Cogent Economics & Finance

ISSN: 2332-2039 (Online) Journal homepage: www.tandfonline.com/journals/oaef20

Pricing contracts on a tradable asset and a non-tradable mean-reverting benchmark: closed-form formulae and applications

Paulo Viegas Carvalho, Carlos Manuel Pinheiro & João Miguel Reis

To cite this article: Paulo Viegas Carvalho, Carlos Manuel Pinheiro & João Miguel Reis (2026) Pricing contracts on a tradable asset and a non-tradable mean-reverting benchmark: closed-form formulae and applications, Cogent Economics & Finance, 14:1, 2740368, DOI: 10.1080/23322039.2026.2740368

To link to this article: https://doi.org/10.1080/23322039.2026.2740368

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

Published online: 08 Oct 2026.

Submit your article to this journal

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=oaef20

COGENT ECONOMICS & FINANCE
2026, VOL. 14, NO. 1, 2740368
https://doi.org/10.1080/23322039.2026.2740368

FINANCIAL ECONOMICS | RESEARCH ARTICLE

OPEN ACCESS Check for updates

Pricing contracts on a tradable asset and a non-tradable mean-reverting benchmark: closed-form formulae and applications

Paulo Viegas Carvalho[a,b,c], Carlos Manuel Pinheiro[b,c,d,e] and João Miguel Reis[b,d]

[a]Department of Finance, Instituto Universitário de Lisboa (ISCTE-IUL), Lisbon, Portugal; [b]Business Research Unit (BRU-IUL), Lisbon, Portugal; [c]ISCAL–Lisbon Accounting and Business School, Lisbon, Portugal; [d]Universidade Europeia, Lisbon, Portugal; [e]Center for Transdisciplinary Development Studies (CETRAD), Vila Real, Portugal

ABSTRACT

We develop a closed-form pricing framework for contracts whose payoff depends on the difference between a tradable underlying asset and a non-tradable mean-reverting benchmark. The framework is designed for applications in which the relevant exercise benchmark or reference variable evolves over time, as in volatility-linked, inflation-linked, weather-related, and real-option settings. The aim is to obtain analytically tractable formulae for an exchange-style payoff involving a tradable underlying asset and a non-tradable mean-reverting benchmark. As shown, this is relevant in sensitive contexts where risks cannot be fully hedged through trading. The framework is applied to real option valuation, examining investment when project value follows geometric Brownian motion and investment cost follows a mean-reverting process. Sensitivity analysis reveals that cost volatility and mean-reversion speed exert distinct and opposing influences on the value of investment options, while the correlation between project value and cost dynamics plays a central role. An empirical calibration using S&P 500 and VIX data shows that benchmark risk premia and mean-reverting benchmark uncertainty can materially affect contract values, particularly when contracts are away from the money.

IMPACT STATEMENT

This study develops closed-form valuation formulae for contracts whose payoff depends on the difference between a tradable asset and a non-tradable, mean-reverting benchmark. This structure is relevant when the payoff reference variable, such as an investment cost, volatility index, inflation measure, or weather indicator, changes over time and cannot be traded directly. The study shows how benchmark volatility, mean-reversion speed, correlation, and market risk premia affect contract values and investment decisions. An application using S&P 500 and VIX data illustrates that accounting for benchmark uncertainty and its risk premium can materially alter valuations, particularly for contracts that are away from the money. By combining analytical tractability with explicit recognition of market incompleteness, the framework provides a useful basis for valuing financial and real-option exposures that conventional fixed-strike models do not capture.

ARTICLE HISTORY
Received 3 June 2026
Revised 3 August 2026
Accepted 23 September 2026

KEYWORDS
Derivatives pricing; mean-reverting benchmark; real options; stochastic investment; investment under uncertainty

SUBJECTS
Finance; Economics; Business, Management and Accounting

1. Introduction

Traditional option pricing models, based on the pioneering work of Black and Scholes (1973) and Merton (1973), assume a payoff based on the difference between an underlying asset and a fixed strike price. While this assumption is appropriate for standard equity options, it becomes restrictive for a growing class of contracts and real-option problems in which the relevant payoff benchmark, reference level, or investment cost evolves over time and cannot be traded directly. Examples include volatility-linked contracts, inflation-linked instruments, weather-related exposures, and investment-timing decisions under stochastic costs.

Managing exposure to such instruments is central to modern, integrated risk frameworks. Financial institutions, energy firms, agricultural companies and sovereign debt managers all face risks that cannot be fully hedged using liquid, tradable assets. The mean-reversion property and non-tradability of many

CONTACT Paulo Viegas Carvalho jpvfc@iscte-iul.pt Department of Finance, Instituto Universitário de Lisboa (ISCTE-IUL), Lisbon, Portugal.

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

financial variables create pricing and hedging challenges that fall outside the scope of conventional derivative pricing models.

Building on existing literature, this paper presents a closed-form financial contract pricing framework. In this framework, one of the underlying assets, henceforth, the stochastic benchmark, is subject to an Ornstein–Uhlenbeck (OU) process embedded within a lognormal structure to ensure strict positivity. There is still an underlying asset that continues to follow geometric Brownian motion (GBM), and the two processes may be correlated. The term "benchmark" is used deliberately: the stochastic process considered here is not a fixed contractual strike in the standard vanilla-option sense, but rather a second payoff-relevant state variable observed at maturity. Accordingly, the framework should not be interpreted as a pricing model for standard vanilla options with fixed strikes, except in the limiting cases in which the benchmark becomes deterministic.

A key theoretical outcome is that the mean-reversion property of the benchmark makes it non-tradable; it is impossible to replicate its payoff at the risk-free rate using a self-financing strategy without creating arbitrage opportunities.

The resulting contract pricing formulae, while distinct, can be linked to the Black-Scholes-Merton (BSM) structure by replacing the classical strike with a risk-neutral expected value. The payoff structure is also closely related to the exchange-option framework of Margrabe (1978), since the contract depends on the difference between two stochastic quantities; the distinctive feature here is that the second process is non-tradable and mean-reverting.

Two main applications are developed. First the framework is applied to real option valuation, specifically to decisions regarding whether to invest immediately or defer when investment costs are stochastic and mean-reverting rather than fixed. Second, it is demonstrated that the framework can be applied directly to the cases of pricing of weather derivatives, volatility index options and inflation-linked instruments, all of which have in common the feature of a non-tradable, mean-reverting exercise benchmark.

In terms of the paper organization, Section 2 develops the model, establishes the non-tradability of the mean-reverting benchmark, and derives closed-form pricing formulae. Section 3 applies the framework to real option valuation under stochastic investment cost. Section 4 explores applications to weather, volatility, and inflation-linked derivatives. Section 5 presents the empirical calibration using S&P 500 and VIX data. Section 6 concludes and outlines directions for future research.

2. The model

2.1. Model formulation

We consider a contract whose value at inception is given by

$$P_0(V_0, I_0; \phi) = e^{-rT}\mathbb{E}\left[(\phi I_T - \phi V_T)^+\right], \tag{1}$$

where $\phi = -1$ is analogous to the case of a call option and $\phi = 1$ to the case of a put option. V_T is the value of the classical underlying asset at maturity T, I_T is the stochastic benchmark at maturity, and r is the risk-free rate. The expected value is taken under the appropriate risk-neutral measure. We work on a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0\le t\le T}, \mathbb{P})$, where $\mathbb{P}$ is the real-world (physical) probability measure, $\mathcal{F}_t = \sigma((\hat{W}_s^V, \hat{W}_s^Y)_{0\le s\le t})$ is the natural filtration generated by the Brownian motions and investment cost processes, and $\mathcal{F}_T = \mathcal{F}$ is the terminal σ-algebra.

The underlying asset V_t follows geometric Brownian motion under the real-world measure $\mathbb{P}$:

$$dV_t = (\mu - \delta)V_t dt + \sigma_V V_t d\hat{W}_t^V, \tag{2}$$

where μ is the expected return of the asset, $\sigma_V > 0$ is the volatility, δ is the cost of carry (dividend yield in the classical case of stock options), and $\hat{W}_t^V$ is a Wiener process under $\mathbb{P}$.

The stochastic benchmark level is defined as

$$I_t = e^{\hat{\theta}+Y_t}, \tag{3}$$

where Y_t follows an OU process

$$dY_t = -\hat{\beta}Y_t\ dt + \sigma_Y\ d\hat{W}_t^Y. \tag{4}$$

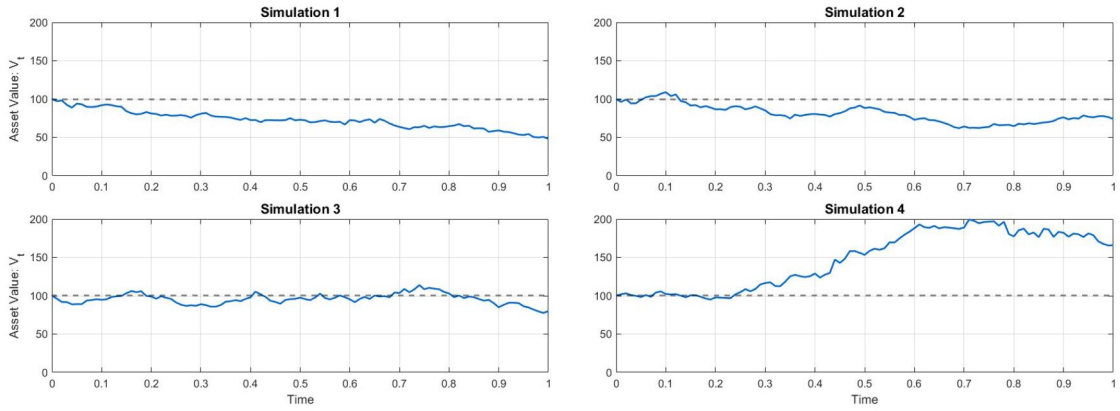


Figure 1. Simulated paths for the GBM process V_t.

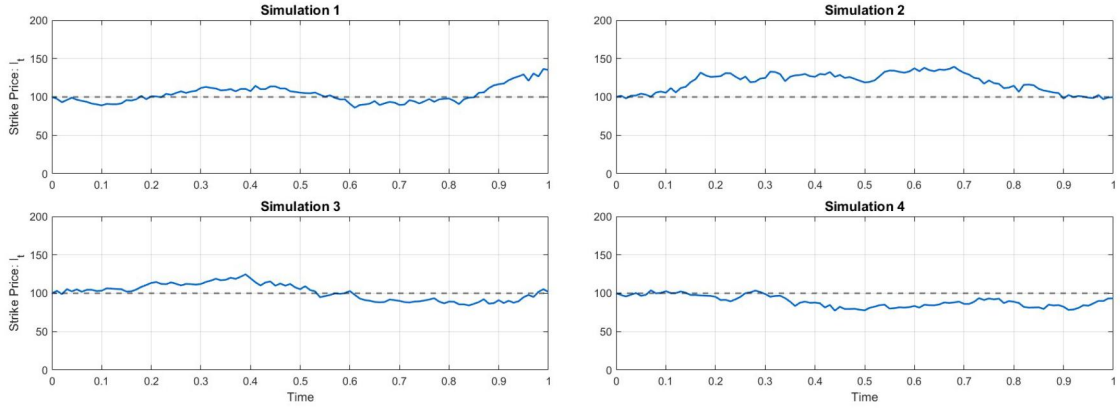


Figure 2. Simulated paths for the mean-reverting process I_t. Starting value matches the long-term average.

The parameter $\hat{\beta}(>0)$ governs the mean reversion speed, $\sigma_Y(>0)$ is the benchmark volatility, $\hat{\theta}$ denotes the long-run log-level of the benchmark and $\hat{W}_t^Y$ is another Wiener process under $\mathbb{P}$. The exponential function ensures that the benchmark remains strictly positive. The two Brownian motions are correlated through their Wiener processes

$$Corr(d\hat{W}_t^V, d\hat{W}_t^Y) = \rho, \tag{5}$$

with $\rho \in [-1, 1]$.

