mathematics

MDPI

Article

Designed Dependence for Tractable Two-Asset First-Passage Pricing

Marcos Escobar-Anel

Department of Statistical and Actuarial Sciences, Western University, London, ON N6A 5B7, Canada; marcos.escobar@uwo.ca

Abstract

Closed-form multi-asset first-passage prices become rare once dependence is allowed to vary with the state. We build a two-asset local-volatility/local-correlation family in which relative-performance passage problems stay one-dimensional by construction. Under a common instantaneous volatility, the log-ratio and log-product are orthogonal, and their variance rates form a simple budget. Holding the variance rate of the log-ratio fixed while leaving the orthogonal index variance free produces the Brownian-dispersion model: the log-ratio is an exact Brownian motion under the pricing measure, whereas the geometric-mean index carries the state-dependent volatility and correlation. Cash-settled and one-homogeneous claims whose path dependence is carried by the ratio are then invariant to the local-correlation specification; the second property holds more generally under constant relative variance, while the first is specific to the orthogonal design. We derive double-knock-out exchange-option, cash-rebate and relative-performance lookback formulas and identify an index-dependent subclass in which ratio and index are independent. There, a two-dimensional exit law factorises and an exchange option knocked out on a curvilinear region is priced as a ratio double-knock-out value times a one-dimensional index survival probability. Simulations under sharply different local-correlation functions confirm the invariant prices while generating very different fixed-strike basket prices.

Keywords: local correlation; dependence design; first-passage times; barrier options; exchange options

MSC: 91G20; 60J65; 60G40

1. Introduction

Closed-form solutions to multi-asset first-passage problems are rare. Even for a pair of correlated Brownian motions, the laws of joint hitting times and double extrema typically require special geometry, image arguments, or spectral expansions [1–3]. Homogeneous claims form an important special case: change of numeraire together with scale invariance identifies the relative price as the natural state variable [4–6]. In the constant-parameter Black–Scholes–Margrabe model, this yields Margrabe's exchange-option formula as well as one-dimensional ratio-barrier formulas [7–9].

The question addressed here is whether such tractability can survive when instantaneous volatility and dependence are state dependent. Local volatility and Markovian projection reproduce one-time marginal laws [10–13], whereas first passage is a property of the path. Local-correlation and cross-dependent-volatility models provide flexible joint dynamics [14–17], and the barrier-exchange literature also clarifies the replication structure

Check for updates

Academic Editor: J.E. Trinidad-Segovia

Received: 28 September 2026
Revised: 20 October 2026
Accepted: 22 October 2026
Published: 25 October 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

of special ratio barriers [18]. In foreign exchange, where a cross rate is exactly the ratio of two rates against a common currency, the dependence between the two main pairs must be designed so that the cross is consistent with its own market [19,20]; these constructions target the one-time law of the ratio. The approaches above are not usually built so that a path coordinate has an exact prescribed law.

We rely on a variance-allocation construction. With $U = \log(S_1/S_2)$, $V = \log(S_1S_2)$ and $P = \sqrt{S_1S_2}$, equal instantaneous volatilities make U and V orthogonal for any correlation process. We therefore specify their variance rates directly. Fixing the variance rate of U gives the Brownian-dispersion (BD) model: U is an exact Brownian motion, while the orthogonal index component carries the state-dependent volatility and correlation.

The central pricing result is an invariance identity,

$$\mathbb{E}[e^{-r\tau}P_\tau G] = P_0\mathbb{E}[e^{-\sigma_U^2\tau/8}G],$$

for ratio stopping times and ratio-path payoffs. For one-homogeneous claims, this identity is an instance of the change-of-numeraire mechanism and, as we note, holds in any model with constant relative variance, with or without equal volatilities. What the orthogonal design adds is twofold: the ratio is a Brownian motion under the pricing measure itself, so cash-settled ratio claims are invariant as well; and, in an index-autonomous subclass, ratio and index are independent. We retain two representative applications—a double-knock-out exchange option and a floating-strike relative-performance lookback—together with a cash rebate at the ratio barrier, rather than cataloguing every barrier and path-dependent formula implied by the same reduction. In the index-autonomous subclass, the exit law from a region bounded by two ratio rays and two index hyperbolas factorises, and the corresponding knock-out exchange option is priced as a ratio double-knock-out value times a one-dimensional index survival probability. A numerical experiment compares two sharply different local-correlation functions: the invariant prices coincide, whereas the fixed-strike basket prices differ substantially.

The design viewpoint is motivated by Escobar-Anel [21]. Related analytical work on multi-asset barriers includes structured stochastic covariance [22], special correlation geometries in three dimensions [23] and multiple barriers under stochastic volatility [24]. Classical curved-boundary and path-dependent techniques include Kunitomo and Ikeda [25] and Goldman et al. [26]. The main contribution of this paper is the construction of an admissible state-dependent dependence family that preserves an exactly solvable first-passage coordinate and its separation from the orthogonal index risk. We distinguished three levels of results next.

- First, the established theory.The reduction in one-homogeneous claims by change of numeraire and scale invariance is classical [4–6]. Margrabe pricing, ratio-barrier reduction and the Brownian image and maximum laws used below are applications of that theory, not new first-passage identities.
- Second, consequences of the BD specification.Equal instantaneous volatilities cancel the pricing-measure log-ratio drift. Fixing its variance rate then fixes its entire Brownian path law under $\mathbb{Q}$; this also makes cash-settled ratio claims invariant. Orthogonality alone does not imply independence of the state processes.
- Third, the main contribution.The variance-budget parameterisation preserves these path laws with an automatically admissible, state-dependent correlation. In its index-autonomous subclass, the chosen ratio–index coordinates and matching boundary geometry make a curvilinear two-dimensional exit problem separable. The distinction between an invariant ratio book and an index-sensitive book also gives a pricing, calibration and hedging interpretation.

In particular, Haug and Haug [8] treat ratio barriers in the constant-parameter exchange-option setting, whereas we preserve that coordinate under state-dependent common volatility and correlation. Guyon [16] calibrates local-correlation models to basket smiles; our family instead fixes the relative path law and leaves the index variance to be constrained by additional instruments. Relative to the European-claim design in Escobar-Anel [21], the target here is a path law, and the variance parameterisation enforces correlation admissibility rather than repairing it after inversion. These are differences in model construction; the familiar homogeneous reduction is not claimed as a new general principle.

Sections 2 and 3 develop the model and the invariance result, Section 4 gives the path-dependent applications, Section 5 gives the ratio–index factorisation, and Section 6 provides numerical evidence.

2. Orthogonal Coordinates and the Variance Budget

2.1. Setting

Let $T^* \in (0, \infty]$; an unbounded stopping-time result below presumes $T^* = \infty$. All fixed-maturity applications require only $T^* \geq T$. When $T^* < \infty$, stopping rules below are understood to satisfy $\tau \leq T^*$ almost surely.

Let $(\Omega, \mathcal{F}, \mathbb{Q}, \{\mathcal{F}_t\}_{t\in[0,T^*]})$ be a filtered probability space satisfying the usual conditions and carrying two standard Brownian motions W_1, W_2 with $\mathrm{d}\langle W_1, W_2\rangle_t = \rho_t \,\mathrm{d}t$ for a predictable process ρ taking values in $[-1, 1]$. We work directly under a risk-neutral measure $\mathbb{Q}$; $r \geq 0$ is a constant short rate. The two traded assets pay no dividends and satisfy

$$\frac{\mathrm{d}S_{i,t}}{S_{i,t}} = r\,\mathrm{d}t + \sigma_{i,t}\,\mathrm{d}W_{i,t}, \qquad S_{i,0} > 0, \qquad i = 1, 2, \tag{1}$$

with σ_1, σ_2 being predictable, strictly positive and locally square integrable. Set $X_i = \ln S_i$ and define the rotated coordinates

$$U_t := X_{1,t} - X_{2,t} = \ln \frac{S_{1,t}}{S_{2,t}}, \qquad V_t := X_{1,t} + X_{2,t} = \ln\left(S_{1,t}S_{2,t}\right). \tag{2}$$

We write $R_t = e^{U_t} = S_{1,t}/S_{2,t}$ for the ratio and $P_t = e^{V_t/2} = \sqrt{S_{1,t}S_{2,t}}$ for the geometric-mean index. The map $(S_1, S_2) \mapsto (R, P)$ is a smooth bijection of $(0, \infty)^2$ onto itself, whose inverse is $S_1 = P\sqrt{R}$, $S_2 = P/\sqrt{R}$.

