency. In the former case of a positive fluency-good reputation PC, participants should predict positive reputations for sellers with high fluency names. In the latter case of a negative fluency-good reputation PC, participants should predict positive reputations for sellers with disfluent names. Such a reversal would be in line with the hypothesis that the experience of fluency itself is (evaluatively) non-specific and malleable, and that the attribution of meaning to fluency experiences may be supported by a PC mechanism.

If, however, fluency is an inherently positive experience that consistently leads to more positive evaluations (Winkielman et al., 2003), then the association of fluency with positive outcomes should be hard to reverse by the mere observation of oppositely skewed distributions of the two attributes. We test these competing predictions in the following experiment.

2. Method

2.1 Participants and Design

To determine the appropriate number of participants to detect the effect of interest – an interaction of the PCs' direction and the fluency of seller name - we performed an a priori power analysis using G*Power (Faul et al., 2007). Setting the parameters to a small effect size of $f = 0.1$ (given the lack of a reference effect size from previous studies manipulating PC learning and name fluency), $\alpha = .05$, $1 - \beta = .80$, $\rho = .50$, the analysis showed $N = 305$ participants to be a sufficient sample size. We were able to recruit 406 participants ($M_{age} = 23.01$, $SD = 3.77$; one participant did not report demographic data) on a university campus. Participants took part voluntarily and received candy bars as a "thank you" for their participation. Participants were randomly assigned to one of 5 PC conditions, illustrated in Table 1 (PC condition: FluentHigh-GoodHigh vs. FluentHigh-BadHigh vs. DisfluentHigh-GoodHigh vs. DisfluentHigh-BadHigh vs. Control - same frequency of Fluent-Disfluent and

Good-Bad attributes). Username fluency (fluent vs. disfluent) of the sellers rated at the end varied within-participants.

Table 1. Frequency of the Fluent vs. Disfluent usernames and Good vs. Bad reputation per PC condition (top row) and respective PC Inference (bottom row).

PC Condition	FluentHigh-GoodHigh	FluentHigh-BadHigh	DisfluentHigh-GoodHigh	DisfluentHigh-BadHigh	Control
Frequent Usernames (Fluent vs. Disfluent)	Fluent	Fluent	Disfluent	Disfluent	Same
Frequent Reputation (Good vs. Bad)	Good	Bad	Good	Bad	Same
PC Inference	Fluency = Positivity	Fluency = Negativity	Fluency = Negativity	Fluency = Positivity	None

Note. There were only clearly positive and negative reputations, and there were only clearly fluent and disfluent seller names. Thus, the PC predictions are always based upon the alignment of high frequencies, that is, the alignment of the base-rates.

2.2 Materials and Procedure

2.2.1 Seller usernames. As usernames, we used 140 stimuli by Silva and colleagues (2017a), who investigated username fluency as a predictor of seller trustworthiness. These stimuli consisted of letter strings forming meaningless words in the German language (the language used in the experiment) that were previously pre-tested as easy-to-pronounce or fluent (70 stimuli; e.g., galmug, gekikite, usanitido) and as difficult-to-pronounce or disfluent

(70 stimuli; e.g., eakrtb, tsrcneha, vlegtiqclapl). Table 1 in Appendix A contains the complete list of stimuli (for details on the pre-tests, see Silva et al., 2017a; Topolinski et al., 2016).

2.2.2 Reputation. We used a five-star rating system to manipulate seller reputation (e.g., Silva et al., 2017a; Silva & Topolinski, 2018), as this system is widely used in different commercial and non-commercial contexts (e.g., hotel ratings, book reviews). To emphasize the distinction between sellers with good and bad reputation, we used five golden yellow stars to indicate sellers with good reputation and one star to indicate sellers with bad reputation.

2.2.3 Procedure. We programmed the experiment with the online survey platform Qualtrics. The Qualtrics survey administered all instructions, materials, and dependent measures. Research assistants used laptops to present the survey and register participants' responses. They approached participants at different university campus sites (e.g., library, study areas) and if they consented to participate in the experiment, a laptop was handed to them. There were two sequential tasks in the experiment, the PC induction phase and the target evaluation phase. For each between-participants condition inducing the different PCs between name fluency and reputation (see below), we created two different stimuli lists varying the assignment of the fluent and disfluent usernames to the two phases of the experiment, as well as the pairing of the usernames with the reputation levels in the PC induction phase. Both the assignment of usernames to the two phases of the experiment and to the reputation levels was done randomly. The stimuli lists from each PC condition can be found at https://tinyurl.com/49s2wcet.

***2.2.3.1 PC induction*.** The survey informed participants that they would see slides containing different eBay seller usernames and their average seller reputation. Participants' task was to read each of the seller slides. Then, the survey presented 60 eBay seller usernames and their reputation. The username–reputation pairs were shown on the computer screen one by one, with the username placed in the center of the screen and the reputation stars placed below. Figure 1 shows examples of these slides. Presentation was self-paced and participants had to click a button placed at the bottom of the screen to go from one seller to the next.