Figures 1 and 2 show the possible paths for the GBM process, V_t and the mean-reverting process I_T, respectively. While the GBM process has deviations, it follows a trend set by μ and δ, whereas the mean-reverting process tends to revert to an average value (which is also the starting value for the parameterization).

2.2. Mean reversion and non-tradability

The mean-reversion property of Y_t has a fundamental consequence: the discounted benchmark process cannot be a $\mathbb{Q}$-Martingale and hence I_t cannot be interpreted as the price of a traded asset. Formally:

Lemma 1. *Let $I_t = e^{\hat{\theta}+Y_t}$ where Y_t follows the OU process*

$$dY_t = -\hat{\beta}Y_t\ dt + \sigma_Y d\hat{W}_t, \quad \hat{\beta} > 0.$$

Then I_t cannot be a traded asset.

Proof. The proof is provided in Appendix A.

Intuitively, when Y_t is above its long-run level, I_t appreciates at a rate below r, enabling arbitrage via a long risk-free position and a short I_t position. Conversely, when Y_t is below its long-run level, I_t appreciates above r, again enabling arbitrage. Tradability requires $\hat{\beta} = 0$, which reduces the OU process to a random

walk and thus contradicts the mean-reversion assumption. This restriction on tradability has significant implications: pure replication of instruments with mean-reverting exercise thresholds is infeasible, and pricing requires the specification of a market price of risk. In other words, in this case there exists no self-financing trading strategy that can be used to replicate this with a risk-free asset and a tradable risky asset.

2.3. Risk-neutral measure

We use risk-neutral valuation to determine the payoff. For V_t, there are two different cases, depending on whether the asset is tradable. For I_t, given its non-tradability, established in Lemma 1, only the non-tradable case is presented.

Assumption 1. *When V_t represents a tradable asset, the risk-neutral dynamics follow the diffusion:*

$$dV_t = (r-\delta)V_t\ dt + \sigma_V V_t\ dW_t^V, \tag{6}$$

under the risk-neutral measure $\mathbb{Q}$. The market price of risk λ is uniquely determined by no-arbitrage: $\lambda = (\mu - \delta - r)/\sigma_V$.

This can be seen by noting that the market price of risk λ is uniquely determined by no-arbitrage. Thus, through Equation (2), with respect to the physical measure $\mathbb{P}$, the process:

$$W_t^V = \hat{W}_t^V + \int_0^t \lambda\ ds$$

is a standard Brownian motion under the risk-neutral measure $\mathbb{Q}$.

Substituting $d\hat{W}_t^V = dW_t^V - \lambda\ dt$ into the dynamics, and considering the Martingale Property, gives the market price of risk

$$\lambda = \frac{\mu - \delta - r}{\sigma_V}.$$

For the tradable V_t case, the change of measure that fixes the risk-neutral drift at V_t is $(r-\delta)$; no additional exogenous market price of risk remains to be specified.

Assumption 2. *When V_t is a non-tradable asset, it follows the diffusion:*

$$dV_t = (r - \delta + \lambda\sigma_V)V_t\ dt + \sigma_V V_t\ dW_t^V, \tag{7}$$

under the risk-neutral measure $\mathbb{Q}$.

Here, where V_t is non-tradable, a market price of risk λ must be specified exogenously. This is noted by regarding that the market price of risk, λ, cannot be determined by arbitrage. By Girsanov's theorem, with market price of risk λ:

$$W_t^V = \hat{W}_t^V + \int_0^t \lambda\ ds$$

is a standard Brownian motion under the risk-neutral measure $\mathbb{Q}$.

Substituting $d\hat{W}_t^V = dW_t^V - \lambda\ dt$ into the physical dynamics leads to the expression in the assumption. Thus, the risk-neutral drift contemplates both the dividend adjustment $(r-\delta)$ and the risk premium $\lambda\sigma_V$.

We assume for the remainder of the paper $\lambda = 0$ as a simplifying assumption which selects one specific martingale measure out of infinitely many.[1] Setting $\lambda = 0$ selects one specific martingale measure and is a common simplification in the incomplete markets literature (see, e.g. Jaimungal et al. (2013)). The implications of setting $\lambda > 0$ are elaborated in Section 3, in the context of real options.

So, in the remainder of this paper, for the tradable and non-tradable asset cases, we refer to the model presented by V_T with Equation (6) (which matches Equation (7) for the case where $\lambda = 0$). Thus, in this case, the expression for Assumption 2 matches the one from Assumption 1.

Assumption 3. *The non-tradable mean-reverting process I_t is given by*

$$I_t = \exp\{\theta + Y_t\}, \tag{8}$$

$$dY_t = -\beta Y_t\ dt + \sigma_Y\ dW_t^Y, \tag{9}$$

in the risk-neutral measure $\mathbb{Q}$, with the parameters

$$\beta = \hat{\beta} - b_y, \tag{10}$$

$$\theta = \hat{\theta} + \frac{a_y}{\beta}, \tag{11}$$

where a_y and b_y are constants.

Arising from assuming an affine market price of risk that depends on the state Y_t:

$$\lambda_t^Y = \frac{a_y + b_y Y_t}{\sigma_Y}, \tag{12}$$

for the physical measure $\mathbb{P}$ terms presented in Equations (3) and (4).

By Girsanov's theorem:

$$W_t^Y = \hat{W}_t^Y + \int_0^t \lambda_s^Y \, ds$$

is a Brownian motion under $\mathbb{Q}$. In addition, substituting $d\hat{W}_t^Y = dW_t^Y - \lambda_t^Y \, dt$ yields Equation (10) and Equation (11).

For applications where the benchmark is a non-tradable index (VIX, CPI, temperature), λ_t^Y would need to be estimated from market prices. While, for the purpose of understanding the qualitative effects of variables such as mean reversion and correlation, setting $\lambda_t^Y = 0$ is a natural approach, we explore the effects of an estimation with market prices in the contract pricing of Section 5.

Jaimungal et al. (2013) provide further details on the impact of different choices for parameters a_y and b_y. Also following that paper, henceforth, this work focuses on the risk-neutral measure.

The correlation structure between Brownian motions is preserved, so the correlation between W_t^V and W_t^Y under $\mathbb{Q}$ remains ρ, the same as between $\hat{W}_t^V$ and $\hat{W}_t^Y$ under $\mathbb{P}$.

2.4. Closed-form pricing formulae

We derive closed-form expressions for the contracts under $\phi \in \{-1; 1\}$. We begin by presenting the risk-neutral solutions to the diffusion processes. For Equations (6)-(7) (with $\lambda = 0$), the standard solution to the SDE is:

$$V_T = V_0 \exp\left\{\left(r - \delta - \frac{1}{2}\sigma_V^2\right)T + \sigma_V W_T^V\right\}.$$

As for the expected value in relation for the filtration, it is also known:

$$\mathbb{E}[V_T \mid \mathcal{F}_t] = V_t e^{(r-\delta)(T-t)}. \tag{13}$$

The OU process from Equations (8)-(9), starting from time 0, has the known explicit solution:

$$Y_T = e^{-\beta T} Y_0 + \sigma_Y \int_0^T e^{-\beta(T-u)} dW_u^Y, \tag{14}$$

and I_T:

$$\ln I_T = \theta + Y_T \sim N\left(\theta + e^{-\beta T} Y_0, \frac{\sigma_Y^2}{2\beta}(1 - e^{-2\beta T})\right),$$

where Y_0 is the starting value for the Y_t process.

In relation to the filtration, using the moment-generating function of a normal distribution, it yields that

$$\mathbb{E}[I_T \mid \mathcal{F}_t] = \exp\left\{\theta + e^{-\beta(T-t)} Y_t + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta(T-t)})\right\}. \tag{15}$$

We start with the case of the contract where $\phi = -1$. Then, following from Equation (1), we find the value at time 0.

Theorem 1. *The value of a contract where $\phi = -1$ with maturity T, underlying asset V_t, following Assumption 1 or Assumption 2 (case with $\lambda = 0$) and Assumption 3, is given by:*

$$P(V_0, I_0; -1) = e^{-\delta T} V_0 \Phi(d_+) - e^{-rT} \mathbb{E}^{\mathbb{Q}}[I_T] \Phi(d_-),$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function,

$$\mathbb{E}^{\mathbb{Q}}[I_T] = \exp\left\{\theta + e^{-\beta T} Y_0 + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta T})\right\},$$

$$d_\pm = \frac{\ln\left(\frac{V_0 e^{(r-\delta)T}}{\mathbb{E}^{\mathbb{Q}}[I_T]}\right) \pm \frac{1}{2}\tilde{\sigma}^2}{\tilde{\sigma}},$$

and

$$\tilde{\sigma}^2 = \sigma_V^2 T + \frac{\sigma_Y^2}{2\beta}(1 - e^{-2\beta T}) - \frac{2\rho\sigma_V\sigma_Y}{\beta}(1 - e^{-\beta T}).$$

Proof. This proof is provided in Appendix B.

Theorem 2. *The value of a contract where $\phi = 1$ with maturity T, underlying asset V_t, following Assumption 1 or Assumption 2 (case with $\lambda = 0$) and Assumption 3, is given by:*

$$P(V_0, I_0; 1) = e^{-rT} \mathbb{E}^{\mathbb{Q}}[I_T] \Phi(-d_-) - e^{-\delta T} V_0 \Phi(-d_+),$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function,

$$\mathbb{E}^{\mathbb{Q}}[I_T] = \exp\left\{\theta + e^{-\beta T} Y_0 + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta T})\right\},$$

$$d_\pm = \frac{\ln\left(\frac{V_0 e^{(r-\delta)T}}{\mathbb{E}^{\mathbb{Q}}[I_T]}\right) \pm \frac{1}{2}\tilde{\sigma}^2}{\tilde{\sigma}},$$

and

$$\tilde{\sigma}^2 = \sigma_V^2 T + \frac{\sigma_Y^2}{2\beta}(1 - e^{-2\beta T}) - \frac{2\rho\sigma_V\sigma_Y}{\beta}(1 - e^{-\beta T}).$$

Proof. The proof is similar to that of Theorem 1 and therefore omitted. ■

In both theorems, $\tilde{\sigma}^2$ represents what is often called effective variance.

As mentioned in Remarks 1 and 2, the obtained contracts can be reduced to the BSM call and put options. Only under those circumstances do we recover the original BSM model.

Remark 1. *The contracts priced in Theorems 1 and 2 are not standard vanilla options with fixed contractual strikes. They are European-style claims whose payoff depends on the difference between the underlying process V_t and the stochastic benchmark I_t. The standard fixed-strike Black–Scholes–Merton formula is recovered when benchmark uncertainty vanishes, for example when $\sigma_Y = 0$, or when the benchmark becomes effectively deterministic.*

Remark 2. *When $\sigma_Y = 0$, the benchmark has no stochastic component and the formulae collapse to the standard Black–Scholes–Merton expressions with dividends. Similarly, in the limiting case $\beta \to \infty$, mean reversion becomes instantaneous and the benchmark approaches a fixed level. These cases clarify that the standard fixed-strike model is obtained only when the benchmark is deterministic or effectively deterministic.*

In Theorems 1 and 2, the constant correlation parameter ρ, as defined in Equation (5), is a simplifying assumption. This assumption is adopted for tractability, although correlations are known to vary over

time and across regimes, especially during volatile or downturn periods, as documented by Ang and Bekaert (2002) and Engle (2002). In the present framework, correlation enters the price through the effective variance term $\tilde{\sigma}^2$, specifically through the covariance component:

$$-\frac{2\rho\sigma_V\sigma_Y}{\beta}(1-e^{-\beta T}).$$

If correlation is deterministic and time-varying, denoted by $\rho(t)$, the closed-form structure is preserved by replacing this covariance component with:

$$-2\sigma_V\sigma_Y\int_0^T e^{-\beta(T-s)}\rho(s)ds.$$

Thus, the constant-correlation case can be interpreted as using a maturity-weighted average correlation over the life of the contract. Fully stochastic or state-dependent correlation would generally require additional modelling assumptions and may break the simple closed-form expression, although its pricing effect operates through the integrated covariance term. Approximation approaches for exchange-style claims with stochastic volatility and stochastic correlation are discussed by Antonelli et al. (2010) and Kim and Park (2017). We retain the constant-correlation specification to preserve analytical tractability and examine its empirical stability in Section 5, where the S&P 500 is used as the underlying process and the VIX as the stochastic benchmark.