2.2. Orthogonalisation

Lemma 1. *In the market* (1)*,*

$$\mathrm{d}\langle U, V\rangle_t = \left(\sigma_{1,t}^2 - \sigma_{2,t}^2\right)\mathrm{d}t.$$

Consequently, $\langle U, V\rangle \equiv 0$ *if and only if* $\sigma_{1,t} = \sigma_{2,t}$ *for* $\mathrm{d}t \otimes \mathrm{d}\mathbb{Q}$*—almost every* (t, ω)*. Writing* $\sigma := \sigma_1 = \sigma_2$ *under that condition, we have*

$$\mathrm{d}U_t = \sigma_t\left(\mathrm{d}W_{1,t} - \mathrm{d}W_{2,t}\right), \qquad \mathrm{d}\langle U\rangle_t = v_{U,t}^2\,\mathrm{d}t, \quad v_{U,t}^2 := 2\sigma_t^2(1 - \rho_t), \tag{3}$$

$$\mathrm{d}V_t = \left(2r - \sigma_t^2\right)\mathrm{d}t + \sigma_t\left(\mathrm{d}W_{1,t} + \mathrm{d}W_{2,t}\right), \quad \mathrm{d}\langle V\rangle_t = v_{V,t}^2\,\mathrm{d}t, \quad v_{V,t}^2 := 2\sigma_t^2(1 + \rho_t). \tag{4}$$

In particular, U *is a driftless local martingale, and the* variance budget

$$v_{U,t}^2 + v_{V,t}^2 = 4\sigma_t^2 \tag{5}$$

holds identically.

Proof. Applying Itô's formula to $X_i = \ln S_i$ gives

$$\mathrm{d}X_{i,t} = \left(r - \tfrac{1}{2}\sigma_{i,t}^2\right)\mathrm{d}t + \sigma_{i,t}\,\mathrm{d}W_{i,t}.$$

Hence, $\mathrm{d}\langle X_1, X_1\rangle = \sigma_1^2\mathrm{d}t$, $\mathrm{d}\langle X_2, X_2\rangle = \sigma_2^2\mathrm{d}t$ and $\mathrm{d}\langle X_1, X_2\rangle = \rho\,\sigma_1\sigma_2\,\mathrm{d}t$. Bilinearity of the bracket then gives

$$\mathrm{d}\langle U, V\rangle = \mathrm{d}\langle X_1 - X_2,\, X_1 + X_2\rangle = \mathrm{d}\langle X_1\rangle - \mathrm{d}\langle X_2\rangle = \left(\sigma_1^2 - \sigma_2^2\right)\mathrm{d}t,$$

the two cross terms cancelling identically irrespective ofρ. This establishes the first claim, and the "if and only if" follows since $\sigma_1^2 - \sigma_2^2$ is a predictable process whose time integral vanishes identically exactly when it vanishes almost everywhere.

Now assume $\sigma_1 = \sigma_2 = \sigma$. The drift of $U = X_1 - X_2$ is then $\left(r - \frac{1}{2}\sigma^2\right) - \left(r - \frac{1}{2}\sigma^2\right) = 0$, which yields (3), and the drift of $V = X_1 + X_2$ is $2r - \sigma^2$, which yields (4). The variance rates are

$$\mathrm{d}\langle U\rangle = \left(\sigma^2 - 2\rho\sigma^2 + \sigma^2\right)\mathrm{d}t = 2\sigma^2(1-\rho)\,\mathrm{d}t, \qquad \mathrm{d}\langle V\rangle = 2\sigma^2(1+\rho)\,\mathrm{d}t,$$

and they sum to $4\sigma^2\,\mathrm{d}t$. □

Although Lemma 1 is elementary, it carries the entire structure of this paper, so a few comments are in order.

Remark 1 (The restriction is on volatilities, not correlation)**.** *Orthogonality requires $\sigma_1 = \sigma_2$ at each time and state, not merely similar historical or option-implied volatilities. The common volatility may vary with both prices, so this is compatible with non-Black–Scholes marginal laws, but two arbitrary single-name volatility surfaces cannot be imposed independently. A common instantaneous variance may be a useful approximation for closely matched sector exposures or two similarly constructed portfolios after checking their variance differential. Sector membership or dual listing alone does not establish it; dual-listed prices may also have currency, liquidity and trading-hour effects. The restriction is less plausible for a low-volatility/high-volatility pair or commodity maturities with a pronounced Samuelson effect. Different carry rates, dividends, jumps and materially unequal instantaneous variances also change the ratio law and are outside the present closed-form specification. Therefore, the BD family is a structural benchmark for suitable pairs, not a universal joint calibration model.*

Remark 2 (U has no drift)**.** *The vanishing of the drift of U in (3) is an identity, not an assumption: the convexity corrections $-\frac{1}{2}\sigma_i^2$ of the two log prices are equal and cancel. It is what makes the exit problem for U scale-free, and it is why the scale function of U is the identity, so that the exit* probabilities *of an interval coincide with those of a Brownian motion whatever the local correlation. This statement presumes that the interval is exited almost surely, as is guaranteed, for example, by a variance rate bounded below by a positive constant. On the other hand, zero drift alone does not guarantee a finite exit time.*

Proposition 1 (Admissibility-free parameterisation)**.** *The map*

$$\Lambda : (0,\infty)\times(-1,1) \longrightarrow (0,\infty)^2, \qquad \Lambda(\sigma,\rho) = \left(\sqrt{2\sigma^2(1-\rho)},\, \sqrt{2\sigma^2(1+\rho)}\right) = (v_U, v_V),$$

is a bijection, with inverse

$$\sigma = \tfrac{1}{2}\sqrt{v_U^2 + v_V^2}, \qquad \rho = \frac{v_V^2 - v_U^2}{v_V^2 + v_U^2}. \tag{6}$$

Proof. Λ is well defined because $1 \mp \rho > 0$ on $(-1, 1)$. Given $(v_U, v_V) \in (0, \infty)^2$, Formulas (6) return $\sigma > 0$ and, since $v_U^2 > 0$ and $v_V^2 > 0$,

$$-1 < \frac{v_V^2 - v_U^2}{v_V^2 + v_U^2} < 1,$$

so $\rho \in (-1, 1)$. Substituting (6) into Λ recovers (v_U, v_V), and substituting $\Lambda(\sigma, \rho)$ into (6) recovers (σ, ρ); each is a one-line computation. Hence, Λ is a bijection. □

The practical content of Proposition 1 is that (v_U, v_V) can be specified with complete freedom. Every pair of strictly positive processes is realised by a legitimate correlation, and the constraint $|\rho| \leq 1$ can never bind. This is the parameterisation in which we shall design. This is pointwise covariance admissibility. It does not by itself ensure existence, uniqueness, non-explosion, or true martingales for an arbitrary state-dependent specification; those issues are treated in Assumption 1 and the sufficient conditions below.

2.3. Feasibility of Pathwise Design

The distinction between marginal and pathwise matching matters here. By Gyöngy's mimicking theorem, under suitable conditions, one can construct a one-dimensional Markov diffusion with the same one-time marginal laws as a given Itô process [12]; the Markovian-projection approach of Piterbarg [13] exploits this fact for the calibration of European options. Such a projection preserves expectations of functions of Y_T, but in general it does not preserve the joint law of $(Y_t)_{t \leq T}$ and hence not its hitting times, extrema, barriers, or lookbacks. For these claims, we need the stronger pathwise identity

$$\mathrm{d}\langle Y \rangle_t = \sigma_Y^2(t, Y_t)\,\mathrm{d}t \qquad \text{along the original multi-variate paths.} \tag{7}$$

The feasibility question below concerns this stronger condition, not Markovian projection.

We now formalise the resulting obstruction. It is stated for a general one-homogeneous functional, which covers spreads, baskets and weighted linear combinations.