Depending on the PC condition, the frequency of fluent/disfluent seller names and seller reputation varied. For example, in the PC condition "FluentHigh-GoodHigh", participants saw 50 fluent usernames and 10 disfluent usernames. Within username fluency, most reputations were good: 41 good and 9 bad reputations, and 9 good and 1 bad reputation, respectively. These distributions actually indicated a small negative correlation: sellers with disfluent/fluent names were more likely of positive/negative reputation, respectively (r = -.08). A PC illusion should emerge due to the alignment of the skewed base-rates; people tend to associate frequent with frequent, and infrequent with infrequent. Thus, participants should infer that sellers with *fluent* names are more likely to have a positive reputation and that sellers with disfluent names are more likely to have a negative reputation.

Conversely, in the PC condition "DisfluentHigh-GoodHigh", the base-rate of disfluent names aligned with the frequency of good reputations, which should lead to a PC illusion that sellers with *disfluent* names are more likely of positive reputation, and vice versa. The factual correlation is again in the opposite direction.

Tables 2 to 6 in Appendix A present the base rates of fluent vs. disfluent usernames and of good vs. bad reputation occurrences in each PC condition and in the control condition (in which the levels of the two attributes co-occurred exactly the same number of times, i.e., 15 presentations per each of the four possible fluency-reputation combinations).

The survey randomized order of presentation of the fluent vs. disfluent – good vs. bad reputation pairs anew for each participant.

2.2.3.2 Target evaluation. Next, the target evaluation phase started. The survey informed participants that they were now going to evaluate a different set of eBay sellers based on their usernames. Their task was to indicate what they thought was the average reputation of each seller on a scale ranging from 1 to 5 stars. Participants judged the reputation of 20 new sellers, half with fluent usernames and half with disfluent usernames, randomly ordered anew for each participant. In each trial, the username was presented in the screen center with the 5-star rating scale below it. Participants clicked on as many stars as they thought the seller reputation was. This action colored the stars in golden yellow. To proceed, participants had to click a button placed at the bottom of the screen. After the evaluation of the 20 sellers, participants indicated their age, gender ("female", “male”, or “other”), and whether German was their native language. Finally, participants were thanked and debriefed. The experiment lasted approximately 7 minutes.

Figure 1. Examples of the slides with seller names (disfluent names on the left; fluent names on the right) and reputation (good reputation on top; bad reputation on lower slides).

2.2.4 Predictions. The two critical conditions in our experiment are those when the PC inference leads to a negative correlation between fluency and positivity; that is, when disfluent names and good reputations are both frequent, or when fluent names and bad reputations are frequent. The two lines of research discussed in this work make opposite predictions for these two negative PCs conditions: for the perspective proposing that the interpretation of fluency experiences is malleable, participants in the negative PCs conditions should evaluate target sellers with fluent names more negatively than sellers with disfluent names, showing evidence of a reversal of the typical “fluency is positive” effect. For the perspective proposing that fluency experiences are inherently positive, such a reversal should be difficult to occur through the simple experience of oppositely skewed base-rates, and participants should always evaluate target sellers with fluent names more positively than sellers with disfluent names.

3. Results

Before analyzing the data, we excluded 68 participants because they indicated that German was not their native language, which could interfere with the effectiveness of our fluency manipulation based on ease of pronouncing German words; however, the pattern of results did not differ when these non-native participants were included in the analyses.

We analyzed participants’ ratings of seller reputation of the remaining 338 participants with a PC condition (FluencyHigh-GoodHigh vs. FluencyHigh-BadHigh vs. DisfluencyHigh-GoodHigh vs. DisfluencyHigh-BadHigh vs. Control) × Username (fluent vs. disfluent) mixed ANOVA, with repeated measures on the latter factor. The main effect of username fluency was significant, $F(1, 333) = 107.97$, $p < .001$, $\eta^2_p = .245$; participants assigned higher reputation ratings to sellers with fluent compared to disfluent usernames (see Table 2, last column). The PC condition effect was also significant, $F(4, 333) = 60.28$, $p < .001$, $\eta^2_p = .420$. As Table 3's bottom row shows, high frequency of good reputations led to higher reputation ratings in the target evaluation phase, while high frequency of low reputation ratings led to low ratings in the target evaluation phase. This result suggests that participants were sensitive to the manipulation of the frequency of good vs. bad reputation sellers in the PC induction phase and adjusted their target evaluations accordingly.

The interaction between the two factors was also significant, $F(4, 333) = 5.31$, $p < .001$, $\eta^2_p = .060$. The interaction reflects that although fluent usernames were always rated more positively than disfluent usernames, this difference was not significant in the "DisfluencyHigh-BadHigh" condition, $t(333) = 0.82$, $p = .412$ (for all other PC conditions, $p < .001$). This result does not support learning based on PCs, which would predict an amplification of the fluency effect, not an attenuation in this condition. Additionally, as Table 3 shows, we also did not find the expected reversals in the "FluencyHigh-BadHigh" and "DisfluencyHigh-GoodHigh" PC conditions.