Remark 3. *The underlying asset is modelled as a GBM, a choice made for analytical convenience. Equity indices are well known to exhibit fat tails, stochastic volatility and regime shifts, as the descriptive statistics of Section 5 confirm. Alternative specifications are available, including* Cox (1975)*,* Cox and Ross (1976)*,* Heston (1993) *and the regime-switching dynamics documented by* Ang and Bekaert (2002)*. The assumption is nonetheless confined to a single step of the argument: neither Lemma 1 nor the change of measure of Appendix A depend on how the underlying evolves – the specification matters only in determining the distribution of the ratio of the two terminal values. Alternative dynamics are therefore accommodated by replacing the dynamics of V_t alone. Jump and regime-switching specifications can preserve the closed form as a probability-weighted average of the present expressions, while specifications such as that of* Heston (1993) *require a numerical treatment. Since the effects documented here operate through the expected benchmark level and the effective variance, both governed by the benchmark process, richer dynamics for the underlying would be expected to alter contract values, most noticeably away from the money, while preserving the ordering of prices across specifications and the direction of the sensitivities to benchmark volatility, reversion speed and correlation.*

3. Real option evaluation

Real option theory provides a rigorous framework for valuing managerial flexibility under uncertainty. The baseline setting asks whether an investment should be undertaken to acquire a project of uncertain value. As pioneering approaches in using option formulations to analyse investments, Brennan and Schwartz (1985) and McDonald and Siegel (1986) evaluate decisions on projects subject to stochastic variables. Considering sequential investments, Majd and Pindyck (1987) also rely on option formulations. Trigeorgis (1993) explores the flexibility provided by real options in overcoming limitations of the traditional net present value approach, while Dixit and Pindyck (1994) provide a comprehensive treatment of investment under uncertainty. Schwartz (2004) applies the real option formulation to R&D and patents, where the investor must decide whether to advance a project with uncertain costs, also contemplating the possibility of abandonment. Décamps et al. (2006) explore the choice among mutually exclusive projects.

In our case, we specifically address uncertainty arising from the investment cost that must be incurred to start the project. In Pindyck (1993), uncertainty stems from two sources, one of which concerns input cost uncertainty. In Sarkar (2003), uncertainty comes from the operational cost, modelled by a mean-reverting variable. Jaimungal et al. (2013) also consider mean-reverting processes in the cost, while additionally modelling the value of the investment project as mean-reverting. In Mac Cawley et al. (2020),

Table 1. Model parameters and baseline values.

Parameter	Variable	Values
Starting value for the project	V_0	$[0, 30]$
Risk-free interest rate	r	5%
Project value cost of carry	δ	3%
Project value volatility	σ_V	30%
Long-term investment cost parameter	θ	$[\ln(10^{-6}), \ln(30)]$
Starting value for Y_t	Y_0	0
Mean reversion parameter	β	1
Investment cost volatility	σ_Y	20%
Correlation between processes	ρ	-30%
Maturity	T	[0, 5]

output commodity prices are modelled as mean-reverting processes, driving the timing of investment decisions. Alexander and Chen (2021) treat mean reversion as one dimension of risk, examining how sensitive real option values are to the choice of price process. Briest et al. (2022) fit an exponential Ornstein–Uhlenbeck process to empirically derived power plant revenues and use it as the state variable in their real options model. The present paper stands out by modelling project value as a geometric Brownian motion and investment cost as an OU process.

The studied investment decision presented is as follows. An agent may invest immediately, paying cost I_0 to receive project value V_0, yielding net payoff $(V_0 - I_0)^+$, or defer investment to a horizon T, at which point the decision is revisited. The option to defer is typically valued by a European call option on V_t with a given strike. Here, it is so with the case of the developed contract where $\phi = -1$, using the stochastic benchmark I_t to represent the investment cost, following the formula from Equation (1) to evaluate the value of investing in the future. Crucially, we consider the cost of investment a value which, while stochastic, cannot be traded. Thus, in the possibility to delay the investment, the value to the investor is given by $\max((V_0 - I_0)^+, P(V_0, I_0; -1))$.

We explore the model with the parameters presented in Table 1. Note that setting $Y_0 = 0$ in Equation (14) implies that the initial value I_0 equals its long-run level e^{θ}.

This is mapped by Figure 3, displaying the immediate investment payoff and the option-to-defer value across the (V_0, I_0) parameter space, for $T = 5$. Combinations of high V_0 and low I_0 favour immediate investment – waiting is not expected to lead to a more advantageous situation. As I_0 rises and V_0 falls, investing immediately becomes suboptimal. In many out-of-the-money cases, there is, as expected, value in waiting, since the stochastic evolution of both variables may allow profitable investment at a later date. For instance, when the project is worth 25, and the investment cost is 30, the agent should not invest immediately, while the option of investing in five years retains positive value. Conversely, on some occasions, where the option to invest is initially at-the-money, there is value in waiting. The plotted exercise boundary signals the combinations of V_0 and I_0 at which the agent is indifferent between investing immediately and deferring to T.

3.1. Sensitivity analysis

We can also examine how each variable affects the value of the option to invest in the future in relation to various model variables.

Figure 4 examines the value of investing at T for different starting values of the investment cost, with $\theta = 15$ and $V_0 = 30$. For early maturities, a lower initial cost I_0 yields a higher option value. That is, when the time comes for the option to be exercised, a lower starting cost translates directly into a higher net payoff. For higher maturities, the option values for the three starting levels of I_0 converge, as the mean-reverting process steers investment costs toward the long-run average $\theta = 15$, progressively diminishing the impact of the initial level. This is reflected analytically in the $e^{-\beta T} Y_0$ term in the expression for $\mathbb{E}^{\mathbb{Q}}[I_T]$, where the weight on Y_0 decays as T grows.

Figure 5 explores the effects of the project-value parameters. As in the standard Black–Scholes–Merton model, increases in project value volatility σ_V raise the option to invest, since greater dispersion steers project values toward more extreme realizations. Crucially, in an at-the-money scenario, downside cases already yield zero payoff, so only the upside matters. Higher σ_V also reduces the rate

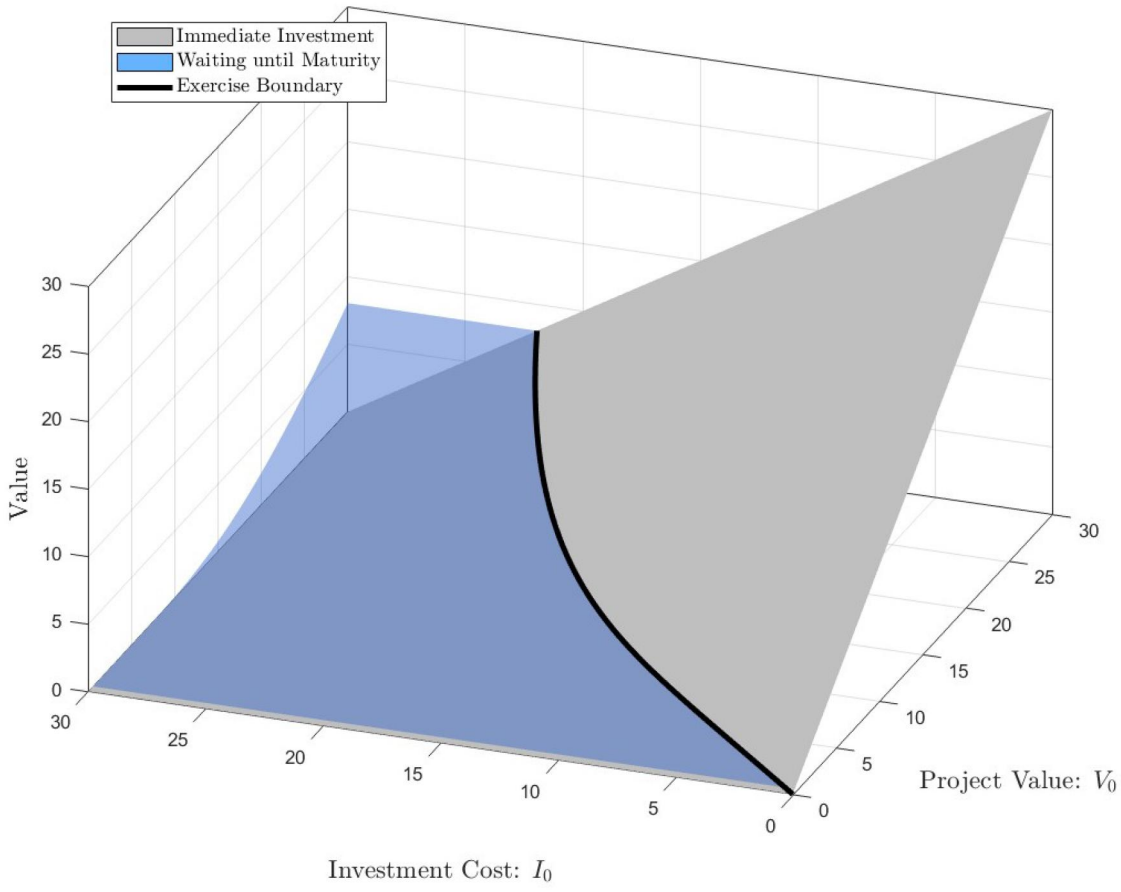


Figure 3. The value of investing immediately versus waiting (the European option to defer). Table 1 parameters and $T = 5$.

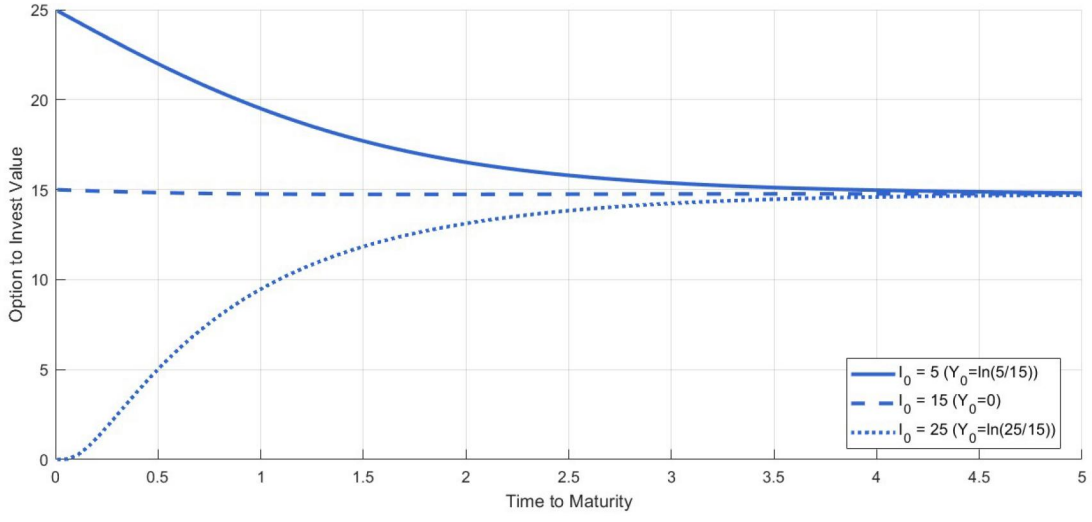


Figure 4. The value of the option to invest at maturity. Table 1 parameters, while $\theta = 15$, $I_0 = [5, 15, 25]$, $V_0 = 30$, and $0 < T \leq 5$.

of time decay (Theta in financial option terms), so that higher-volatility options are more valuable as maturity approaches. Increases in the cost of carry δ, by contrast, unambiguously reduce the option to invest. If the value of the project erodes over time, the option to invest in the future becomes less attractive, in direct analogy with the effect of dividend yield on call options in the Black–Scholes–Merton framework.