Proposition 2 (Attainable instantaneous volatility)**.** *Let $F \in C^2((0, \infty)^2)$ be positively homogeneous of degree one, write $F_i = \partial_i F$, and set $Y_t = F(S_{1,t}, S_{2,t})$. Suppose that the marginals in (1) have constant volatilities σ_1, σ_2. Define*

$$a(s_1, s_2) = \sigma_1 s_1 F_1(s_1, s_2), \qquad b(s_1, s_2) = \sigma_2 s_2 F_2(s_1, s_2).$$

Then:

(i) as ρ ranges over $[-1, 1]$, the attainable instantaneous volatility of Y at (s_1, s_2) is exactly

$$\mathcal{U}(s_1, s_2) = \Big[\big||a| - |b|\big|,\ |a| + |b|\Big];$$

(ii) the interval is homogeneous of degree one: $\mathcal{U}(\lambda s_1, \lambda s_2) = \lambda \mathcal{U}(s_1, s_2)$;
(iii) if $\sigma_1 = \sigma_2 = \sigma$ and $F(s_1, s_2) = \alpha_1 s_1 + \alpha_2 s_2$ with $\alpha_1 \alpha_2 \neq 0$, then for every y in the range of F,

$$\bigcap_{\{F = y\}} \mathcal{U}(s_1, s_2) = \{\sigma |y|\}. \tag{8}$$

The value $\sigma|y|$ is generated by $\rho = 1$.

Consequently, for a non-degenerate linear spread or basket with fixed equal marginal volatility, a pathwise target of the form $\sigma_Y = \sigma_Y(t, Y_t)$ that must hold at every state in every level set can only have $\sigma_Y(t, y) = \sigma|y|$, corresponding to the comonotonic limit $\rho \equiv 1$.

Proof. By Itô's formula, the diffusion term of Y is

$$a\,\mathrm{d}W_1 + b\,\mathrm{d}W_2,$$

so its variance rate is $a^2 + b^2 + 2\rho ab$. As a function of ρ, this quantity is affine, with endpoint values $(a+b)^2$ and $(a-b)^2$. Taking square roots yields (i). Because the derivatives of a degree-one homogeneous function are homogeneous of degree zero, a and b are homogeneous of degree one, which proves (ii).

For (iii), $a = \sigma\alpha_1 s_1$ and $b = \sigma\alpha_2 s_2$. At $\rho = 1$, the volatility is

$$|a + b| = \sigma|\alpha_1 s_1 + \alpha_2 s_2| = \sigma|y|$$

at every point of the level set, so $\sigma|y|$ lies in the intersection. To see that no other value can lie in it, approach the boundary of the positive orthant along the level set. For a basket-type combination, one asset can be sent to zero while the other converges to the value required by $F = y$, and both endpoints of $\mathcal{U}$ then converge to $\sigma|y|$. For a spread-type combination, the same argument applies, the asset whose coefficient has the sign of y being the one that converges while the other is sent to zero; when $y = 0$, the level set is a ray, and scaling the ray to the origin shrinks $\mathcal{U}$ to $\{0\}$. The intersection is therefore the singleton in (8). Since $ab \neq 0$ in the interior, the value $|a + b|$ is generated uniquely by $\rho = 1$ there. □

Remark 3 (Why admissibility should be built in)**.** *The proposition provides a simple diagnostic for proposed one-dimensional targets. For instance, take $Y = S_1 - S_2$ with equal constant marginal volatility σ and suppose one attempts to impose*

$$\mathrm{d}Y_t = rY_t\,\mathrm{d}t + \sigma_3\,\mathrm{d}B_t, \qquad c := \sigma_3/\sigma.$$

Variance matching would require

$$\rho^*(s_1, s_2) = \frac{s_1^2 + s_2^2 - c^2}{2s_1 s_2}.$$

The condition $|\rho^*| \leq 1$ is equivalent to

$$|s_1 - s_2| \leq c \leq s_1 + s_2.$$

Hence, the target can hold only on a proper subset of the state space. Because the intended Gaussian spread has unbounded support at every positive horizon, it cannot remain in this feasible set almost surely. Clipping ρ^* to $[-1, 1]$ restores an admissible correlation but alters the spread variance and hence the target law. This is exactly the kind of direct inversion used as a design device for European claims in Escobar-Anel [21]; here, the example serves to isolate the pathwise admissibility issue. The lesson is not specific to that construction: admissibility is safest when the parameterisation itself enforces it.

Remark 4 (What homogeneity does and does not imply)**.** *Part (ii) shows that scale matters: a target that ignores the scale of a degree-one functional is unlikely to be globally compatible with fixed marginal volatilities. Homogeneity of the target is thus a useful necessary design clue, but for a general non-linear F it is not sufficient. The strong singleton conclusion of part (iii) is specific to linear spreads and baskets and should not be read as a statement about every one-homogeneous functional.*

2.4. A Rigidity Result

The next proposition explains why one cannot have everything simultaneously.

Proposition 3 (Rigidity). *Assume $\sigma_1 = \sigma_2 = \sigma$ with σ constant, so that both marginals are geometric Brownian motions with volatility σ. Then*

$$\langle U\rangle_t + \langle V\rangle_t = 4\sigma^2 t \qquad \text{for all } t,\ \mathbb{Q}\text{-a.s.}$$

Consequently, U is a Brownian motion and V is an arithmetic Brownian motion with deterministic drift and constant variance rate if and only if ρ is constant. More generally, if U is required to be a Brownian motion with variance rate σ_U^2, then $\rho \equiv 1 - \sigma_U^2/(2\sigma^2)$ is constant.

Proof. Integrating the variance budget (5) yields $\langle U\rangle_t + \langle V\rangle_t = 4\int_0^t \sigma^2 \mathrm{d}s = 4\sigma^2 t$. If U is a Brownian motion with variance rate σ_U^2, then $\langle U\rangle_t = \sigma_U^2 t$, so by (3) $2\sigma^2(1-\rho_t) = \sigma_U^2$ for a.e. t, that is, $\rho_t = 1 - \sigma_U^2/(2\sigma^2)$, a constant. The converse is immediate. □

Proposition 3 identifies the trade-off behind the BD construction. With a common constant marginal volatility, making the log-ratio a Brownian motion with a fixed variance rate pins the correlation to a constant. To retain a genuinely state-dependent correlation while keeping an exact Brownian ratio, the common marginal volatility must be allowed to vary with the state. This is the route we take below.

3. The Brownian-Dispersion Model

3.1. Definition and Canonical Representation

Definition 1 (BD model). *Fix a constant $\sigma_U > 0$ and a Borel function*

$$\rho : [0, T^*] \times (0,\infty)^2 \longrightarrow (-1, 1).$$

A Brownian-dispersion *(BD) model is any well-posed solution of (1) with*

$$\sigma_{1,t} = \sigma_{2,t} = \sigma(t, S_{1,t}, S_{2,t}) := \frac{\sigma_U}{\sqrt{2\big(1-\rho(t, S_{1,t}, S_{2,t})\big)}}, \qquad \rho_t = \rho(t, S_{1,t}, S_{2,t}). \tag{9}$$

The finite-horizon martingale and integrability requirements used for pricing are imposed explicitly in Assumption 1 below.

Definition 1 lets the correlation approach, but not attain, ±1. This is useful for index-dependent specifications such as those of Section 5. The canonical representation below is the constructive formulation: it parametrises the same model through a strictly positive index-volatility process ν and makes the correlation bound automatic.

Theorem 1 (Canonical representation). *Let β^U and β^V be independent standard Brownian motions and define*

$$U_t = u_0 + \sigma_U \beta_t^U, \qquad \frac{\mathrm{d}P_t}{P_t} = \Big(r - \frac{\sigma_U^2}{8}\Big)\mathrm{d}t + \nu_t\, \mathrm{d}\beta_t^V, \qquad P_0 = \sqrt{S_{1,0}S_{2,0}}, \quad u_0 = \ln\frac{S_{1,0}}{S_{2,0}}, \tag{10}$$

where ν is a strictly positive, locally square integrable predictable process. Set

$$S_{1,t} := P_t\, e^{U_t/2}, \qquad S_{2,t} := P_t\, e^{-U_t/2}. \tag{11}$$

Then:

(i) (S_1, S_2) solves (1) with a common volatility and a correlation given by

$$\sigma_t = \sqrt{\nu_t^2 + \frac{\sigma_U^2}{4}}, \qquad \rho_t = \frac{4\nu_t^2 - \sigma_U^2}{4\nu_t^2 + \sigma_U^2} \in (-1, 1); \tag{12}$$

(ii) conversely, in the BD model of Definition 1, the processes $U = \ln(S_1/S_2)$ *and* $P = \sqrt{S_1 S_2}$ *satisfy* (10) *with*

$$\nu_t = \frac{\sigma_U}{2}\sqrt{\frac{1+\rho_t}{1-\rho_t}}, \tag{13}$$

and the driving Brownian motions

$$\beta_t^U = \int_0^t \frac{\mathrm{d}W_{1,s} - \mathrm{d}W_{2,s}}{\sqrt{2(1-\rho_s)}}, \qquad \beta_t^V = \int_0^t \frac{\mathrm{d}W_{1,s} + \mathrm{d}W_{2,s}}{\sqrt{2(1+\rho_s)}}$$

are independent;

(iii) if, for every finite T,

$$\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^T \nu_t^2\,\mathrm{d}t\right)\right] < \infty, \tag{14}$$

then $e^{-rt}S_{i,t}$, $i = 1, 2$, *are true* $\mathbb{Q}$*-martingales on finite horizons;*

(iv) if $\nu_t = \nu(t, U_t)$ *for a positive, locally bounded Borel function, then* (10)–(11) *define* (S_1, S_2) *explicitly as a functional of* (β^U, β^V)*; by (ii), every solution of* (1)–(9) *admits this representation, so the BD model is unique in law. The corresponding local correlation follows from* (12).