Overall, these results suggest that the association of high fluency (vs. disfluency) with more positive (negative) evaluations did not change via negative PCs between the seller usernames' fluency and their reputation score.

Table 2. Mean reputation ratings (SE) given to fluent and disfluent names (lines) per PC condition (columns).

PC Condition	Fluency-Good	Fluency-Bad	Disfluency-Good	Disfluency-Bad	Control	Overall
Fluent Username	3.84 (.10)	2.68 (.09)	3.76 (.09)	2.29 (.10)	3.20 (.09)	3.16 (.04)
Disfluent Username	3.26 (.10)	1.89 (.09)	3.26 (.09)	2.20 (.10)	2.61 (.09)	2.64 (.04)
Overall	3.55 (.08)	2.28 (.07)	3.51 (.08)	2.25 (.08)	2.91 (.08)	

4. Follow-up study: Learning of Negative PCs

While the significant PC conditions effect showed that participants encoded the base-rates, one might argue that the present setup does not lead to PCs despite the encoding of the base-rates. To address this point, we conducted a follow-up study testing whether participants do learn the intended negative PCs when an attribute other than fluency is paired with good vs. bad reputation. Thus, instead of manipulating the frequency of seller username fluency, we manipulated the online market the sellers belonged to. This was necessary so that we could test that the absence of negative PCs in our study was not due to the experimental procedure employed and that they can be observed when a neutral stimulus dimension is manipulated.

In the follow-up study (described in detail in Appendix B), participants (N = 99) learned whether sellers belonged to one of two artificial online markets – Batrek vs. Galmug – and their average seller reputation (1 vs. 5 stars, as in the main study). Keeping Batrek as the most frequent online market sellers belonged to, there were two PC conditions: one in which good reputation sellers were the most frequent, thus inducing a PC illusion between Batrek sellers and positive reputation (i.e., a positive PC condition); and one in which bad reputation sellers were the most frequent, thus inducing a PC illusion between Batrek sellers and negative reputation (i.e., a negative PC condition). The base rates of the two attributes – online market and reputation – in each condition were the same as in the correspondent conditions of the main study. After the PC induction phase, participants evaluated a new set of sellers from the two markets. If this experimental setting allows participants to infer positive vs. negative PCs, then participants in the positive PC condition should evaluate Batrek sellers more positively than Galmug sellers but participants in the negative PC condition should evaluate Batrek sellers more negatively than Galmug sellers. We found exactly this PC pattern, indicated by a significant interaction between the two factors, $F(1, 97) = 10.64$, $p = .002$, $\eta^2_p = .099$: participants in the positive PC condition gave more positive evaluations to Batrek ($M = 4.23$, $SE = .12$) than to Galmug sellers ($M = 4.07$, $SE = .16$), but this pattern reversed for participants in the negative PC condition (Batrek, $M = 2.19$, $SE = .12$; Galmug, $M = 2.73$, $SE = .15$).

These results indicate that our experimental setting allows learning of negative PCs, and thus the reasons why fluency’s meaning could not be changed through the induction of negative PCs is more likely to lie in the specificities of the fluency experience than in a shortcoming of the experimental design.

5. Discussion

We discussed two lines of research that may explain fluency effects on judgments. First, fluency may be an inherently positive experience and should thus strengthen judgments on positive dimensions about a fluently-processed stimulus (e.g., Reber et al., 2004). Second, the interpretation of fluency may be malleable and learned in a given judgment context (e.g., Unkelbach, 2006, Hertwig et al., 2008). A challenge for the latter research line is to explain how people actually learn the relation of fluency with a given stimulus attribute in a complex world with multiple contingencies.

Here, we suggested pseudocontingencies (PCs) as a simple mechanism that may explain how people, for example, infer a relation between fluency and positivity. However, we found no evidence for the predicted reversals of fluency's positive meaning in neither the condition favoring the inference of PCs between fluent names and bad reputation nor the condition favoring the inference of PCs between disfluent names and good reputation. Rather, across all conditions of our experiment, fluent names were judged more positively than disfluent names. Only in the condition in which disfluent names and bad reputation occurrences were the most frequent (the "disfluent-bad" condition) did the difference between fluent and disfluent usernames not reach statistical significance. One explanation for this unexpected result (the PC prediction was an amplification of the fluency-positivity effect) may be that in the "DisfluentHigh-BadHigh" condition, disfluent usernames were presented more frequently, increasing their familiarity, which in turn could have decreased the negative impact of disfluency on reputation ratings. However, one should then expect the same result in the "DisfluentHigh-GoodHigh" condition, which did not happen; in this condition, participants rated fluent names more positively than the disfluent ones. This suggests that it was the combination of the two attributes' base rates being skewed to the negative pole that led participants to attribute low reputation values to all targets.

As such, the results of our study support the hedonic fluency model. The data are in line with assuming that fluency is inherently positive, while the pseudocontingencies had no effect. However, there are different possible reasons for why the PC induction implemented in this study was not effective in modulating the interpretation of fluency. Below, we discuss them and their implications for the line of research suggesting the malleability of fluency's meaning.