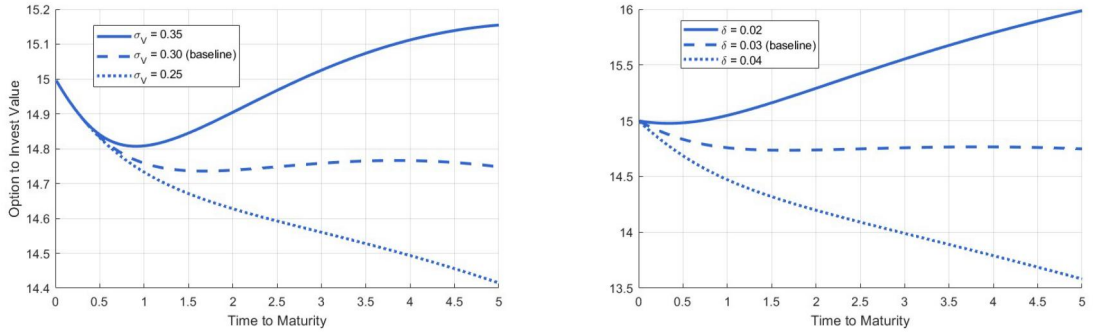


Figure 5. The value of the option to invest at maturity. Table 1 parameters, with $\sigma_V = [0.25, 0.30, 0.35]$ on the left-hand side, with $\delta = [0.02, 0.03, 0.04]$ on the right-hand side, $I_0 = 15$, $V_0 = 30$ and $0 < T \leq 5$.

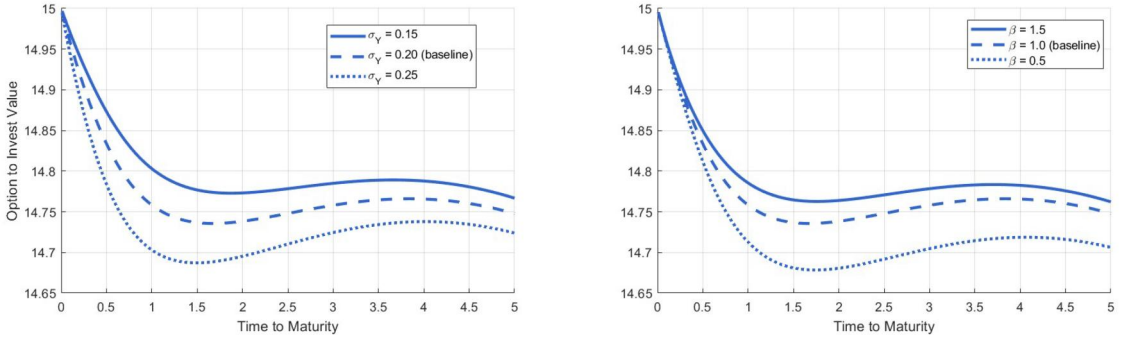


Figure 6. The value of the option to invest at maturity. Table 1 parameters, with $\sigma_Y = [0.20, 0.25, 0.30]$ on the left-hand side, with $\beta = [0.5, 1, 1.5]$ on the right-hand side, $I_0 = 15$, $V_0 = 30$, $\rho = 0$ and $0 < T \leq 5$.

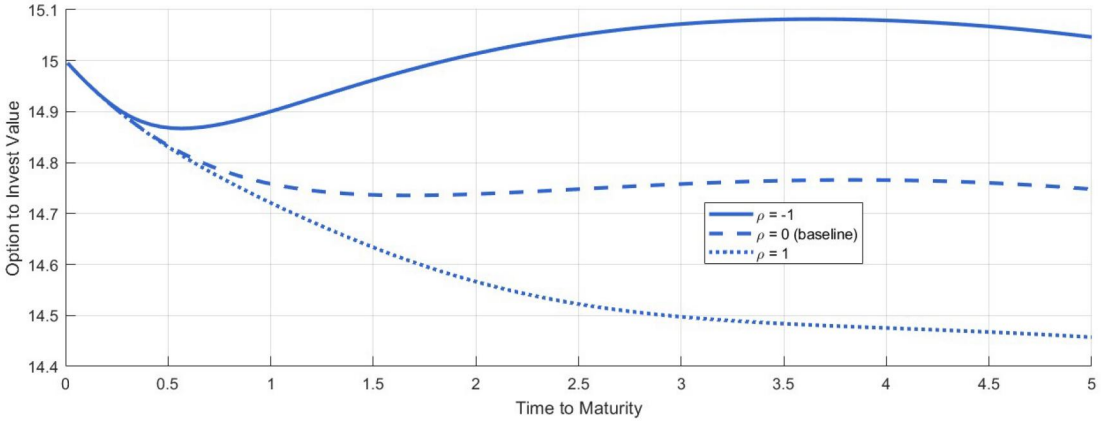


Figure 7. The value of the option to invest at maturity. Table 1 parameters, with $\rho = [-1, 0, 1]$, $I_0 = 15$, $V_0 = 30$ and $0 < T \leq 5$.

Figure 6 considers the sensitivity to the investment cost process parameters, setting $\rho = 0$. On the left-hand side, higher cost volatility σ_Y is shown to reduce the option to defer. This can be interpreted as holding a short position in the investment cost. In this case, right-tail realizations of I_t eliminate gains in project value, while left-tail realizations are bounded by zero – effectively imposing a short investment cost volatility position on the option to invest in the future. On the right-hand side, higher mean-reversion speed β increases the option value. The speed of mean reversion in the OU process acts as a counterweight to cost volatility, steering the investment cost toward its long-run level e^{θ} and reducing the probability of extreme realizations at maturity. The effect of β is therefore the reverse of σ_Y.

Figure 7 analyses the role of correlation. Higher (positive) correlations reduce the value of the option to invest in the future, while lower (negative) correlations increase it. Intuitively, when correlation is

strongly negative, a sharp rise in project value is accompanied by a fall in investment cost, amplifying the net payoff. Conversely, a sharp fall in project value is accompanied by a rise in investment cost, but since losses are bounded by zero, the option holder is shielded from the downside. When correlation is close to unity, the situation is reversed as high project values are offset by high costs, and situations where both are low are bounded by zero. This highlights the importance of not only considering how an investment's value and its associated cost behave individually but also their relation. Moving in synchrony or asynchrony has opposite implications for the project evaluation.

While not explored numerically, the role of the market price of risk is addressed over both diffusions. On the project value side, if V_t were assigned a market price of risk λ, its risk-neutral drift would become $r-\delta+\lambda\sigma_V$ (see Equation (7)). This would lead the deferral option to increase with a positive premium, as $\lambda>0$ raises the option to invest; conversely, a negative premium would lower it.. Crucially, this sign ambiguity arises only when the project value is itself treated as non-tradable. When V_t is tradable (see Assumption 1), λ is pinned by no-arbitrage at $(\mu-\delta-r)/\sigma_V$ and $\lambda=0$ is the standard normalisation that leaves the drift at $r-\delta$.

The assumption is more consequential on the cost side. Setting $\lambda^Y=0$ on the mean-reverting investment would force, through Equations (10)–(11), the risk-neutral reversion speed and long-run level to coincide with their physical counterparts ($\beta=\hat{\beta}$ and $\theta=\hat{\theta}$). Since these parameters enter both the expected cost $\mathbb{E}^{\mathbb{Q}}[I_T]$ and the effective variance $\tilde{\sigma}^2$ of Theorem 1, suppressing the cost's market price of risk feeds directly into the deferral value and hence the invest-versus-wait conclusion.

The direction of the effect follows the same mechanism as the sensitivity analysis of the figures above: a premium that slows risk-neutral reversion ($\beta<\hat{\beta}$) enlarges $\tilde{\sigma}^2$ and raises the option to defer, exactly as a lower reversion speed does in Figure 6 (left-hand side panel), whereas a premium that speeds reversion ($\beta>\hat{\beta}$) compresses $\tilde{\sigma}^2$ and lowers it. The magnitude of the resulting mispricing is therefore an empirical question. It is regarded in more detail for the case of the contract calibration for the S&P 500 and the VIX in Section 5.

4. Non-tradable benchmark contract applications

We will explore various types of contracts to evaluate whether the developed setup can be applied. After that, we will evaluate the contracts using the respective formulae.

4.1. Volatility derivatives

Contracts on volatility indices – most prominently VIX options – represent an application domain for our formulation.[2] The VIX is a non-tradable index, meaning that it cannot be purchased directly in the spot market, though futures and options written on it are actively traded. Grünbichler and Longstaff (1996) and Detemple and Osakwe (2000) establish foundational pricing frameworks for volatility options using mean-reverting processes; Mencía and Sentana (2013) provide a comprehensive empirical review.

In the proposed framework, the VIX acts as the stochastic benchmark and a broad equity index (such as the S&P 500) serves as the underlying asset. The OU structure for the VIX is empirically well supported, reflecting the well-documented tendency of implied volatility to revert to long-run means following periods of market stress. The well-studied leverage effect – the negative correlation between equity returns and changes in implied volatility, formalized by Christie (1982) and confirmed empirically by Fleming et al. (1995) and Whaley (2000) – is captured directly in the model through the negative correlation parameter ρ.

When a volatility index (e.g. the VIX) replaces the fixed reference level and a stock (or stock index) replaces the underlying asset in a spread-style payoff, the contract can be priced directly with our formulae. Crucially, the volatility index is set to follow a mean-reverting process while the correlation parameter sets a potentially negative relationship between the two processes.

4.2. Inflation-linked derivatives

Inflation-linked derivatives – caps, floors, and options referenced to consumer price indices – constitute the second domain of application with growing importance for liability-driven investment and

macroeconomic risk management. The Consumer Price Index (CPI) is, like the VIX, a process whose dynamics are often modelled as mean-reverting toward a long-run target consistent with central bank objectives. [3] The structural analogies between interest rates and price indices enabled the foundational work of Hull and White (1990) and Heath et al. (1992) on interest rate derivatives to provide the conceptual scaffolding upon which inflation derivative pricing was subsequently built. Jarrow and Yildirim (2003) and Mercurio (2005) develop foundational models for inflation derivative pricing by treating instantaneous inflation as a mean-reverting process in line with interest-rate frameworks.

In the proposed model, the CPI index takes the place of the stochastic benchmark, while equity claims serve as the typical underlying asset. Crucially, the inflation index is modelled as following a mean-reverting process, reflecting the anchoring role of central bank inflation targets, while the correlation parameter captures the potentially negative relationship between inflation and the returns on the equity product.

The correlation parameter captures the short-run negative relationship between inflation and equity returns documented by Fama and Schwert (1977), and Fama (1981), as well as the long-run positive relationship identified by Boudoukh and Richardson (1993). The appropriate sign and magnitude of ρ therefore depend on the investment horizon of the derivative, which provides a natural calibration anchor.

4.3. Weather derivatives

Weather derivatives represent a notable class of incomplete-market instruments. Temperature, precipitation, and related meteorological variables are not tradable assets, yet they generate economically significant exposures for energy utilities and agricultural firms. Alaton et al. (2002), Brody et al. (2002), and Benth and Šaltytė-Benth (2005) establish that temperature dynamics are well-described by mean-reverting processes, motivating a rich literature on the pricing of weather-indexed contracts.

The framework accommodates weather derivatives in which the benchmark is a weather variable modeled as a mean-reverting OU process, while the underlying asset follows a GBM process correlated with weather conditions. This configuration – a tradable GBM underlying with a non-tradable OU benchmark – maps directly onto the developed contract structure.

As for the correlation parameter ρ, it captures the empirically documented relationship between weather realizations and equity returns: Hong et al. (2019) find that food company returns are negatively affected by droughts. Meanwhile, Cheng et al. (2022) identify severe drought effects across seven out of 15 Chinese industry sectors[4], and Addoum et al. (2023) link extreme temperatures to earnings declines in 40% of US industries. In all such cases, a negative correlation between the weather benchmark (in the classic position of the strike) and equity value (underlying) has an impact on the option contracts explored here.