Proof. (i) Assume the BD model. By Lemma 1, U and V are orthogonal, $\mathrm{d}U = v_U\,\mathrm{d}\beta^U$ with $v_U^2 = 2\sigma^2(1-\rho)$, and $\mathrm{d}V = (2r - \sigma^2)\mathrm{d}t + v_V\,\mathrm{d}\beta^V$ with $v_V^2 = 2\sigma^2(1+\rho)$, where

$$\beta^U := \int_0^\cdot \frac{\mathrm{d}W_{1,s} - \mathrm{d}W_{2,s}}{\sqrt{2(1-\rho_s)}}, \qquad \beta^V := \int_0^\cdot \frac{\mathrm{d}W_{1,s} + \mathrm{d}W_{2,s}}{\sqrt{2(1+\rho_s)}}.$$

Each is a continuous local martingale with $\langle\beta^U\rangle_t = \langle\beta^V\rangle_t = t$ and therefore a Brownian motion by Lévy's theorem; moreover, $\langle\beta^U, \beta^V\rangle \equiv 0$ by Lemma 1, so a two-dimensional application of Lévy's theorem shows that (β^U, β^V) is a two-dimensional Brownian motion, and its components are independent.

Substituting the BD specification (9),

$$v_{U,t}^2 = 2\sigma_t^2(1-\rho_t) = 2\cdot\frac{\sigma_U^2}{2(1-\rho_t)}\cdot(1-\rho_t) = \sigma_U^2,$$

a constant, so U is a driftless Brownian motion with variance rate σ_U^2. Similarly

$$v_{V,t}^2 = 2\sigma_t^2(1+\rho_t) = \sigma_U^2\,\frac{1+\rho_t}{1-\rho_t},$$

which is (13) once we set $\nu := v_V/2$. It remains to compute the dynamics of $P = e^{V/2}$. By Itô's formula and (4),

$$\frac{\mathrm{d}P_t}{P_t} = \frac{1}{2}\,\mathrm{d}V_t + \frac{1}{8}\,\mathrm{d}\langle V\rangle_t = \frac{1}{2}\big(2r - \sigma_t^2\big)\mathrm{d}t + \frac{v_{V,t}^2}{8}\mathrm{d}t + \frac{v_{V,t}}{2}\,\mathrm{d}\beta_t^V.$$

Using the variance budget (5) in the form $\sigma^2 = \frac{1}{4}\big(\sigma_U^2 + v_V^2\big)$, the drift becomes

$$r - \frac{\sigma_U^2 + v_{V,t}^2}{8} + \frac{v_{V,t}^2}{8} = r - \frac{\sigma_U^2}{8},$$

a constant, and the diffusion coefficient is $v_V/2 = \nu$. This is precisely (10).

(ii) Conversely, start from (10). By (11), $\ln S_1 = \ln P + \frac{1}{2}U$ and $\ln S_2 = \ln P - \frac{1}{2}U$, so

$$\mathrm{d}\ln S_{1,t} = \Big(r - \frac{\sigma_U^2}{8} - \frac{\nu_t^2}{2}\Big)\mathrm{d}t + \nu_t\,\mathrm{d}\beta_t^V + \frac{\sigma_U}{2}\,\mathrm{d}\beta_t^U,$$

with the analogous expression for $\ln S_2$ carrying $-\sigma_U/2$. Because $\beta^U \perp \beta^V$,

$$\mathrm{d}\langle \ln S_1\rangle_t = \mathrm{d}\langle \ln S_2\rangle_t = \Big(\nu_t^2 + \frac{\sigma_U^2}{4}\Big)\mathrm{d}t =: \sigma_t^2\,\mathrm{d}t, \qquad \mathrm{d}\langle \ln S_1, \ln S_2\rangle_t = \Big(\nu_t^2 - \frac{\sigma_U^2}{4}\Big)\mathrm{d}t,$$

which gives the common volatility in (12) and the correlation

$$\rho_t = \frac{\nu_t^2 - \sigma_U^2/4}{\nu_t^2 + \sigma_U^2/4} = \frac{4\nu_t^2 - \sigma_U^2}{4\nu_t^2 + \sigma_U^2} \in (-1, 1).$$

Finally , returning to prices,

$$\frac{\mathrm{d}S_{1,t}}{S_{1,t}} = \mathrm{d}\ln S_{1,t} + \frac{1}{2}\mathrm{d}\langle \ln S_1\rangle_t = \Big(r - \frac{\sigma_U^2}{8} - \frac{\nu_t^2}{2} + \frac{\nu_t^2}{2} + \frac{\sigma_U^2}{8}\Big)\mathrm{d}t + \cdots = r\,\mathrm{d}t + \sigma_t\,\mathrm{d}W_{1,t},$$

where the Brownian motions are defined differentially by

$$\mathrm{d}W_{1,t} = \frac{\nu_t\,\mathrm{d}\beta_t^V + (\sigma_U/2)\,\mathrm{d}\beta_t^U}{\sigma_t}, \qquad \mathrm{d}W_{2,t} = \frac{\nu_t\,\mathrm{d}\beta_t^V - (\sigma_U/2)\,\mathrm{d}\beta_t^U}{\sigma_t}.$$

Each has unit quadratic-variation rate, and their covariation rate is ρ_t in (12); Lévy's theorem therefore provides the required Brownian representation. The drift is exactly r, since the two convexity corrections cancel.

(iii) By (i), $e^{-rt}S_{i,t} = S_{i,0}\,\mathcal{E}(\int \sigma \mathrm{d}W_i)_t$. Since $\sigma_t^2 = \nu_t^2 + \sigma_U^2/4$, condition (14) implies, for every finite T,

$$\mathbb{E}\Big[\exp\Big(\frac{1}{2}\int_0^T \sigma_t^2\,\mathrm{d}t\Big)\Big] = e^{\sigma_U^2 T/8}\mathbb{E}\Big[\exp\Big(\frac{1}{2}\int_0^T \nu_t^2\,\mathrm{d}t\Big)\Big] < \infty.$$

Novikov's condition therefore makes both discounted asset prices true martingales on finite horizons.

(iv) If $\nu_t = \nu(t, U_t)$, then U is given explicitly as a Brownian motion and

$$P_t = P_0 \exp\Big(\Big(r - \frac{\sigma_U^2}{8}\Big)t - \frac{1}{2}\int_0^t \nu^2(s, U_s)\,\mathrm{d}s + \int_0^t \nu(s, U_s)\,\mathrm{d}\beta_s^V\Big)$$

is an explicit functional of the two independent Brownian paths. Strong existence and pathwise uniqueness relative to (β^U, β^V) are therefore immediate; since (ii) maps every solution of (1)–(9) to this representation, the model is unique in law. □

Remark 5 (Variance decomposition). *Equation (12) is a variance decomposition of a familiar type: individual variance = index variance + dispersion variance,*

$$\sigma^2 = \nu^2 + \frac{\sigma_U^2}{4}, \qquad \textit{covariance} = \rho\sigma^2 = \nu^2 - \frac{\sigma_U^2}{4},$$

so that correlation is high precisely when index volatility dominates dispersion. This mechanical comovement is consistent with empirical evidence that conditional equity correlations rise with market volatility [27] and with the economic importance of correlation risk documented in option markets [28]. Note also the identity $\sigma_U^2 = \sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2$: the dispersion volatility σ_U is exactly the Margrabe volatility of the pair, and it is constant *in the BD model by construction.*

Remark 6 (Where the local structure lives). *Three specifications of ν will recur:*

(a) $\nu_t = \nu(t, U_t)$, *equivalently* $\rho = \rho(t, S_1/S_2)$: *the correlation depends on relative performance and is constant along rays through the origin of the price plane;*

(b) $\nu_t = \nu(t, P_t)$, *equivalently* $\rho = \rho(t, \sqrt{S_1 S_2})$: *the correlation depends on the index level and is constant along hyperbolas $s_1 s_2 = const$. This case yields exact independence of the two coordinates and is developed in Section 5;*

(c) $\nu_t = \nu(t, U_t, P_t)$, *the general Markovian case.*

Theorem 2 below covers all three. Only in case (b) is P autonomous.