5.1 Comparison with Other Fluency-Reversal Paradigms

When comparing our experiment with studies that were successful in demonstrating the malleability of the interpretation of fluency experiences, there are several factors that differ. First, for studies suggesting that the interpretation of fluency is supported by lay theories, the information is provided explicitly. For example, in the Briñol and colleagues' (2006) study, participants were explicitly told that a feeling of ease while generating arguments is a sign that the arguments are bad. For studies suggesting that the interpretation is learned (e.g., Unkelbach, 2006; Westermann & Olds, 2012), there was a factual correlation between the fluency experience and its meaning. In addition, these studies provided feedback, which generally supports contingency learning. In our study, despite it presenting a more subtle, realistic, and more ecologically valid learning context than those paradigms, participants needed to infer the contingency from the base-rates. If they learned the factual correlation, it was basically zero. Thus, and especially if one considers the hedonic model of fluency as an inherently positive experience (Winkielman et al., 2003), learning based on PCs might not be strong enough to reverse the association of fluency with positivity.

Independent of why the PC manipulation failed here, the results present some relevant insights, namely that the fluency's positive connotation did have an influence across all PC conditions. This leads to straightforward hypotheses for future research. Fluency's learned interpretation and its inherent positivity may independently contribute to judgments, which could be examined in a Process-Dissociation paradigm (e.g., Jacoby, 1991). Alternatively, people may learn what the interpretation of positive feelings is in a given context.

5.2 Diagnosticity of Fluency Cues

Although we treat processing fluency as a unitary construct in the present research, in line with Alter and Oppenheimer (2009), there is substantial research suggesting that different fluency sources may lead to differential effects on judgments (e.g., Lanska et al., 2014; Lanska & Westerman, 2018; Silva et al., 2015; Silva et al., 2017b; Vogel et al., 2020; Whittlesea, 1993). In this line of research, reliance on a given fluency form is dependent on its perceived diagnostic value regarding the stimulus attribute being judged. Recently, Vogel et al. (2020) proposed a fluency-specificity hypothesis (see also Silva et al., 2016). According to this hypothesis, the processing experiences evoked by different fluency forms (e.g., perceptual vs. conceptual vs. linguistic fluency) are not uniform and individuals differentiate between them. As such, individuals can make a selective use of fluency experiences in contexts where different fluency cues are available. Judgments will then be primarily affected by the fluency experience that is more relevant and diagnostic of the target stimulus dimension, depending on the processing operations that are perceived as relevant for the judgment formation.

It is thus possible that the modulation of fluency effects by PCs is more or less likely to appear depending on how diagnostic the fluency cue is for the judgment. When a given fluency cue is highly diagnostic of the target stimulus-dimension, the interpretation given to the fluency experience may be more difficult to reverse than when fluency stems from a less diagnostic cue (for a demonstration in the context of judgments of truth and repetition vs. color contrast as fluency instantiations, see Silva et al., 2016). Thus, perhaps PCs were not successful in modulating fluency's interpretation in the specific setting of our experiment, but may still be a possible mechanism for extracting the meaning of fluency in judgmental contexts with different stimulus attributes. Here, the fluency associated with the pronounceability of online sellers' names may be so strongly associated with perceptions of their general credibility as sellers, perhaps due to the characteristics of the natural ecology in which online sellers mostly have pronounceable names and an average to high reputation (see Resnick et al., 2006), that this interpretation is very difficult to change through illusory, not-reinforced, PC inferences (for a set of studies showing the consistent and strong effect of pronounceability on judgements of seller trustworthiness, see Silva et al., 2017a). However, if a fluency instantiation that is less diagnostic of seller reputation, for example, the visual contrast in which seller names were presented, had been used, the predicted reversals of the positive interpretation of fluency might have been observed.

We strongly believe that experimental variations along the lines sketched here will provide a deeper insight how the experience of ongoing mental processes influences people's judgments and decisions.

6. Conclusion

Both fluency's inherent positivity as well as its interpretation have been supported by numerous studies. The present results indicate that PCs are not the learning mechanism proper and support the hedonic fluency model. In particular for clearly evaluative judgments such as those in our study, fluency's inherent positivity seems to dominate participants' judgments. Yet, this does not imply that the interpretation of fluency effects is not malleable, but the experience is probably less neutral than previously argued (e.g., Mandler et al., 1987; Unkelbach, 2006, 2007). Thus, similarly to Landwehr and Eckmann's (2020) recent work testing and showing the simultaneous contributions of the hedonic and the amplification models of fluency experiences to individuals' judgments, future research may investigate the interplay of these two explanatory approaches, rather than pitting them against each other.

Data Availability and Open Practices Statement

All the materials and datasets generated and analyzed during the course of this work are available at (https://tinyurl.com/49s2wcet). The experiments were not preregistered.

Footnotes

1: These studies showed that fluent processing increases activity in the “smiling muscle” zygomaticus major. However, the most frequently cited studies providing psycho-physical evidence have low power and small effects; we believe that these results should be treated with caution.