5. Empirical calibration

This Section estimates the dynamics from data, documents their stability, recovers the market price of risk from a traded forward curve rather than assuming it away, and quantifies the pricing consequences of the mean-reverting, non-tradable structure of the payoff of Equation (1).

We use a particular case for this flagship application of our framework: the underlying V_t is a broad equity index, the S&P 500, and the non-tradable benchmark I_t is the VIX, as introduced in Subsection 4.1.

5.1. Data

Data for the two processes are obtained from the daily Yahoo Finance series,[5] spanning 2015–2024. All series are aligned on common trading dates, providing 2515 daily observations and 2514 log-returns. Two series for the S&P 500 are obtained, one considering the returns with dividends and the other without, allowing the deduction of the dividend yield (cost-of-carry), thus obtaining the cost of carry for the pricing model parameter δ.

Table 2. Descriptive statistics, daily 2015–2024 data. S&P return is the daily equity log-return in percent. VIX stands for the benchmark level, while ln VIX is the log-benchmark.

Statistic	S&P return (%)	VIX level	ln vIX
Observations	2514	2515	2515
Mean	0.042	18.269	2.844
Std. deviation	1.127	7.284	0.333
Minimum	−12.77	9.14	2.21
Maximum	8.97	82.69	4.42
Skewness	−0.81	2.60	0.76
Excess kurtosis	15.72	12.92	0.86
AR(1)	−0.135	0.965	0.971

Since the price index excludes dividends while the total-return index reinvests them, their log-increment gap is the dividend accrual. Averaging it over the whole sample and annualising yields the estimated value $\hat{\delta} \approx 1.82\%$.

In order to obtain a risk-free rate, we used the FRED[6] series for the three-month U.S. Treasury yield over 2015–2024 ($n = 2515$ daily observations) and used its sample average, resulting in the estimated risk-free rate $\hat{r} \approx 1.83\%$.

The risk-neutral dynamics are identified from the VIX futures term structure: we use the CBOE VX (Volatility Index) monthly futures settlement curve from CBOE's delayed-quote page.

For the two processes, we construct the equity log-return $R_t^V = \ln(V_t/V_{t-1})$ and the log-benchmark $\ln I_t = \ln \mathrm{VIX}_t$. Table 2 reports the descriptive statistics.

The futures snapshot is taken as of a recent date (3 July 2026) [7], keeping the standard monthly contracts and discarding the weekly contracts and any zero-volume (untraded) lines. This leaves a liquid strip of seven maturities, from the front contract to roughly seven months out, whose settlement prices are taken as the market forwards. [8] The benchmark spot close on the valuation date ($I_0 = 16.15$) is retained alongside the strip: it anchors the risk-neutral fit. The estimation of market price of risk parameters is developed in Appendix C.

5.2. Physical measure parameters

Under the physical measure, the benchmark follows the log-normal, mean-reverting specification of Equations (3)–(4). Hence $X_t \equiv \ln I_t$ is an Ornstein–Uhlenbeck process reverting to $\hat{\theta}$, and over a step Δ its transition density is Gaussian and exact, requiring no discretization approximation at the daily sampling frequency, giving an AR(1) representation that avoids the Euler discretisation bias (explored in Iacus (2008)). Thus:

$$X_{t+\Delta} = c + \varphi\, X_t + \varepsilon_{t+\Delta}, \qquad \varphi = e^{-\hat{\beta}\Delta}, \quad c = \hat{\theta}(1-\varphi), \quad \varepsilon_{t+\Delta} \sim \mathcal{N}\left(0, \frac{\sigma_Y^2}{2\hat{\beta}}(1-\varphi^2)\right). \tag{16}$$

The structural parameters are recovered from the ordinary least squares fit of Equation (16) – which yields the estimated autoregressive coefficient $\hat{\varphi}$ (the sample counterpart of the population coefficient φ) and the residual variance $\hat{\sigma}_\varepsilon^2$. Thus:

$$\hat{\beta} = -\frac{\ln \hat{\varphi}}{\Delta}, \qquad \hat{\theta} = \overline{X}, \qquad \hat{\sigma}_Y = \sqrt{\hat{\sigma}_\varepsilon^2 \, \frac{2\,\hat{\beta}}{1-\hat{\varphi}^2}}, \qquad \text{half-life} = \frac{\ln 2}{\hat{\beta}}, \tag{17}$$

where the long-run level is estimated by the sample mean $\overline{X}$, which avoids the ratio $c/(1-\hat{\varphi})$ and its instability when $\hat{\varphi}$ is near unity.

Because the AR(1) coefficient is downward biased in finite samples, the mean-reversion speed is biased upward; the bias is especially pronounced for mean-reverting diffusions, as documented by Yu (2012). We therefore report the raw estimate alongside the bias-corrected value, following Kendall (1954): $\hat{\varphi}_c = \hat{\varphi} + (1+3\hat{\varphi})/n$.

The measure-invariant diffusion parameters are estimated directly. With equity log-returns $R_{t+\Delta}^V = \ln(V_{t+\Delta}/V_t)$:

$$\hat{\sigma}_V = \frac{\mathrm{sd}(R^V)}{\sqrt{\Delta}}, \qquad \hat{\rho} = \mathrm{corr}(R_{t+\Delta}^V, \hat{\varepsilon}_{t+\Delta}), \tag{18}$$

Table 3. Physical-measure ($\mathbb{P}$) daily estimates for the period 2015–2024 ($n = 2515$). standard errors for $\hat{\beta}, \hat{\theta}, \hat{\sigma}_Y$ are calculated using the Delta-method; $\hat{\sigma}_V$ and $\hat{\rho}$ are sample moments.

Parameter	Symbol	Estimate	Std. error	*AR(1) fit diagnostics*	
Reversion speed (raw)	$\hat{\beta}$	7.40	1.23	Adjusted R^2	0.9429
Reversion speed (bias-corr.)	$\hat{\beta}_c$	7.00	1.23	Log-likelihood	2794.2
Half-life	$\ln 2/\hat{\beta}_c$	1.19 mo.	—	AIC / BIC	−5582.5 / −5565.0
Long-run benchmark level	$\hat{\theta}$	2.844	0.055	*Random-walk null* $\varphi = 1$	
(level $e^{\hat{\theta}}$)		17.19	—	Log-likelihood (RW)	2775.9
Benchmark volatility	$\hat{\sigma}_Y$	1.282	0.018	LR statistic (ref.)	36.7
Underlying volatility	$\hat{\sigma}_V$	0.179	—	Dickey–Fuller t	−6.07
Correlation (leverage)	$\hat{\rho}$	−0.725	—	Dickey–Fuller p	< 0.001

Table 4. Parameter stability. Rolling windows report mean ± standard deviation across windows; sub-periods are contiguous calendar splits with Delta-method standard errors in parentheses and AR(1) fit diagnostics.

Specification	$\hat{\beta}$	$\hat{\theta}$	$\hat{\sigma}_Y$	Adj. R^2	Log-lik.
Rolling (2y/6m, 16 windows)	9.69±5.12	2.864±0.207	1.270±0.163	—	—
Sub-period 2015–2018	10.82 (2.54)	2.673 (0.056)	1.346 (0.031)	0.9069	1077
Sub-period 2019–2021	6.11 (2.25)	2.995 (0.105)	1.352 (0.035)	0.9425	800
Sub-period 2022–2024	6.46 (2.31)	2.922 (0.083)	1.119 (0.029)	0.9400	937

where R^V and the OU innovation $\hat{\varepsilon}$ span the same interval, so their sample correlation consistently estimates the Brownian correlation ρ of the model.

The regressions of Equations (16)–(18) are estimated on realised historical paths and thus deliver the physical ($\mathbb{P}$) dynamics $(\hat{\beta}, \hat{\theta}, \sigma_Y, \sigma_V, \rho)$. By Girsanov's theorem, the change of measure $\mathbb{P} \to \mathbb{Q}$ alters only drifts: σ_Y, σ_V, ρ transfer to the pricing formulae of Theorems 1 and 2, whereas the drift parameters β and θ differ between measures through the mapping $\beta = \hat{\beta} - b_y$, $\theta = \hat{\theta} + a_y/\beta$ of Equations (8)–(9). Throughout, as before, $(\hat{\beta}, \hat{\theta})$ denote physical-measure quantities and (β, θ) their risk-neutral counterparts. The risk-neutral drift will be obtained from the futures curve.

Table 3 reports the point estimates. The benchmark is a high "vol-of-vol" object, $\hat{\sigma}_Y \approx 1.28$, with fast mean reversion (half-life ≈ 1.2 months) and a strongly negative correlation to the underlying ($\hat{\rho} \approx -0.73$). The negative sign of the correlation validates the aforementioned negative correlation documented by the leverage-effect literature. Table 3 also reports Delta-method standard errors for the estimated parameters and the AR(1) fit diagnostics. The regression for $\ln I$ attains an adjusted R^2 of approximately 0.94.

The log-likelihood is used to test the property on which the whole framework rests – that the benchmark genuinely mean-reverts rather than following a random walk. The mean reversion the model relies on is strongly supported by the data: the Dickey–Fuller statistic is well below the 1% critical value, decisively rejecting the random walk, and the mean-reverting specification raises the log-likelihood relative to the restricted random walk while lowering both AIC and BIC.

5.3. Parameter stability

We are also interested in studying whether the parameters of the VIX estimation are stable across time. Table 4 summarises the results of the rolling-window estimation (two-year windows, six-month step) for three contiguous sub-periods, reported with Delta-method standard errors and the AR(1) fit diagnostics (adjusted R^2 and log-likelihood). The mean-reversion speed ($\hat{\beta}$) is stable across sub-periods, but the rolling standard deviation is a large fraction of its level, which is consistent with the known difficulty of estimating the reversion speed from short samples, as documented in Yu (2012).

The long-run level and the benchmark volatility, by contrast, are estimated precisely and shift visibly across sub-periods – the average benchmark level $e^{\hat{\theta}}$ moves from about 14.5 in the calm 2015–2018 window to roughly 20 around the 2019–2021 turmoil, and the vol-of-vol falls in the final block – a movement we test formally next.

We test stability formally in Table 5, which reports Wald tests of the null that each parameter is equal across the three sub-periods. The standard error of $\hat{\theta}$ accounts for the strong serial dependence of $\ln I$: because the benchmark is highly persistent ($\hat{\varphi} \approx 0.97$), the variance of its sample mean is far larger than

Table 5. Cross-period parameter-stability tests. Each row is a wald test of equality of the indicated parameter across the three Sub-periods; the statistic is χ^2 with the stated degrees of freedom under the null of equality. Standard errors are calculated using the Delta-method, except the correlation, which is tested via the fisher z-transform ($\mathrm{Var}(z) = 1/(n-3)$). the joint row covers the OU benchmark parameters $(\hat{\beta}, \hat{\theta}, \hat{\sigma}_Y)$; $\hat{\rho}$ and $\hat{\sigma}_V$ are estimated on partly overlapping innovations and are reported separately. *, **, ***: rejection at 10%, 5%, 1%.

Null hypothesis	Parameter	Wald χ^2	d.f.	p-value
Equal reversion speed	$\hat{\beta}$	2.29	2	0.319
Equal long-run level	$\hat{\theta}$	10.77***	2	0.005
Equal benchmark vol.	$\hat{\sigma}_Y$	37.70***	2	0.000
Equal correlation	$\hat{\rho}$	38.82***	2	0.000
Equal underlying vol.	$\hat{\sigma}_V$	197.24***	2	0.000
Joint equality (OU params)	$\hat{\beta}, \hat{\theta}, \hat{\sigma}_Y$	50.75***	6	0.000

the i.i.d. formula would imply, and following Newey and West (1987), we use the corresponding long-run (autocorrelation-corrected) variance.

With the correct standard errors, the mean-reversion speed cannot be distinguished across sub-periods ($p \approx 0.32$), but every other parameter shifts significantly: the long-run level ($p \approx 0.005$), the benchmark volatility ($p < 0.001$), the correlation ($p < 0.001$) [9], and the underlying volatility ($p < 0.001$) all reject equality, as does the joint null on the OU parameters ($p < 0.001$).