3.2. The Invariance Theorem

We now turn to the pricing engine. Let $\mathcal{F}^U = \{\mathcal{F}_t^U\}$ denote the $\mathbb{Q}$-augmented filtration generated by U, equivalently by β^U, equivalently by the ratio process $R = e^U$.

Assumption 1. *In the canonical construction, the filtration is the usual augmentation of that generated by (β^U, β^V), with deterministic positive initial asset values. The process ν is predictable, strictly positive and satisfies $\int_0^T \nu_t^2 \, \mathrm{d}t < \infty$ almost surely for every finite $T \le T^*$. If specified through the state, the canonical system has a unique global strong solution. For every such T,*

$$\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^T \nu_t^2 \, \mathrm{d}t\right)\right] < \infty. \tag{15}$$

The natural-filtration condition makes the conditioning argument below explicit; more general filtrations require that β^V remain a Brownian motion after the entire U path is revealed. No positive lower bound on ν is needed for the invariance identity.

Primitive sufficient conditions.

For the Markovian cases, write $Y = \log P$, $c = r - \sigma_U^2/8$ and $n(t,u,y) = \nu(t,u,e^y)$. A convenient sufficient set of conditions is as follows: n is Borel measurable in t, locally Lipschitz in (u,y) uniformly on compact time intervals, and, for each finite T, $0 < n \le C_T < \infty$ on $[0,T] \times \mathbb{R}^2$. Then

$$\mathrm{d}U = \sigma_U \, \mathrm{d}\beta^U, \qquad \mathrm{d}Y = \left(c - \tfrac{1}{2}n^2(t,U,Y)\right)\mathrm{d}t + n(t,U,Y)\mathrm{d}\beta^V$$

has locally Lipschitz drift and diffusion and bounded coefficients on each finite time interval. The standard local-existence/pathwise-uniqueness argument, followed by the bounded-coefficient non-explosion estimate, gives a unique global strong solution [29]. Both log coordinates remain finite; hence, P, S_1, S_2 are strictly positive and cannot explode or hit zero in finite time. Moreover, (15) is bounded by $e^{C_T^2 T/2}$, so the index stochastic exponential and both discounted assets are true martingales. These are sufficient, not necessary, conditions. In the ratio-only case $\nu(t, U_t)$, local boundedness suffices for the explicit strong construction in Theorem 1(iv); the exponential integrability requirement must still be checked separately. The bounded smooth ratio specifications and the bounded, locally Lipschitz clipped index specification used in Section 6 satisfy the displayed sufficient conditions.

Theorem 2 (Index invariance)**.** *Let Assumption 1 hold and set*

$$\alpha^* := \frac{\sigma_U^2}{8}.$$

Let τ be an almost surely finite $\mathcal{F}^U$-stopping time and let G be $\mathcal{F}_\tau^U$-measurable with $\mathbb{E}[e^{-\alpha^*\tau}|G|] < \infty$. Then

$$\mathbb{E}[e^{-r\tau}P_\tau G] = P_0 \, \mathbb{E}[e^{-\alpha^*\tau}G]. \tag{16}$$

In particular, once σ_U is fixed, the value does not depend on the local-correlation specification.

Proof. Set $c = r - \alpha^*$ and

$$Z_t := \frac{e^{-ct}P_t}{P_0} = \exp\left\{\int_0^t \nu_s \, \mathrm{d}\beta_s^V - \frac{1}{2}\int_0^t \nu_s^2 \, \mathrm{d}s\right\}.$$

This is a positive local martingale, and (15) makes it a true martingale on every fixed finite horizon. We need the stronger conditional statement at a ratio stopping time, not just its unconditional expectation at a deterministic time.

Let $\mathcal{H} = \sigma(\beta_s^U : 0 \le s \le T^*)$ and condition on an entire U path using a regular conditional probability on the Brownian path space. Independence of the two driving Brownian motions and the filtration condition in Assumption 1 imply that β^V is still Brownian in the conditional filtration. The originally predictable ν remains non-anticipating in β^V; any dependence on U is now a fixed path input. Put $t_m = m \wedge T^*$. For every integer $m \ge 1$,

$$\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^{t_m} \nu_s^2 \, ds\right)\middle| \mathcal{H}\right] < \infty \quad \text{almost surely.}$$

Indeed, the unconditional expectation is finite. Taking the intersection of these probability-one sets over the countably many integers ensures that conditional Novikov holds on *all* these horizons for the same set of U paths. Thus, Z is a mean-one martingale on each $[0, t_m]$ under almost every conditional probability.

Because τ is an $\mathcal{F}^U$-stopping time, revealing $\mathcal{H}$ fixes τ to a finite deterministic value $t(u)$. Under that conditional probability, apply bounded optional sampling to $\tau \wedge t_m$; on every path with $t(u) \le t_m$, this is just evaluation at $t(u)$. Choose an integer m with $t_m \ge t(u)$ using the simultaneous conditional martingale property established above. This yields

$$\mathbb{E}[Z_\tau \mid \mathcal{H}] = 1 \qquad \mathbb{Q} \text{ almost surely.} \tag{17}$$

There is no optional-stopping assertion at an unbounded stopping time under the original measure and no passage to Z_∞; uniform integrability over an infinite horizon is not being assumed.

Finally, $e^{-\alpha^* \tau} G$ is $\mathcal{H}$-measurable. Conditional Tonelli first proves

$$\mathbb{E}[e^{-r\tau} P_\tau |G|] = P_0 \mathbb{E}[e^{-\alpha^* \tau} |G|] < \infty.$$

The tower property can therefore be applied to the signed payoff, and (17) gives (16). □

Primitive payoff and stopping-time integrability.

The condition on G concerns the ratio payoff and cannot follow from a restriction on ν alone. Useful sufficient cases are:

- $\tau \le T$ and $|G| \le C \exp\{q \sup_{t \le T} |U_t|\}$ for fixed $C, q < \infty$; finite-horizon Brownian maxima have all such exponential moments. This covers the exchange and lookback payoffs below.
- Any almost surely finite ratio stopping time with bounded G; then $e^{-\alpha^* \tau} \le 1$, so no moment of τ is needed.
- The first exit τ from a bounded ratio-log interval of width ℓ, and $|G| \le C(1 + \tau^q) e^{b\tau}$ with $q \ge 0$ and

$$b - \alpha^* < \lambda_1, \qquad \lambda_1 = \frac{\pi^2 \sigma_U^2}{2\ell^2}. \tag{18}$$

  The killed Brownian kernel gives an exit-time tail of order $e^{-\lambda_1 t}$, hence finiteness under the strict inequality. In particular, interval exits are almost surely finite and have every polynomial moment; bounded boundary rebates are covered even though their stopping times are unbounded.

Together with, for example, the finite-horizon bounds on ν above, these conditions justify the pricing identity without an infinite-horizon Novikov condition.