References

Albrecht, S., & Carbon, C. C. (2014). The Fluency Amplification Model: Fluent stimuli show more intense but not evidently more positive evaluations. *Acta Psychologica*, *148*, 195-203.

Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review, 13,* 219–235.

Bornstein, R. F., & D'Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition*, *12*, 103-128.

Briñol, P., Petty, R. E., & Tormala, Z. L. (2006). The malleable meaning of subjective ease. *Psychological Science, 17*, 200-206.

Corneille, O., Mierop, A., & Unkelbach, C. (2020). Repetition increases both the perceived truth and fakeness of information: An ecological account. *Cognition, 205*, 104470.

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behavior Research Methods, 39, 175–191.

Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General, 144*, 993–1002.

Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology, 87*, 453–467.

Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review, 116*, 187–206.

Fiedler, K., Freytag, P., & Unkelbach, C. (2007). Pseudocontingencies in a simulated classroom. *Journal of Personality and Social Psychology, 92*, 665–677.

Fiedler, K., Kutzner, F., & Vogel, T. (2013). Pseudocontingencies: Logically unwarranted but smart inferences. *Current Directions in Psychological Science*, *22*, 324-329.

Graf, L. K., Mayer, S., & Landwehr, J. R. (2018). Measuring processing fluency: One versus five items. *Journal of Consumer Psychology*, *28*, 393-411.

Greifeneder, R., & Unkelbach, C. (2013). Experiencing thinking. In C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking: How the fluency of mental processes influences cognition and behavior* (pp. 1–7). New York, NY: Psychology Press.

Harmon-Jones, E., & Allen, J. J. (2001). The role of affect in the mere exposure effect: Evidence from psychophysiological and individual differences approaches. *Personality and Social Psychology Bulletin*, *27*, 889-898.

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 1191.1206.

Herzog, S. M., & Hertwig, R. (2013). The ecological validity of fluency. In C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking: How feelings from mental processes influences cognition and behaviour* (pp.190–219). London, England: Psychology Press.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541

Kutzner, F., Freytag, P., Vogel, T., & Fiedler, K. (2008). Base-rate neglect as a function of base rates in probabilistic contingency learning. *Journal for the Experimental Analysis of Behavior, 90*, 23–32.

Kutzner, F., Vogel, T., Freytag, P., & Fiedler, K. (2011). Contingency inferences driven by base rates: Valid by sampling. *Judgment and Decision Making, 6*, 211-221.

Landwehr, J. R., & Eckmann, L. (2020). The nature of processing fluency: Amplification versus hedonic marking. *Journal of Experimental Social Psychology*, *90*, 103997.

Lanska, M., Olds, J. M., & Westerman, D. L. (2014). Fluency effects in recognition memory: Are perceptual fluency and conceptual fluency interchangeable? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 40,* 1–11.

Lanska, M., & Westerman, D. (2018). Transfer appropriate fluency: Encoding and retrieval interactions in fluency-based memory illusions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*, 1001-1012.

Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 646-648.

Newell, B. R., & Shanks, D. R. (2007). Recognizing what you like: Examining the relation between the mere-exposure effect and recognition. *European Journal of Cognitive Psychology, 19*, 103-118.

Olds, J. M., & Westerman, D. L. (2012). Can fluency be interpreted as novelty? Retraining the interpretation of fluency in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38*, 653-664.

Oppenheimer, D. M. (2006). Consequences of erudite vernacular utilized irrespective of necessity: Problems with using long words needlessly. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, *20*, 139-156.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness & Cognition, 8*, 338-342.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure:

Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review*, *8*, 364-382.

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science, 9*, 45-48.

Resnick, P., Zeckhauser, R., Swanson, J., & Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics, 9*, 79–101.

Schwarz, N (2015). Metacognition. In M. Mikulincer, P.R. Shaver, E. Borgida, & J. A. Bargh (Eds.), *APA Handbook of Personality and Social Psychology: Attitudes and Social Cognition* (pp. 203-229). Washington, DC: APA

Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology*, *14*, 332-348.

Silva, R. R., Chrobot, N., Newman, E., Schwarz, N., & Topolinski, S. (2017a). Make it short and easy: Username complexity determines trustworthiness above and beyond objective reputation. *Frontiers in Psychology*, 8:2200.

Silva, R. R., Garcia-Marques, T., & Mello, J. (2016). The differential effects of fluency due to repetition and fluency due to color contrast on judgments of truth. *Psychological Research, 80,* 821–837.

Silva, R. R., Garcia-Marques, T., & Reber, R. (2017b). The informative value of type of repetition: Perceptual and conceptual fluency influences on judgments of truth. *Consciousness and Cognition, 51*, 53-67.

Silva, R. R., & Topolinski, S. (2018). My username is IN! The influence of inward vs. outward wandering usernames on judgments of online seller trustworthiness. *Psychology & Marketing*, *35*, 307–319.