The resulting rejection indicates that the constant correlation assumption of the framework does not hold over a decade spanning several regimes. The pattern can be interpreted as an economic regime effect: the estimation window covers the relatively calm mid-2010s, the COVID-19 shock, and the 2022 inflation context. The magnitudes support this interpretation: the long-run level $e^{\hat{\theta}}$ ranges from about 14.5 to 20 across the blocks, and the vol-of-vol ranges from about 1.12 to 1.35, differences that are economically meaningful and not merely statistically detectable.

One of these rejections points to a modelling extension: considering a time-varying correlation. As argued in Subsection 2.4, the pricing impact of relaxing it is typically limited; the rejection of constancy therefore points to modest gains from a stochastic-correlation extension. Whether these gains are material enough to justify the added complexity is ultimately an empirical question, and a formal treatment of stochastic correlation (for instance, along the lines of Antonelli et al. (2010)) remains a natural direction for future work.

5.4. Contract pricing

With the obtained estimations, we price the contracts where $\phi = 1$, as in Theorem 2. This is done with the following specifications: the mean-reverting, non-tradable benchmark with $\lambda^Y = 0$; the futures-calibrated risk-neutral drift of Appendix C; and the fixed-strike Black–Scholes benchmark obtained when $\sigma_Y = 0$, as mentioned in Remark 1.

Initial levels are normalized to 100 for exposition. The calibrated $\hat{\sigma}_V$, $\hat{\sigma}_Y$, $\hat{\rho}$ of Table 3 are used throughout, and Table C.1 (in Appendix C) estimations are used when considering the market price of risk.

Two effects are visible in the results of Table 6. First, the mean-reverting benchmark structure raises value above the fixed-strike benchmark, and this difference increases sharply as the contract moves out of the money. This is because the benchmark volatility $\hat{\sigma}_Y$ enters the effective variance $\tilde{\sigma}^2$ and, through the convex exponential form of I_T, keeps the expected benchmark well above the fixed strike. Ignoring the stochastic, the mean-reverting nature of the exercise threshold hence underestimates these instruments, and this difference becomes significantly substantial as the contracts move further from the money.

Second, the futures-calibrated prices exceed the $\lambda^Y = 0$ prices, quantifying the benchmark risk premium. In intuitive terms, the market prices the benchmark as reverting toward a higher long-run level, and somewhat more slowly, than the physical dynamics imply, so the risk-neutral expected benchmark and the effective variance are larger and contract values are higher. The size of this gap is the practical limitation of the $\lambda^Y = 0$ simplification for this application.

Lastly, Table 7 isolates the two parameters that drive the departure from the BSM setup, reporting the at-the-money price across a grid of the benchmark volatility $\hat{\sigma}_Y$ and the correlation ρ. Value increases monotonically in $\hat{\sigma}_Y$ and rises as ρ becomes more negative, consistent with the covariance term in $\tilde{\sigma}^2$.

Table 6. Contract prices where $\phi = 1$ under three specifications, with the mean-reverting benchmark premium made explicit. MR($\lambda^Y = 0$): mean-reverting benchmark, physical benchmark drift. MR(futures fit): risk-neutral drift from the futures fit. BS: fixed-strike Black–Scholes ($\sigma_Y = 0$, $K = 100$). The last two columns give the MR($\lambda^Y = 0$) price relative to the futures-calibrated price (the sign of the persistence premium) and relative to the fixed-strike benchmark (the value of the mean-reverting structure). Normalised levels, $r = 1.83\%$ (sample-average 3-month Treasury yield), dividend yield $\delta = 1.82\%$ (from the price vs total-return gap).

		Contract price ($\phi = 1$)			MR($\lambda^Y = 0$) vs.	
T	V_0	MR($\lambda^Y = 0$)	MR(futures fit)	BS	% vs. Fut.	% vs. BS
1	70	38.86	56.17	29.57	−31%	+31%
	85	28.75	44.94	16.32	−36%	+76%
	100	20.87	35.52	6.99	−41%	+199%
	115	14.96	27.85	2.33	−46%	+542%
5	70	38.82	54.37	30.41	−29%	+28%
	85	30.54	45.03	21.26	−32%	+44%
	100	23.93	37.19	14.44	−36%	+66%
	115	18.74	30.68	9.61	−39%	+95%

Table 7. Sensitivity of the at-the-money mean-reverting-benchmark contract ($\phi = 1$, $V_0 = 100$, $T = 1$, $\lambda^Y = 0$) to the benchmark volatility σ_Y and the correlation ρ, the two parameters that govern the departure from Black–Scholes. Other inputs as in Table 3. Value rises with σ_Y and as ρ becomes more negative.

	$\rho = -0.9$	$\rho = -0.5$	$\rho = 0.0$
$\sigma_Y = 0.5$	11.03	10.26	9.21
$\sigma_Y = 0.9$	15.80	14.77	13.36
$\sigma_Y = 1.3$	21.66	20.49	18.91

6. Conclusions

This paper has developed an analytically tractable valuation framework for contracts whose payoff depends on the difference between an underlying asset and a non-tradable, mean-reverting benchmark, with applications to the integrated management of financial, operational, and emerging risks. The paper makes several key theoretical contributions.

First, it provides a formal proof that mean-reversion precludes tradability of the benchmark, requiring the specification of a market price of risk. Second, it derives closed-form pricing formulae for contracts that are linked to the Black–Scholes–Merton structure through a risk-neutral expected benchmark and an expanded effective variance. Importantly, these contracts should not be interpreted as standard vanilla options with fixed contractual strikes; the fixed-strike Black–Scholes–Merton case is recovered only when benchmark uncertainty vanishes, or the benchmark becomes effectively deterministic. Third, it provides a systematic analysis of how cost volatility, mean-reversion speed, and cross-process correlation jointly determine investment option value in a real options setting. Fourth, it applies the formulation to three economically important classes of derivatives: weather derivatives, volatility index options, and inflation-linked instruments.

In the proposed contract-pricing framework, the non-tradability of the exercise benchmark and its mean-reverting dynamics are key features that standard fixed-strike models cannot accommodate. The correlation parameter plays a central risk-management role, capturing the negative co-movement between the underlying asset and the risk benchmark that is empirically documented across all three application domains. From a risk and resilience perspective, the framework provides financial institutions, corporates, and regulators with a tractable tool to price and manage exposures to non-tradable benchmarks. The inability to replicate these instruments perfectly – a consequence of market incompleteness arising from non-tradability – implies that pricing requires the judgement-based specification of a market price of risk, an undertaking for which the present model provides a controlled foundation.

The empirical calibration using S&P 500 and VIX data illustrates the practical relevance of the framework. The results show how the physical dynamics of a non-tradable mean-reverting benchmark can be estimated from historical data, how risk-neutral benchmark dynamics can be inferred from traded futures prices, and how benchmark risk premia can materially affect valuation. In particular, the comparison with the fixed-strike Black–Scholes–Merton benchmark shows that ignoring stochastic benchmark uncertainty may understate contract values, especially away from the money.

Several extensions are suggested for future research. In the real options domain, a perpetual American-style option under stochastic mean-reverting cost – relevant to open-ended project commitments – would require numerical methods, but the present framework offers a foundation for the endeavor. Extensions to regime-switching, CEV or Heston dynamics for the underlying asset, or to time-varying correlation structures, would further broaden the model's empirical applicability. Model calibration to other contexts – such as inflation and weather as benchmarks – constitutes a further avenue of practical relevance.

Notes

1. This is equivalent to selecting the minimal martingale measure as in Föllmer and Schweizer (1991). For empirical applications, non-zero λ would be calibrated from observed option prices; the closed-form structure remains unchanged, with parameters adjusted as in equations (10)-(11).
2. The VIX, published by the CBOE, measures the market's expectation of 30-day volatility in the S&P 500, derived from S&P 500 option prices.
3. The CPI, published by national statistical agencies, measures the weighted average change over time in the prices paid by consumers for a representative basket of goods and services.
4. Agriculture, forestry, animal husbandry, and fishery; manufacturing; information transmission, software, and information technology services; finance; real estate; leasing and business services; water, environment, and public facilities management.
5. For the S&P 500, ^SP500TR is the total-return index ticker, which reinvests dividends, and ^GSPC is the price index ticker, which excludes them. ^VIX is the volatility index ticker.
6. Ticker DGS3MO.
7. from .https://www.cboe.com/tradable_products/vix/vix_futures/
8. The window therefore does not coincide with the other parameters. The risk-neutral speed thus reflects current term-structure pricing, while the physical speed reflects the historical sample. This does not affect the qualitative result and allows a replication with time-aligned data.
9. The correlation is tested through the Fisher *z*-transform.
10. The fit treats the coefficients (a_y, b_y) as constant, so λ_t^Y varies only through the state Y_t.

Author contributions

CRediT: **Paulo Viegas Carvalho**: Formal analysis, Investigation, Methodology, Supervision, Validation, Visualization, Writing – review & editing; **Carlos Manuel Pinheiro**: Formal analysis, Investigation, Methodology, Supervision, Validation, Visualization, Writing – review & editing; **João Miguel Reis**: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Disclosure statement

The authors report that there are no competing interests to declare.

Funding

No funding was received for this work.

About the authors

Paulo Viegas de Carvalho is Associate Professor of Finance at ISCTE Business School and Adjunct Professor at ISCAL, teaching and coordinating courses in financial risk, corporate finance, and financial markets. His research focuses on financial markets, risk management, banking regulation, sustainable finance, and risk models, with publications in high-impact journals. He holds a PhD in Finance from ISCTE Business School and has extensive experience teaching executives from central banks, banks, corporations, and governments.

Carlos Manuel Pinheiro is an Associate Professor of Finance with habilitation in Management at Universidade Europeia, as well as a Visiting Adjunct Professor at the Lisbon Accounting and Business School. He holds a PhD in Finance from ISCTE Business School, and he combines over 30 years of senior banking management experience

with his academic career. His work focuses on international financial institutions and markets, governance, and behavioural finance.

João Miguel Reis is an Assistant Professor at Universidade Europeia. He holds a PhD in Finance, with Credit Risk and Financial Markets as his main research areas. He is an integrated researcher at the Business Research Unit of ISCTE.

ORCID

Paulo Viegas Carvalho http://orcid.org/0000-0002-4026-1112
Carlos Manuel Pinheiro http://orcid.org/0000-0001-8897-0332
João Miguel Reis http://orcid.org/0000-0002-4700-424X

Data availability statement

The data that support the findings of this study are available from the corresponding author, PVC, upon reasonable request.

References

Addoum, J. M., Ng, D. T., & Ortiz-Bobea, A. (2023). Temperature shocks and industry earnings news. *Journal of Financial Economics*, *150*(1), 1–45. https://doi.org/10.1016/j.jfineco.2023.07.002

Alaton, P., Djehiche, B., & Stillberger, D. (2002). On modelling and pricing weather derivatives. *Applied Mathematical Finance*, *9*(1), 1–20. https://doi.org/10.1080/13504860210132897

Alexander, C., & Chen, X. (2021). Model risk in real option valuation. *Annals of Operations Research*, *299*(1–2), 1025–1056. https://doi.org/10.1007/s10479-019-03273-4

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies*, *15*(4), 1137–1187. https://doi.org/10.1093/rfs/15.4.1137

Antonelli, F., Ramponi, A., & Scarlatti, S. (2010). Exchange option pricing under stochastic volatility: A correlation expansion. *Review of Derivatives Research*, *13*(1), 45–73. https://doi.org/10.1007/s11147-009-9043-4

Bakshi, G., & Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, *16*(2), 527–566. https://doi.org/10.1093/rfs/hhg002

Benth, F. E., & Šaltytė-Benth, J. (2005). Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, *12*(1), 53–85. https://doi.org/10.1080/1350486042000271638

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654. https://doi.org/10.1086/260062

Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, *22*(11), 4463–4492. https://doi.org/10.1093/rfs/hhp008

Boudoukh, J., & Richardson, M. (1993). Stock returns and inflation: A long-horizon perspective. *The American Economic Review*, *83*, 1346–1355.