Corollary 1 (Cash-settled ratio claims)**.** *Let τ be an almost surely finite $\mathcal{F}^U$-stopping time and let G be $\mathcal{F}_\tau^U$-measurable with $\mathbb{E}[e^{-r\tau}|G|] < \infty$. Then, $\mathbb{E}[e^{-r\tau} G]$ is computed with $U = u_0 + \sigma_U \beta^U$ and does not depend on the local-correlation specification.*

Proof. By Theorem 1(ii), $U = u_0 + \sigma_U \beta^U$ under $\mathbb{Q}$ in every BD specification. Since τ and G are functionals of the path of U, the joint law of (τ, G) does not involve ν. □

Remark 7 (Change-of-numeraire interpretation). *For bounded stopping times, Theorem 2 also follows from the classical change-of-numeraire and scale-invariance mechanism [4–6]. Under the S_2 numeraire, $R = S_1/S_2$ is a martingale; deterministic relative quadratic variation makes* log *R a Brownian motion with the usual log drift, and an exponential tilt removes that drift. We retain the conditional proof above because it also covers almost surely finite, unbounded ratio stopping times.*

Remark 8 (What equal volatilities buy). *With constant relative variance but unequal instantaneous volatilities, Itô's formula and Lévy's characterisation give*

$$\mathrm{d}U_t = \tfrac{1}{2}(\sigma_{2,t}^2 - \sigma_{1,t}^2)\mathrm{d}t + \sigma_U \,\mathrm{d}B_t^{\mathbb{Q}}. \tag{19}$$

Thus, constant relative variance fixes the Brownian noise, *not the $\mathbb{Q}$-drift. A driftless Brownian log-ratio under $\mathbb{Q}$ requires $\sigma_1^2 = \sigma_2^2$ almost everywhere. If the variance differential is a non-zero constant, U is instead an arithmetic Brownian motion with that constant drift divided by two; cash-settled ratio prices then depend on this additional drift. With a state-dependent differential, even that fixed-drift Brownian law is lost.*

For bounded ratio stopping times, the degree-one prices nevertheless remain the same under constant relative variance and true asset martingales: under the S_2 numeraire, $\mathrm{d}U = -\sigma_U^2 \mathrm{d}t/2 + \sigma_U \,\mathrm{d}B^{(2)}$, independently of the variance differential. Equivalently, the reduced formula is

$$\mathbb{E}_{\mathbb{Q}}[e^{-r\tau} P_\tau G] = P_0 \mathbb{E}^{(0)}[e^{-\alpha^* \tau} G],$$

where $\mathbb{E}^{(0)}$ evaluates the same stopping rule and payoff under the driftless *Brownian ratio law. Outside BD, this is generally not the original $\mathbb{Q}$-law of U. What equal volatilities buy is invariance of the cash book under $\mathbb{Q}$ as well as the degree-one book and the orthogonal structure used in Section 5.*

Remark 9 (Local dispersion). *The proof of Theorem 2 uses the constancy of σ_U only through the drift of P. Suppose σ_U in (9) is replaced by $\sigma_U(t, U_t)$ for a positive function such that $\mathrm{d}U_t = \sigma_U(t, U_t)\,\mathrm{d}\beta_t^U$ is well posed. The computation in the proof of Theorem 1(ii) then gives $\mathrm{d}P_t/P_t = \big(r - \sigma_U^2(t, U_t)/8\big)\mathrm{d}t + \nu_t \,\mathrm{d}\beta_t^V$, and the same conditioning argument yields*

$$\mathbb{E}[e^{-r\tau} P_\tau G] = P_0 \,\mathbb{E}\Big[\exp\Big(-\frac{1}{8}\int_0^\tau \sigma_U^2(s, U_s)\,\mathrm{d}s\Big) G\Big],$$

with U now a driftless one-dimensional local-volatility diffusion under the obvious modification of the integrability requirement. Corollary 1 is unchanged. Invariance to the index-variance allocation therefore survives a state-dependent dispersion volatility; only the one-dimensional problem for U becomes a Feynman–Kac problem with potential $\sigma_U^2/8$. We keep σ_U constant below because it delivers closed forms.

Remark 10 (Identification and observational equivalence). *Fix σ_U. Every claim in the invariant classes of Theorem 2 and Corollary 1 has the same price as in any constant-parameter Black–Scholes–Margrabe model with equal marginal volatilities $\bar{\sigma}$ and constant correlation $\bar{\rho}$ satisfying $2\bar{\sigma}^2(1 - \bar{\rho}) = \sigma_U^2$. This observational equivalence is intentional: the relative-performance book identifies the dispersion variance, not the orthogonal index-variance allocation. The BD models remain distinguishable through claims outside the invariant class, including fixed-strike baskets and, in general, the single-name marginals.*

Remark 11 (What the theorem does and does not cover)**.** *The invariant class of Theorem 2 is* $\{P_\tau G : G \in \mathcal{F}_\tau^U\}$*, and that of Corollary 1 is its degree-zero counterpart* $\{G : G \in \mathcal{F}_\tau^U\}$*. Translated back into prices, for the former, a claim qualifies when its payoff is the product of a function of the terminal prices that is* positively homogeneous of degree one *and a functional of the* ratio path*. Indeed, if* π *is one-homogeneous, then by* (11)*,*

$$\pi\big(S_{1,t}, S_{2,t}\big) = \pi\big(P_t e^{U_t/2}, P_t e^{-U_t/2}\big) = P_t\, \widehat{\pi}(U_t), \qquad \widehat{\pi}(u) := \pi\big(e^{u/2}, e^{-u/2}\big), \tag{20}$$

and $\widehat{\pi}(U_\tau)$ *is* $\mathcal{F}_\tau^U$*-measurable. Claims carrying a power of P other than zero or one—anything with a fixed strike in currency units, such as* $(S_{1,T} + S_{2,T} - K)^+$*—are* not *invariant in general and do depend on* $\rho(\cdot)$*: for* $\gamma \notin \{0, 1\}$*,* P_τ^γ *carries the factor* $\exp\{\frac{1}{2}(\gamma^2 - \gamma)\int_0^\tau \nu_t^2\, \mathrm{d}t\}$ *relative to a stochastic exponential. This is not a defect but a decomposition of the calibration problem: relative-performance products identify* σ_U*, while fixed-strike basket and other index-sensitive products constrain* ν *and hence* $\rho(\cdot)$*, without generally identifying an unrestricted surface uniquely.*

Corollary 2 (Synthetic Black–Scholes reduction)**.** *Let* Φ *be a payoff functional of the ratio path* $\{R_t\}_{t\le T}$*,* $R = e^U$*, with* $\mathbb{E}[e^{-\alpha^* T}|\Phi|] < \infty$*, and consider the claim* $\Pi_T = P_T\, \Phi\big(\{R_t\}_{t\le T}\big)$*. Its time-0 value in the BD model is then*

$$\Pi_0 = P_0 \cdot \mathrm{BS}_{\hat{r},\hat{q},\,\sigma_U}[\Phi],$$

where $\mathrm{BS}_{\hat{r},\hat{q},\sigma_U}[\Phi]$ *denotes the Black–Scholes value of the claim* Φ *written on a* single *asset with spot* $R_0 = S_{1,0}/S_{2,0}$*, volatility* σ_U*, constant interest rate and continuous dividend; this yields*

$$\hat{r} = \frac{\sigma_U^2}{8}, \qquad \hat{q} = \hat{r} - \frac{\sigma_U^2}{2} = -\frac{3\sigma_U^2}{8}.$$

For a one-homogeneous terminal payoff π*, the precise dictionary is*

$$\Phi(R) = \widehat{\pi}(\log R) = \pi(\sqrt{R}, 1/\sqrt{R}).$$

Thus, for instance, $\pi(s_1, s_2) = (s_1 - \kappa s_2)^+$ *corresponds to* $\Phi(R) = (\sqrt{R} - \kappa/\sqrt{R})^+$*, not to* $(R - \kappa)^+$*.*

Proof. By Theorem 2 with $\tau \equiv T$, $\Pi_0 = P_0 e^{-\alpha^* T}\mathbb{E}[\Phi]$, where the expectation is taken with $\ln R = U$ a driftless Brownian motion with volatility σ_U. In a Black–Scholes market with rate $\hat{r}$ and yield $\hat{q}$, the risk-neutral log-price satisfies $\mathrm{d}\ln R = (\hat{r} - \hat{q} - \frac{1}{2}\sigma_U^2)\mathrm{d}t + \sigma_U \mathrm{d}W$, which is driftless iff $\hat{r} - \hat{q} = \frac{1}{2}\sigma_U^2$, and the discount factor $e^{-\hat{r}T}$ equals $e^{-\alpha^* T}$ iff $\hat{r} = \sigma_U^2/8$. Solving the two equations yields the stated $(\hat{r}, \hat{q})$. □