Topolinski, S., Bakhtiari, G., & Erle, T. M. (2016). Can I cut the Gordian tnok? The impact of pronounceability, actual solvability, and length on intuitive problem assessments of anagrams. *Cognition*, *146*, 439–452.

Unkelbach, C. (2006). The learned interpretation of cognitive fluency. *Psychological Science, 12*, 339-345.

Unkelbach, C. (2007). Reversing the truth effect: Learning the interpretation of processing fluency in judgments of truth. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 33*, 219-230.

Unkelbach, C., & Greifeneder, R. (2013). *The experience of thinking: How the fluency of mental processes influences cognition and behaviour*. London, England: Psychology Press.

Unkelbach, C., Alves, H., & Koch, A., (2020). Negativity bias, positivity bias, and valence assymetries: Explaining the differential processing of positive and negative information. In B. Gawronski (Ed.), *Advances in Experimental Psychology* (pp. 115-187). Cambridge, MA: Academic Press.

Unkelbach, C., Koch, A., & Alves, H. (2019). The evaluative information ecology: On the frequency and diversity of "good" and "bad". *European Review of Social Psychology, 30*, 216-270.

Vogel, T., & Kutzner, F. (2017). Pseudocontingencies in consumer choice: Preference for prevalent product categories decreases with decreasing set quality. *Journal of Behavioral Decision Making, 30*(5), 1193-1205.

Vogel, T., Silva, R. R., Thomas, A., & Wänke, M. (2020). Truth is in the mind, but beauty is in the eye: Fluency effects are moderated by a match between fluency source and judgment dimension. *Journal of Experimental Psychology: General. 149*, 1587–1596.

Wänke, M. (2013). 10. Almost everything you always wanted to know about ease-of-retrieval effects. In C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking: How feelings from mental processes influences cognition and behaviour* (pp.151–170). London, England: Psychology Press.

Whittlesea, B. W. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition,19*, 1235-1253.

Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: psychophysiological evidence that processing facilitation elicits positive affect. *Journal of Personality and Social Psychology*, *81*, 989-1000.

Winkielman, P., & Schwarz, N. (2001). How pleasant was your childhood? Beliefs about memory shape inferences from experienced difficulty of recall. *Psychological Science, 12*, 176-179.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 189-217). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, 9,* 1–27.

Zajonc, R. B. (2001). Mere Exposure: A Gateway to the Subliminal. *Current Directions in Psychological Science*, *10*, 224–228.

Appendix A

Table 1

Word stimuli used as fluent and disfluent names (ordered alphabetically).

Fluent Names		Disfluent Names	
Batrek	Lapusaw	Agnrdfnduoe	Fgzutkjmnuis
Beketegineb	Lavukef	Aiugnhcfbe	Hguehjkopwmn
Bulo	Lehginitc	Alpwsua	Hhnklcii
Cildenug	Lerbewnut	Atrwgua	Hlucnta
Davo	Madi	Auegnltnvra	Hqjthixeszquj
Ekisepo	Mernigwuki	Ausngdvfre	Idnstiaou
Fabu	Naba	Aztcsu	Ieclhfrne
Fechliren	Napges	Bmjukpolrasd	Iekbgbenete
Fehstoctrab	Neblugrob	Cfguftakkgit	Ietkkge
Fola	Negravanlut	Dtrguoplsmnd	Ignltjeieu
Fudnegrodan	Nekaripo	Eaintvschge	Ilrjpae
Gahbufneic	Nerdet	Eakrtb	Insrdmioe
Galmug	Nilegituje	Eargzpi	Ipskeoe
Gekikite	Nimisedor	Easuanmc	Juwtgfopkjas
Getareigo	Notlinfu	Efhuwopklmns	Kqizutvgpapt
Getelegim	Padasnut	Egncldui	Lkjwuiksguaa
Gima	Pezarig	Eihntclgi	Llopljuzkbqi
Gjawusentu	Poda	Eilegmtge	Luwgthgsnmse
Gubrevnev	Rerubel	Elnbrtweu	Mpupklzcoxik
Hiknilch	Reso	Eoiggrtae	Oacdhblrte
Hira	Rifo	Erdnte	Oaenpkri
Jaliper	Saka	Erlmlkeoi	Oesrsprpo
Jela	Sborntrat	Esgnpa	Otrsdo
Kagu	Scaztu	Etrmknvuu	Pfiumkn
Kelmerilo	Secanamu	Eutngjwsua	Plkopuardghn
Keva	Segu	Evnrbvgue	Qoltdrivsgtu

Table 1 (cont.)

Word stimuli used as fluent and disfluent names (ordered alphabetically).

Fluent Names		Disfluent Names	
Knupfim	Sesu	Evrbiueretjn	Rfeiklxoxjkl
Lapnuk	Sibo	Eztsdo	Sobrntrta
Snacthec	Usanitido	Tohcbtsrcea	Uflvkea
Sordot	Vagfurnesd	Tqefqaycortc	Ugnblrbeo
Sorpesrop	Vanigetesch	Tsrcneha	Unklpa
Tabrelchod	Vijuntreber	Uadnstpa	Uotfnlin
Tagu	Wagratu	Uaglgm	Vlegtiqclapl
Tahnluc	Zemi	Ueignkrwmi	Yqriklbvuogt
Tenkuvmur	Zetdos	Uerbrle	Zhngadruismn

Table 2

Base rates of fluent vs. disfluent seller names and of good vs. bad seller reputation occurrences in the PC fluent name – good reputation condition (n = 80).