Brennan, M. J., & Schwartz, E. S. (1985). Evaluating natural resource investments. *The Journal of Business*, *58*(2), 135–157. https://doi.org/10.1086/296288

Briest, G., Lauven, L. P., Kupfer, S., & Lukas, E. (2022). Leaving well-worn paths: Reversal of the investment-uncertainty relationship and flexible biogas plant operation. *European Journal of Operational Research*, *300*(3), 1162–1176. https://doi.org/10.1016/j.ejor.2021.10.035

Brody, D. C., Syroka, J., & Zervos, M. (2002). Dynamical pricing of weather derivatives. *Quantitative Finance*, *2*(3), 189–198. https://doi.org/10.1088/1469-7688/2/3/302

Carr, P., & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, *22*(3), 1311–1341. https://doi.org/10.1093/rfs/hhn038

Cheng, X., Wang, Y., & Wu, X. (2022). The effects of drought on stock prices: An industry-specific perspective. *Frontiers in Environmental Science*, *10*, 978404. https://doi.org/10.3389/fenvs.2022.978404

Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of Financial Economics*, *10*(4), 407–432. https://doi.org/10.1016/0304-405X(82)90018-6

Cox, J. C. (1975). Notes on option pricing I: Constant elasticity of variance diffusions. Working Paper, Stanford University. Reprinted in *Journal of Portfolio Management*, *23*(1996), 15–17.

Cox, J. C., & Ross, S. A. (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, *3*(1–2), 145–166. https://doi.org/10.1016/0304-405X(76)90023-4

Décamps, J. P., Mariotti, T., & Villeneuve, S. (2006). Irreversible investment in alternative projects. *Economic Theory*, *28*(2), 425–448. https://doi.org/10.1007/s00199-005-0629-2

Detemple, J., & Osakwe, C. (2000). The valuation of volatility options. *Review of Finance*, *4*(1), 21–50. https://doi.org/10.1023/A:1009814324980

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*(3), 339–350. https://doi.org/10.1198/073500102288618487

Fama, E. F. (1981). Stock returns, real activity, inflation, and money. *The American Economic Review*, *71*, 545–565.

Fama, E. F., & Schwert, G. W. (1977). Asset returns and inflation. *Journal of Financial Economics*, *5*(2), 115–146. https://doi.org/10.1016/0304-405X(77)90014-9

Fleming, J., Ostdiek, B., & Whaley, R. E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, *15*(3), 265–302. https://doi.org/10.1002/fut.3990150303

Föllmer, H., & Schweizer, M. (1991). Hedging of contingent claims under incomplete information. In Davis, M.H.A., Elliott, R.J. (Eds.), *Applied Stochastic Analysis. Gordon and Breach., volume 5 of* Stochastics Monographs, pp. 389–414.

Grünbichler, A., & Longstaff, F. A. (1996). Valuing futures and options on volatility. *Journal of Banking & Finance*, *20*(6), 985–1001. https://doi.org/10.1016/0378-4266(95)00034-8

Heath, D., Jarrow, R., & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, *60*(1), 77–105. https://doi.org/10.2307/2951677

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, *6*(2), 327–343. https://doi.org/10.1093/rfs/6.2.327

Hong, H., Li, F. W., & Xu, J. (2019). Climate risks and market efficiency. *Journal of Econometrics*, *208*(1), 265–281. https://doi.org/10.1016/j.jeconom.2018.09.015

Hull, J., & White, A. (1990). Pricing interest-rate-derivative securities. *Review of Financial Studies*, *3*(4), 573–592. https://doi.org/10.1093/rfs/3.4.573

Iacus, S. M. (2008). *Simulation and Inference for Stochastic Differential Equations: With R Examples*. Springer.

Jaimungal, S., de Souza, M. O., & Zubelli, J. P. (2013). Real option pricing with mean-reverting investment and project value. *The European Journal of Finance*, *19*(7–8), 625–644. https://doi.org/10.1080/1351847X.2011.601660

Jarrow, R., & Yildirim, Y. (2003). Pricing treasury inflation protected securities and related derivatives using an HJM model. *The Journal of Financial and Quantitative Analysis*, *38*(2), 337–358. https://doi.org/10.2307/4126754

Kendall, M. G. (1954). Note on bias in the estimation of autocorrelation. *Biometrika*, *41*(3/4), 403–404. https://doi.org/10.2307/2332720

Kim, J. H., & Park, C. R. (2017). A multiscale extension of the Margrabe formula under stochastic volatility. *Chaos, Solitons & Fractals*, *97*, 59–65. https://doi.org/10.1016/j.chaos.2017.02.006

Lin, Y. N. (2007). Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. *Journal of Futures Markets*, *27*(12), 1175–1217. https://doi.org/10.1002/fut.20291

Mac Cawley, A., Cubillos, M., & Pascual, R. (2020). A real options approach for joint overhaul and replacement strategies with mean reverting prices. *Annals of Operations Research*, *286*(1–2), 303–324. https://doi.org/10.1007/s10479-018-2906-z

Majd, S., & Pindyck, R. S. (1987). Time to build, option value, and investment decisions. *Journal of Financial Economics*, *18*(1), 7–27. https://doi.org/10.1016/0304-405X(87)90059-6

Margrabe, W. (1978). The value of an option to exchange one asset for another. *The Journal of Finance*, *33*(1), 177–186. https://doi.org/10.1111/j.1540-6261.1978.tb03397.x

McDonald, R. L., & Siegel, D. R. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, *101*(4), 707–728. https://doi.org/10.2307/1884175

Mencía, J., & Sentana, E. (2013). Valuation of vix derivatives. *Journal of Financial Economics*, *108*(2), 367–391. https://doi.org/10.1016/j.jfineco.2012.12.003

Mercurio, F. (2005). Pricing inflation-indexed derivatives. *Quantitative Finance*, *5*, 289–302.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, *4*(1), 141–183. https://doi.org/10.2307/3003143

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708. https://doi.org/10.2307/1913610

Pindyck, R. S. (1993). Investments of uncertain cost. *Journal of Financial Economics*, *34*(1), 53–76. https://doi.org/10.1016/0304-405X(93)90040-I

Sarkar, S. (2003). The effect of mean reversion on investment under uncertainty. *Journal of Economic Dynamics and Control*, *28*(2), 377–396. https://doi.org/10.1016/S0165-1889(02)00181-1

Schwartz, E. S. (2004). Patents and R&D as real options. *Economic Notes*, *33*(1), 23–54. https://doi.org/10.1111/j.0391-5026.2004.00124.x

Trigeorgis, L. (1993). Real options and interactions with financial flexibility. *Financial Management*, *22*(3), 202–224. https://doi.org/10.2307/3665939

Whaley, R. E. (2000). The investor fear gauge. *The Journal of Portfolio Management*, *26*(3), 12–17. https://doi.org/10.3905/jpm.2000.319728

Yu, J. (2012). Bias in the estimation of the mean reversion parameter in continuous time models. *Journal of Econometrics*, *169*(1), 114–122. https://doi.org/10.1016/j.jeconom.2012.01.004

Zhang, J. E., & Zhu, Y. (2006). VIX futures. *Journal of Futures Markets*, *26*(6), 521–531. https://doi.org/10.1002/fut.20209

Zhu, Y., & Zhang, J. E. (2007). Variance term structure and VIX futures pricing. *International Journal of Theoretical and Applied Finance*, *10*(01), 111–127. https://doi.org/10.1142/S0219024907004123

Appendix A:

Proof of Lemma 1

Proof. For I_t to be a traded asset, the discounted process $e^{-rt}I_t$ must be a $\mathbb{Q}$-martingale. This would require:

$$\mathbb{E}^{\mathbb{Q}}[I_T|\mathcal{F}_t] = e^{r(T-t)}I_t \quad \text{for all } t < T. \tag{A.1}$$

Recalling the OU dynamics then presented for the risk-neutral measure parameters in (15), for tradability the expression must equal (A.1), thus

$$\exp\left\{\hat{\theta} + e^{-\hat{\beta}(T-t)}Y_t + \frac{\sigma_Y^2}{4\hat{\beta}}(1 - e^{-2\hat{\beta}(T-t)})\right\} = e^{r(T-t)}e^{\hat{\theta}+Y_t}.$$

Taking logarithms and rearranging:

$$Y_t\left(e^{-\hat{\beta}(T-t)} - 1\right) = r(T-t) - \frac{\sigma_Y^2}{4\hat{\beta}}(1 - e^{-2\hat{\beta}(T-t)}).$$

Since this equality must hold $\mathbb{Q}$-almost surely, and Y_t has full support on $\mathbb{R}$ under $\mathbb{Q}$, the coefficient of Y_t on the left-hand side must vanish identically. That is, when:

$$e^{-\hat{\beta}(T-t)} - 1 = 0,$$

which forces $\hat{\beta} = 0$. However, taking the limit $\hat{\beta} \to 0$ (via L'Hôpital's rule), the right-hand side becomes $r(T-t) - \frac{\sigma_Y^2}{2}(T-t)$, which is non-zero for generic parameters r and σ_Y. Hence no valid $\hat{\beta}$ exists that satisfies the martingale condition (A.1), and we conclude that I_t cannot be a traded asset. ■

Appendix B:

Proof of Theorem 1

Proof. Through the Radon-Nikodym derivative, we define the measure $\mathbb{P}^T$ using the strictly positive benchmark payoff I_T as the change-of-measure density.

$$\frac{d\mathbb{P}^T}{d\mathbb{Q}}|_{\mathcal{F}_t} = \frac{\mathbb{E}[I_T \mid \mathcal{F}_t]}{\mathbb{E}[I_T \mid \mathcal{F}_0]}.$$

Knowing that I_t is always greater than zero, we rewrite Equation (1), using the I_T forward measure from above, change the measure to $\mathbb{P}^T$ and define ξ_t:

$$\begin{aligned}
P(V_0, I_0) &= e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[(V_T - I_T)^+\right] \\
&= e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[I_T\left(\tfrac{V_T}{I_T} - 1\right)^+\right] \\
&= e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[\frac{I_T}{\mathbb{E}^{\mathbb{Q}}[I_T]} \cdot \mathbb{E}^{\mathbb{Q}}[I_T] \cdot \left(\frac{V_T}{I_T} - 1\right)^+\right] \\
&= e^{-rT}\mathbb{E}^{\mathbb{Q}}[I_T] \cdot \mathbb{E}^{\mathbb{Q}}\left[\frac{I_T}{\mathbb{E}^{\mathbb{Q}}[I_T]} \cdot \left(\frac{V_T}{I_T} - 1\right)^+\right] \\
&= e^{-rT}\mathbb{E}^{\mathbb{Q}}[I_T] \cdot \mathbb{E}^{\mathbb{P}^T}\left[\left(\tfrac{V_T}{I_T} - 1\right)^+\right] \\
&= e^{-rT}\mathbb{E}^{\mathbb{Q}}[I_T] \cdot \mathbb{E}^{\mathbb{P}^T}\left[(\xi_T - 1)^+\right]
\end{aligned} \tag{B.1}$$

where we use the change of measure formula: for any $\mathcal{F}_T$-measurable random variable X, $\mathbb{E}^{\mathbb{Q}}\left[\frac{I_T}{\mathbb{E}^{\mathbb{Q}}[I_T|\mathcal{F}_0]} \cdot X\right] = \mathbb{E}^{\mathbb{P}^T}[X]$.

To evaluate $\mathbb{E}^{\mathbb{P}^T}[(\xi_T - 1)^+]$, we define the process ξ_t for all $t \in [0, T]$ as the ratio of conditional expectations:

$$\xi_t \triangleq \frac{\mathbb{E}^{\mathbb{Q}}[V_T \mid \mathcal{F}_t]}{\mathbb{E}^{\mathbb{Q}}[I_T \mid \mathcal{F}_t]}.$$

At the terminal time, this reduces to:

$$\xi_T = \frac{\mathbb{E}^{\mathbb{Q}}[V_T \mid \mathcal{F}_T]}{\mathbb{E}^{\mathbb{Q}}[I_T \mid \mathcal{F}_T]} = \frac{V_T}{I_T},$$

since conditioning on all information up to time T gives the realized values.