4. Fixed-Maturity First-Passage Pricing

We retain two representative path-dependent applications and a cash rebate. Fix $0 < L < R_0 < H < \infty$, where $R_0 = S_{1,0}/S_{2,0}$, and set

$$u_- := \ln L, \qquad u_+ := \ln H, \qquad \ell := u_+ - u_-.$$

$$\tau := \inf\{t \ge 0 : U_t \notin (u_-, u_+)\}. \tag{21}$$

For maturity T, write $s = \sigma_U\sqrt{T}$ and

$$\varphi_T(z) = \frac{1}{\sigma_U\sqrt{2\pi T}} \exp\!\left(-\frac{z^2}{2\sigma_U^2 T}\right).$$

Lemma 2 (Gaussian exponential integral)**.** *For* $\theta \in \mathbb{R}$*,* $-\infty \le a < b \le \infty$ *and* $m \in \mathbb{R}$*,*

$$\mathcal{E}_\theta(a, b; m) := \int_a^b e^{\theta z}\varphi_T(z - m)\,\mathrm{d}z = e^{\theta m + \frac{1}{2}\theta^2 s^2}\left[N\!\left(\frac{b - m - \theta s^2}{s}\right) - N\!\left(\frac{a - m - \theta s^2}{s}\right)\right].$$

Proof. Complete the square and change the variable of integration. □

4.1. Knock-Out Claims at a Fixed Maturity

Theorem 3 (Double knock-out)**.** *Let π be continuous and one-homogeneous, and consider the claim paying $\pi(S_{1,T}, S_{2,T})\mathbf{1}_{\{\tau>T\}}$ at T. Its value is*

$$\Pi_0 = P_0\, e^{-\sigma_U^2 T/8} \int_{u_-}^{u_+} \widehat{\pi}(u)\, p_T^{\mathrm{DB}}(u_0, u)\, \mathrm{d}u, \tag{22}$$

where p_T^{DB} is the transition density of a σ_U-Brownian motion killed at $\{u_-, u_+\}$,

$$p_T^{\mathrm{DB}}(u_0, u) = \sum_{n\in\mathbb{Z}} \Big[\varphi_T\big(u - u_0 - 2n\ell\big) - \varphi_T\big(u + u_0 - 2u_- - 2n\ell\big)\Big], \qquad u \in (u_-, u_+). \tag{23}$$

For the single-barrier cases, only two terms are required:

$$p_T^{\mathrm{UO}}(u_0, u) = \varphi_T(u - u_0) - \varphi_T\big(u - (2u_+ - u_0)\big), \qquad p_T^{\mathrm{DO}}(u_0, u) = \varphi_T(u - u_0) - \varphi_T\big(u - (2u_- - u_0)\big).$$

Proof. Since T is a (deterministic, hence $\mathcal{F}^U$-) stopping time and $G := \widehat{\pi}(U_T)\mathbf{1}_{\{\tau>T\}}$ is $\mathcal{F}_T^U$-measurable and integrable (continuity on the closed ratio corridor makes G bounded on survival; it is zero otherwise), Theorem 2 gives

$$\Pi_0 = \mathbb{E}[e^{-rT} P_T\, \widehat{\pi}(U_T)\mathbf{1}_{\{\tau>T\}}] = P_0 e^{-\alpha^* T}\, \mathbb{E}[\widehat{\pi}(U_T)\mathbf{1}_{\{\tau>T\}}],$$

with U a driftless Brownian motion with variance rate σ_U^2 started at u_0. It remains to identify the sub-probability density of U_T on the survival event, which is the classical double-reflection (image) series. For completeness, we verify the boundary conditions. Write $q(u) := p_T^{\mathrm{DB}}(u_0, u)$ as in (23). At $u = u_-$, the first family contributes $\varphi_T(u_- - u_0 - 2n\ell)$ and the second $\varphi_T(u_- + u_0 - 2u_- - 2n\ell) = \varphi_T(u_0 - u_- - 2n\ell) = \varphi_T(u_- - u_0 + 2n\ell)$ by the evenness of φ_T; replacing $n \mapsto -n$ in the second sum shows that the two families cancel term by term, so $q(u_-) = 0$. At $u = u_+$, using $u_+ - 2u_- = 2\ell - u_+$, the second family contributes $\varphi_T(u_+ + u_0 - 2u_- - 2n\ell) = \varphi_T\big(u_0 - u_+ + 2(1-n)\ell\big) = \varphi_T\big(u_+ - u_0 - 2(1-n)\ell\big)$, and the substitution $m = 1-n$ again cancels the first family, so $q(u_+) = 0$. Since each term solves the heat equation $\partial_T q = \frac{1}{2}\sigma_U^2 \partial_{uu} q$ and the series converges absolutely and uniformly on compacts (the shifts grow linearly in n while the Gaussian decays superexponentially), q is the Dirichlet heat kernel on (u_-, u_+) with initial condition δ_{u_0}, that is, the killed transition density. The single-barrier expressions follow from the usual one-sided reflection principle. □

For the exchange payoff, the integral can be computed term by term.

Corollary 3 (Double-knock-out exchange option)**.** *Let $\kappa > 0$ and consider the claim $\big(S_{1,T} - \kappa S_{2,T}\big)^+ \mathbf{1}_{\{\tau>T\}}$. Set $k := \ln\kappa$ and $a := \max(u_-, k)$. If $a \geq u_+$, the value is zero; otherwise*

$$\mathrm{DKO}(\kappa) = P_0\, e^{-s^2/8} \sum_{n\in\mathbb{Z}} \Big[\mathcal{A}\big(u_0 + 2n\ell\big) - \mathcal{A}\big(2u_- - u_0 + 2n\ell\big)\Big], \tag{24}$$

where $s = \sigma_U\sqrt{T}$ and

$$\mathcal{A}(m) := \mathcal{E}_{1/2}\big(a, u_+; m\big) - \kappa\, \mathcal{E}_{-1/2}\big(a, u_+; m\big),$$

with $\mathcal{E}_\theta$ as in Lemma 2. Explicitly,

$$\mathcal{A}(m) = e^{\frac{m}{2}+\frac{s^2}{8}}\left[N\Big(\tfrac{u_+ - m - \frac{s^2}{2}}{s}\Big) - N\Big(\tfrac{a - m - \frac{s^2}{2}}{s}\Big)\right] - \kappa\, e^{-\frac{m}{2}+\frac{s^2}{8}}\left[N\Big(\tfrac{u_+ - m + \frac{s^2}{2}}{s}\Big) - N\Big(\tfrac{a - m + \frac{s^2}{2}}{s}\Big)\right].$$

Proof. Here, $\widehat{\pi}(u) = \left(e^{u/2} - \kappa e^{-u/2}\right)^+$, which is strictly positive iff $e^u > \kappa$, i.e., $u > k$. Hence, the integral in (22) runs over (a, u_+) with $a = \max(u_-, k)$ and integrand $\left(e^{u/2} - \kappa e^{-u/2}\right) p_T^{\mathrm{DB}}(u_0, u)$. Insert (23), interchange sum and integral (justified by absolute convergence), and apply Lemma 2 with $\theta = \frac{1}{2}$ and $\theta = -\frac{1}{2}$ to each shifted Gaussian. Note that $\frac{1}{2}\theta^2 s^2 = s^2/8$ for both values of θ; this is why the common factor $e^{s^2/8}$ appears and partially cancels the discount $e^{-s^2/8}$. □

Certified truncation and narrow corridors.

Let $M = \sup_{u\in[u_-,u_+]} |\widehat{\pi}(u)|$, $\lambda = \ell/s$, and truncate the image sum to $|n| \le N$, $N \ge 1$. For an omitted index $|n| = k$, all Gaussian distances are at least $2(k-1)\ell$ (the reflected family requires this more conservative distance). Summing four absolute Gaussian tails per k and then integrating over the corridor gives the rigorous price bound

$$|\Pi_0 - \Pi_0^{(N)}| \le P_0 e^{-s^2/8} M \, \frac{4\lambda}{\sqrt{2\pi}} \frac{e^{-2N^2\lambda^2}}{1 - e^{-2(2N+1)\lambda^2}}. \tag{25}$$

Here, we used $(N + j)^2 \ge N^2 + (2N + 1)j$ for integers $j \ge 0$. For $N = 4$ and $\lambda \ge 1$, the factor after $P_0 e^{-s^2/8} M$ is at most 2.03×10^{-14}. For the numerical corridor, $\lambda = 1.94203$ and the largest exchange payoff on the corridor occurs at $\kappa = 0.90$; the bound is 4.06×10^{-51} price units. This bounds series truncation, not floating-point cancellation; differences in normal CDFs in the far right tail are evaluated using complementary CDFs.