	Good Reputation	Bad Reputation	Total
Fluent	41	9	50
Disfluent	9	1	10
Total	50	10	60

Table 3

Base rates of fluent vs. disfluent seller names and of good vs. bad seller reputation occurrences in the PC fluent name – bad reputation condition (n = 84).

	Good Reputation	Bad Reputation	Total
Fluent	9	41	50
Disfluent	1	9	10
Total	10	50	60

Table 4

Base rates of fluent vs. disfluent seller names and of good vs. bad seller reputation occurrences in the PC disfluent name – good reputation condition (n = 81).

	Good Reputation	Bad Reputation	Total
Fluent	9	1	10
Disfluent	41	9	50
Total	50	10	60

Table 5

Base rates of fluent vs. disfluent seller names and of good vs. bad seller reputation occurrences in the PC disfluent name – bad reputation condition (n = 76).

	Good Reputation	Bad Reputation	Total
Fluent	1	9	10
Disfluent	9	41	50
Total	10	50	60

Table 6

Base rates of fluent vs. disfluent seller names and of good vs. bad seller reputation occurrences in the control condition (n = 85).

	Good Reputation	Bad Reputation	Total
Fluent	15	15	30
Disfluent	15	15	30
Total	30	30	60

Appendix B

Follow-up Study: Learning Negative Pseudocontingencies (PCs)

To validate the experimental procedure of our main study, we conducted a follow-up experiment testing whether participants learn the intended negative PCs when an attribute other than fluency was paired with good vs. bad seller reputation. Thus, different from the main experiment, instead of manipulating the base rates of seller username fluency, we manipulated the base rates of two artificial online markets (Batrek vs. Galmug) that the sellers belonged to.

1. Method

1.1 Participants and Design

We recruited 99 participants (M_{age} = 36.14, *SD* = 11.63) through Clickworker online data collection platform. Participants took part voluntarily and received €1.70 (approximately 2.00 USD) as compensation for their participation. Participants were randomly assigned to one of two PC conditions: Positive PCs vs. Negative PCs. Online market (Batrek vs. Galmug) of the sellers rated at the end varied within-participants.

1.2 Materials and Procedure

1.2.1 Online Markets. In this experiment, we created a scenario inducing positive vs. negative PCs between the sellers of two different artificial online markets and good vs. bad seller reputation. As names for the two online markets, we selected two of the easy-to-pronounce stimuli that had been used in the main experiment: Batrek and Galmug.

1.2.2 Reputation. As in the main experiment of this work, we used a five-star rating system to manipulate seller reputation: five golden yellow stars were used to signal sellers with good reputation and one star was used to signal sellers with bad reputation.

1.2.3 Procedure. The experiment was programmed with the online survey platform Qualtrics. The Qualtrics survey administered all instructions, materials, and dependent measures. Participants were users registered at the online data collection platform Clickworker who were sent the link to the experiment by the platform. To take part in the experiment, participants first had to read the informed consent form and agree with the terms of the experiment. There were two sequential tasks in the experiment, the PC induction and the target evaluation, which are described below.

***1.2.3.1 PC induction*.** The instructions in this phase asked participants to imagine that they were visiting a planet in the outer space, in which there were some things very similar to planet Earth. One such thing was that there were two online markets where users could buy and sell to each other and that the names of these online markets were Batrek and Galmug. Participants were informed that they would see slides of different sellers from Batrek and Galmug online markets, and that each slide would present the market a seller belonged to, the identity code of the seller in that market, and the seller's average reputation. Figure 1 presents examples of the seller slides. While the seller identification codes were not relevant for the PC induction, they were necessary in order to create diversity in the slide presentation to make clear to participants each slide pertained to a different seller (given that the two relevant attributes vary in only two absolute values in each attribute – Batrek vs. Galmug and 1 vs. 5 stars). To vary the seller identity codes, we used meaningless combinations of three letters and three digits randomly generated in the website Research Randomizer (www.randomizer.org). Table 1 presents the seller identification codes that were paired with each market in the PC induction phase and in the target evaluation phase. After reading the instructions, participants could start the presentation of 60 seller slides. Participants' task was to read each of the seller slides. The slides were shown one by one, presentation was self-paced, and participants had to click a button placed at the bottom of the screen to go from one seller to the next.