We then verify that ξ_t is a martingale under the forward measure $\mathbb{P}^T$. For $0 \leq s < t \leq T$:

$$\begin{aligned}\mathbb{E}^{\mathbb{P}^T}[\xi_t \mid \mathcal{F}_s] &= \mathbb{E}^{\mathbb{Q}}\left[\frac{\mathbb{E}[V_T \mid \mathcal{F}_t]}{\mathbb{E}[I_T \mid \mathcal{F}_t]} \cdot \frac{\mathbb{E}[I_T \mid \mathcal{F}_t]/\mathbb{E}[I_T]}{\mathbb{E}[I_T \mid \mathcal{F}_s]/\mathbb{E}[I_T]} \middle| \mathcal{F}_s\right] \\ &= \frac{1}{\mathbb{E}[I_T \mid \mathcal{F}_s]} \mathbb{E}^{\mathbb{Q}}[\mathbb{E}[V_T \mid \mathcal{F}_t] \mid \mathcal{F}_s] \\ &= \frac{\mathbb{E}[V_T \mid \mathcal{F}_s]}{\mathbb{E}[I_T \mid \mathcal{F}_s]} = \xi_s.\end{aligned} \tag{B.2}$$

Therefore, ξ_t is indeed a $\mathbb{P}^T$-martingale.

With the results presented in Equations (14) and (15), the process ξ_t can be written as:

$$\xi_t = \frac{V_t e^{r(T-t)}}{\exp\left\{\theta + e^{-\beta(T-t)}Y_t + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta(T-t)})\right\}}. \tag{B.3}$$

To show that ξ_T is log-normally distributed under $\mathbb{P}^T$, we apply Itô's lemma to $\ln \xi_t$.

From Equation (B.3), we can write:

$$\ln \xi_t = \ln V_t + r(T-t) - \theta - e^{-\beta(T-t)}Y_t - \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta(T-t)}).$$

Since V_t follows a GBM under $\mathbb{Q}$:

$$d \ln V_t = \left(r - \delta - \frac{1}{2}\sigma_V^2\right)dt + \sigma_V dW_t^V.$$

For the Y_t term, recalling again that:

$$dY_t = -\beta Y_t dt + \sigma_Y dW_t^Y.$$

Thus

$$\begin{aligned} d \ln \xi_t &= d \ln V_t - d\left(e^{-\beta(T-t)}Y_t\right) - d\left(\frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta(T-t)})\right) \\ d \ln \xi_t &= \left(r - \delta - \frac{1}{2}\sigma_V^2 + \frac{\sigma_Y^2}{2\beta}e^{-2\beta(T-t)}\right)dt + \sigma_V dW_t^V - \sigma_Y e^{-\beta(T-t)}dW_t^Y. \end{aligned}$$

Since by Equation (B.2), ξ_t is a $\mathbb{P}^T$-martingale, its drift under $\mathbb{P}^T$ must be zero. Thus, under this measure:

$$d \ln \xi_t = \sigma_V dW_t^{T,V} - \sigma_Y e^{-\beta(T-t)}dW_t^{T,Y},$$

where $W_t^{T,V}$ and $W_t^{T,Y}$ are $\mathbb{P}^T$-Brownian motions with correlation ρ.

Integrating from 0 to T:

$$\ln \xi_T - \ln \xi_0 = \sigma_V W_T^{T,V} - \sigma_Y \int_0^T e^{-\beta(T-s)}dW_s^{T,Y}.$$

The right-hand side is a Gaussian random variable with mean zero. To compute its variance, we apply the variance of a difference formula.

For the first term:

$$\mathrm{Var}(\sigma_V W_T^{T,V}) = \sigma_V^2 T.$$

For the second term, by Itô's isometry:

$$\mathrm{Var}\left(\sigma_Y \int_0^T e^{-\beta(T-s)}dW_s^{T,Y}\right) = \sigma_Y^2 \int_0^T e^{-2\beta(T-s)}ds.$$

As for the covariance between the terms, using the correlation ρ between $W^{T,V}$ and $W^{T,Y}$:

$$\mathrm{Cov}\left(\sigma_V W_T^{T,V}, \sigma_Y \int_0^T e^{-\beta(T-s)}dW_s^{T,Y}\right) = \rho\sigma_V\sigma_Y \int_0^T e^{-\beta(T-s)}ds.$$

Therefore,

$$\tilde{\sigma}^2 = \sigma_V^2 T + \sigma_Y^2 \int_0^T e^{-2\beta(T-s)}ds - 2\rho\sigma_V\sigma_Y \int_0^T e^{-\beta(T-s)}ds. \tag{B.4}$$

and computing the integrals

$$\tilde{\sigma}^2 = \sigma_V^2 T + \frac{\sigma_Y^2}{2\beta}(1 - e^{-2\beta T}) - 2\rho\sigma_V\sigma_Y \frac{1 - e^{-\beta T}}{\beta}.$$

Therefore, ξ_T is log-normally distributed under $\mathbb{P}^T$:

$$\xi_T \overset{d}{=} \xi_0 \exp\left\{-\frac{1}{2}\tilde{\sigma}^2 + \tilde{\sigma}Z\right\}, \tag{B.5}$$

where $Z \sim \mathcal{N}(0,1)$ under $\mathbb{P}^T$.

Given the result from Equation (B.5), we can apply the standard option pricing formula for log-normal variables. In our case, with 1 in place of the strike, we have:

$$\mathbb{E}^{\mathbb{P}^T}\left[(\xi_T - 1)^+\right] = \xi_0\Phi(d_+) - \Phi(d_-),$$

where:

$$d_\pm = \frac{\ln \xi_0 \pm \frac{1}{2}\tilde{\sigma}^2}{\tilde{\sigma}}.$$

Substituting back into Equation (B.1), the contract price is:

$$P(V_0, I_0) = e^{-\delta T}V_0\Phi(d_+) - e^{-rT}\mathbb{E}^{\mathbb{Q}}[I_T]\Phi(d_-)$$

where $\xi_0 = \frac{\mathbb{E}^{\mathbb{Q}}[V_T]}{\mathbb{E}^{\mathbb{Q}}[I_T]} = \frac{V_0 e^{(r-\delta)T}}{\mathbb{E}^{\mathbb{Q}}[I_T]}$, and:

$$d_\pm = \frac{\ln\left(V_0 e^{(r-\delta)T}/\mathbb{E}^{\mathbb{Q}}[I_T]\right) \pm \frac{1}{2}\tilde{\sigma}^2}{\tilde{\sigma}},$$
$$\tilde{\sigma}^2 = \sigma_V^2 T + \frac{\sigma_Y^2}{2\beta}(1 - e^{-2\beta T}) - 2\rho\sigma_V\sigma_Y\frac{1 - e^{-\beta T}}{\beta}.$$

The expectations $\mathbb{E}^{\mathbb{Q}}[V_T]$ and $\mathbb{E}^{\mathbb{Q}}[I_T]$ are given in Equations (13) and (15), respectively.

Appendix C:

Market price of risk

With Subsection 5.2 covering the measure-invariant diffusions $(\sigma_V, \sigma_Y, \rho)$, already estimated under $\mathbb{P}$, we find the risk-neutral drift of the benchmark. For the tradable underlying (S&P 500) the market price of risk is fixed by no-arbitrage, and the drift is $r - \delta$, so no premium is estimated. The question arises in the non-tradable benchmark, whose price of risk is unidentified by arbitrage (as presented in Lemma 1) and must be recovered from market data. This is done with the VIX futures term structure.

Because a futures price is a risk-neutral expectation of the settlement value, the curve $F(0, T_i)$ is the market's $\mathbb{E}^{\mathbb{Q}}[\text{VIX}_{T_i}]$ across maturities, as explored in Grünbichler and Longstaff (1996) and Zhang and Zhu (2006). This makes the futures curve the natural instrument for identifying the risk-neutral drift: the no-arbitrage relation linking the drift of the risk-neutral dynamics to the shape of the forward curve is established for volatility indices by Zhu and Zhang (2007).

Under the log-normal OU dynamics of Equations (8)-(9) – the specification analysed for volatility indices by Detemple and Osakwe (2000) – the conditional forward has the closed form

$$\mathbb{E}^{\mathbb{Q}}[I_T] = \exp\{\theta + e^{-\beta T}Y_0 + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta T})\}. \tag{C.1}$$

This procedure hence obtains the risk-neutral parameters by fitting the model-implied forward from Equation (C.1) to observed futures, so that the physical and risk-neutral measures are estimated jointly rather than in isolation, as per Zhang and Zhu (2006) and Lin (2007).

$$\min_{\beta,\theta}\sum_i\left[\ln F(0,T_i) - (\theta + e^{-\beta T_i}(\ln I_0 - \theta) + \frac{\sigma_Y^2}{4\beta}(1 - e^{-2\beta T_i}))\right]^2. \tag{C.2}$$

The fit anchors the benchmark at its observed spot on the valuation date ($Y_0 = \ln I_0 - \theta$) treated inside the least-squares problem: given β, the minimising θ is available in closed form, and β is found by grid search.

Inverting the mapping Equations (8)-(9) identifies the market price of risk of the benchmark process itself, as presented in Equation (12):

$$\lambda_t^Y = \frac{a_y + b_y\ Y_t}{\sigma_Y},$$

with the two constants [10] recovered as

$$b_y = \hat{\beta} - \beta, \qquad a_y = \beta(\theta - \hat{\theta}).$$

The constant part a_y generates the long-run level shift $\theta = \hat{\theta} + a_y/\beta$ of Equation (11), while the state-dependent slope b_y generates the reversion-speed shift $\beta = \hat{\beta} - b_y$ of Equation (10).

Table C.1: Risk-neutral ($\mathbb{Q}$) calibration to the VIX futures curve, and the implied market price of risk of the benchmark, $\lambda_t^Y = (a_y + b_y Y_t)/\sigma_Y$, via its two constants (a_y, b_y). The calibration anchors the benchmark at its observed spot on the valuation date.

Quantity	Symbol	Value
Benchmark spot (valuation date)	I_0	16.15
Spot anchor	$Y_0 = \ln I_0 - \theta$	−0.207
Risk-neutral reversion speed	β	5.46
Risk-neutral half-life	$\ln 2/\beta$	1.5 months
Risk-neutral long-run level	e^{θ}	19.86
Persistence premium/MPR slope	$b_y = \hat{\beta} - \beta$	+1.54
Level premium/MPR constant	$a_y = \beta(\theta - \hat{\theta})$	+0.787

We calibrate the risk-neutral drift to the aforementioned observed VIX futures term structure (CBOE VIX monthly settlement curve). Each contract's expiration date is converted to a year-fraction $T_i = (\text{expiry} - \text{valuation})/365$, and the settlement price is used as the market forward $F(0, T_i) = \mathbb{E}^{\mathbb{Q}}[\text{VIX}_{T_i}]$.

Table 1 presents the results from solving Equation (C.2) on this curve, with $\hat{\sigma}_Y$ held at its physical estimate. The risk-neutral reversion speed ($\beta \approx 5.5$) remains below the physical speed ($\hat{\beta} \approx 7.0$) – the pattern documented in Carr and Wu (2009), whereby mean reversion under $\mathbb{P}$ becomes slower under $\mathbb{Q}$ – yielding a positive persistence premium ($b_y \approx 1.5$). The spot-anchored fit also assigns an economically meaningful level premium: the risk-neutral long-run level ($e^{\theta} \approx 19.9$) lies well above its physical counterpart ($e^{\hat{\theta}} \approx 17.2$), so $a_y \approx 0.79$. The positive volatility risk premium of Bakshi and Kapadia (2003) and Bollerslev et al. (2009) is thus expressed through both channels: the market prices the benchmark as reverting somewhat more slowly, and toward a higher long-run level, than the physical dynamics imply.