In this experiment, Batrek online market was always the most frequent market to be depicted in the seller slides. Then, depending on the PC condition, the frequency of Good vs. Bad seller reputation varied. In the Positive PC condition, participants saw 50 Batrek sellers and 10 Galmug sellers. Within the markets, most seller reputations were good: 41 good and 9 bad reputations for Batrek sellers, and 9 good and 1 bad reputation for Galmug sellers. A PC illusion should emerge due to the alignment of the skewed base-rates: people tend to associate frequent with frequent, and infrequent with infrequent. Thus, participants should infer that sellers from Batrek market are more likely to have a positive reputation and that sellers from Galmug market are more likely to have a negative reputation. Conversely, in the Negative PC condition, most seller reputations were bad: 41 bad and 9 good reputations for Batrek sellers, and 9 bad and 1 good reputation for Galmug sellers. This should lead to a PC illusion that sellers Batrek are more likely of negative reputation and that sellers from Galmug are more likely to have positive reputation. Order of presentation of the market – reputation pairs was randomized anew for each participant.

Figure 1. Examples of the slides in the PC induction phase.

***1.2.3.2 Target evaluation*.** Next, the target evaluation phase started. The survey informed participants that they were now going to evaluate a different set of Batrek and

Galmug sellers. Their task was to indicate what they thought was the average reputation of each seller on a scale ranging from 1 to 5 stars. Participants judged the reputation of 20 new sellers, half belonging to Batrek market and half to Galmug market (randomly ordered for each participant). In each trial, the online market the seller belonged to and the seller identity code were presented in the screen with the 5-star rating scale below it. Participants clicked on as many stars as they thought the seller reputation was. To proceed, participants had to click a button placed at the bottom of the screen. After the evaluation of the 20 sellers, participants indicated their age and gender ("female", “male”, or “gender variant”). Finally, participants were thanked and debriefed. The experiment lasted approximately 7 minutes.

1.2.4 Predictions. If participants infer positive vs. negative PCs according to their respective conditions, then it is expected that participants in the positive PC condition evaluate Batrek sellers more positively than Galmug sellers, but participants in the negative PC condition should evaluate Batrek sellers more negatively than Galmug sellers.

2. Results

We analyzed participants’ ratings of seller reputation with a PC condition (Positive PC vs. Negative PC) × Online market (Batrek vs. Galmug) mixed ANOVA, with the second factor as repeated measures. The PC condition effect was significant, $F(1, 97) = 97.24$, $p < .001$, $\eta^2_p = .501$, showing that a higher frequency of good reputations led to higher ratings in the target evaluation phase than when low reputation ratings were the most frequent (Positive PC condition, $M = 4.15$, $SE = .12$; Negative PC condition, $M = 2.46$, $SE = .11$) . This main effect suggests that, just as in the main study, participants were sensitive to the manipulation of the frequency of good vs. bad reputation in the PC induction phase and adjusted their evaluations accordingly. The interaction between the two factors was also significant, $F(1, 97) = 10.64$, $p = .002$, $\eta^2_p = .099$, supporting our predictions: in the positive PC condition, sellers belonging to Batrek were evaluated more positively ($M = 4.23$, $SE = .12$) than sellers belonging to Galmug ($M = 4.07$, $SE = .16$), but this pattern reversed in the negative PC condition in which Batrek sellers were evaluated more negatively (Batrek, $M = 2.19$, $SE = .12$; Galmug, $M = 2.73$, $SE = .15$). The effect of Online market was not statistically significant, $F(1, 97) = 3.14$, $p = .079$, $\eta^2_p = .031$ (Batrek, $M = 3.21$, $SE = .08$; Galmug, $M = 3.40$, $SE = .11$).

3. Discussion

Results of this follow-up study suggest that the experimental setting employed in our experiments can promote the learning of negative PCs. Therefore, the absence of a reversal of fluency's positive meaning in the main experiment of this work is more likely to be related to the specificities of fluency experiences than to the specificities of the experimental design that was used to test our hypotheses.

Table 1

Letter and number combinations used as the seller identity codes.

Batrek Sellers		Galmug Sellers	Sellers in Target Evaluation
fk1tho	2kq9vn	6lp1az	10bftn
99cown	bgl9et	3yp91u	n1awc6
2gqpnh	8ja2i5	5q7p4u	051hl1
z8y50w	kqiyev	751niv	4dn0gc
8soti9	25djqj	l29mof	a6x5ny
sg6ln4	z05x7k	az4rc1	vb642d
gy8td2	de07lm	5fqjd1	mipym5
lt1k7x	ff183w	123qig	bphxj8
dprpys	o89xav	35xyif	7yhxb2
2e4b4x	nyr5mr	f0adpl	d6d5km
4qpwkg	hogpeh		wcdvpk
6f6zww	1jbcte		066yjk
mgssdv	a695kw		6lnp5g
os4n7i	2pwu6i		0axxkd
d3d6g9	vrf7t2		m69g0l
77nbzt	mjleui		afm2aq
9g8a9y	16g3n7		2u4666
0he9n3	ms8um6		u1t11j
75ewcb	oig6od		93zhwn
5ksl6g	k8bli1		
sdfb8c	yoatxv		
v7zjb7	2qda5y		
l4ky3c	0oh69q		
fkuo9t	8emboi		
axm4kt	jo3lv3		