For a narrow corridor, use the equivalent eigenfunction expansion

$$p_T^{\mathrm{DB}}(u_0, u) = \frac{2}{\ell} \sum_{j=1}^{\infty} \sin\left(\frac{j\pi(u_0 - u_-)}{\ell}\right) \sin\left(\frac{j\pi(u - u_-)}{\ell}\right) e^{-j^2\pi^2 s^2/(2\ell^2)}. \tag{26}$$

With $a = \pi^2/(2\lambda^2)$, its truncation at $j \le J$ satisfies

$$|\Pi_0 - \Pi_0^{[J]}| \le 2P_0 e^{-s^2/8} M \frac{e^{-a(J+1)^2}}{1 - e^{-a(2J+3)}}. \tag{27}$$

The comparison in Table 1 uses a tolerance of 10^{-10} after dividing prices by $P_0 e^{-s^2/8} M$. Each image index contains a difference of two Gaussian integrals; each eigenfunction term contains one sine projection. The certified lengths explain why the spectral form is preferable for narrow corridors, rather than claiming a universal four-image accuracy rule.

Table 1. Smallest cutoffs certified by (25) and (27) for a normalised absolute price error below 10^{-10}. These are term counts, not measured execution times.

ℓ/s	Image Cutoff N	Image Indices $2N+1$	Eigenfunctions J
1.94203	2	5	4
1.00	4	9	2
0.75	5	11	1
0.50	7	15	1

Corollary 4 (Margrabe holds exactly)**.** *In every admissible BD specification, the unconstrained exchange option* $(S_{1,T} - \kappa S_{2,T})^+$ *is worth*

$$S_{1,0} N(d_1) - \kappa S_{2,0} N(d_2), \qquad d_1 = \frac{\ln\left(R_0/\kappa\right) + \frac{1}{2}\sigma_U^2 T}{\sigma_U \sqrt{T}}, \qquad d_2 = d_1 - \sigma_U \sqrt{T},$$

that is, Margrabe [7]'s formula with volatility σ_U.

Proof. Remove the barriers so that $\mathbf{1}_{\{\tau>T\}} \equiv 1$ and $U_T \sim \mathcal{N}(u_0, s^2)$. Theorem 2 gives $\Pi_0 = P_0 e^{-s^2/8}\mathbb{E}[\left(e^{U_T/2} - \kappa e^{-U_T/2}\right)^+]$. By Lemma 2 with $b = \infty$, $a = k = \ln\kappa$,

$$\mathbb{E}[e^{U_T/2}\mathbf{1}_{\{U_T>k\}}] = e^{\frac{u_0}{2}+\frac{s^2}{8}}\,N\Big(\tfrac{u_0-k+\frac{s^2}{2}}{s}\Big), \qquad \mathbb{E}[e^{-U_T/2}\mathbf{1}_{\{U_T>k\}}] = e^{-\frac{u_0}{2}+\frac{s^2}{8}}\,N\Big(\tfrac{u_0-k-\frac{s^2}{2}}{s}\Big).$$

Multiplying by $P_0 e^{-s^2/8}$ and using $P_0 e^{u_0/2} = S_{1,0}$, $P_0 e^{-u_0/2} = S_{2,0}$ yields the result, with $d_1 = (u_0 - k + \frac{1}{2}s^2)/s$ as stated. □

Corollary 5 (Cash rebate at the ratio barrier)**.** *Let $r > 0$ and let τ be as in (21). In every BD specification, one unit of cash paid at τ if the ratio exits through the upper barrier is worth*

$$\mathbb{E}[e^{-r\tau}\mathbf{1}_{\{U_\tau=u_+\}}] = \frac{\sinh\left(\gamma(u_0-u_-)\right)}{\sinh(\gamma\ell)}, \qquad \gamma = \frac{\sqrt{2r}}{\sigma_U},$$

and the lower-barrier rebate is obtained by replacing $u_0 - u_-$ with $u_+ - u_0$. Likewise, a cash double no-touch paying one unit at T on $\{\tau > T\}$ is worth $e^{-rT}\mathbb{Q}(\tau > T)$, with $\mathbb{Q}(\tau > T)$ given in Remark 12.

Proof. By Corollary 1, it suffices to compute with $U = u_0 + \sigma_U\beta^U$. The rebate formula is the joint Laplace transform of the exit time and exit side of Brownian motion from an interval [30]; the no-touch value is immediate. □

Corollary 4 is best regarded as a benchmark rather than as a stand-alone novelty. Change-of-numeraire arguments already show that exchange-option pricing is naturally governed by relative-price dynamics [4], and, by Remark 8, Margrabe's formula with volatility σ_U holds throughout the larger class of models with constant relative variance. What is useful here is that the BD construction enforces a constant relative quadratic-variation rate σ_U^2 even though the marginal volatility and the instantaneous correlation may vary with the state, so σ_U can be read directly from the exchange-option implied volatility. Corollary 5 is the counterpart that is specific to the BD family: outside it, the $\mathbb{Q}$-drift of U is in general state dependent and cash rebates at ratio barriers are no longer determined by σ_U alone.

4.2. A Relative-Performance Lookback

The invariance is not restricted to barrier survival. Let

$$M_T := \max_{0\le t\le T} U_t, \qquad \mathcal{M}_T := \max_{0\le t\le T} R_t = e^{M_T}.$$

The following standard Brownian maximum transform follows from the reflection principle; see, e.g., Borodin and Salminen [30] and the option-pricing applications in Conze and Viswanathan [31].

Lemma 3 (Terminal value and running maximum)**.** *For $U_t = u_0 + \sigma_U\beta_t$, $m \ge u_0 \vee y$, and $s = \sigma_U\sqrt{T}$,*

$$\mathbb{Q}(U_T \le y, M_T \le m) = N\Big(\frac{y-u_0}{s}\Big) - N\Big(\frac{y-2m+u_0}{s}\Big).$$

Moreover, for $\theta \in \mathbb{R}$, $x = m - u_0 \ge 0$ and $K_\theta = e^{\theta u_0 + \theta^2 s^2/2}$,

$$H_\theta(m) := \mathbb{E}[e^{\theta U_T}\mathbf{1}_{\{M_T\le m\}}] = K_\theta\left[N\Big(\frac{x-\theta s^2}{s}\Big) - e^{2\theta x}N\Big(\frac{-x-\theta s^2}{s}\Big)\right], \tag{28}$$

and

$$\partial_m H_\theta(m) = K_\theta\left[\frac{2}{s}\mathfrak{n}\left(\frac{x-\theta s^2}{s}\right) - 2\theta e^{2\theta x}N\left(\frac{-x-\theta s^2}{s}\right)\right]. \tag{29}$$

Proof. The first formula is the reflection principle. Differentiating it with respect to the terminal state gives the killed density $\varphi_T(y-u_0) - \varphi_T(y-(2m-u_0))$ on $y \le m$. Integrating $e^{\theta y}$ against this density and applying Lemma 2 yields (28). Differentiation then gives (29), using

$$e^{2\theta x}\mathfrak{n}\left(\frac{-x-\theta s^2}{s}\right) = \mathfrak{n}\left(\frac{x-\theta s^2}{s}\right).$$

□

Theorem 4 (Floating-strike relative-performance lookback)**.** *Consider the claim that pays*

$$S_{2,T}\max_{0\le t\le T}\frac{S_{1,t}}{S_{2,t}} - S_{1,T} = S_{2,T}(\mathcal{M}_T - R_T) \ge 0.$$

*Its time-*0 *value is*

$$\Pi_0 = S_{1,0}\left[\left(2+\frac{\sigma_U^2 T}{2}\right)N\left(\frac{\sigma_U\sqrt{T}}{2}\right) + \sigma_U\sqrt{T}\,\mathfrak{n}\left(\frac{\sigma_U\sqrt{T}}{2}\right) - 1\right]. \tag{30}$$

It does not depend on r or on the local-correlation specification.

Proof. Since $S_{2,T} = P_T e^{-U_T/2}$,

$$\Pi_T = P_T\left(e^{M_T - U_T/2} - e^{U_T/2}\right).$$

Theorem 2 thus gives

$$\Pi_0 = P_0 e^{-s^2/8}\left(\mathbb{E}[e^{M_T-U_T/2}] - \mathbb{E}[e^{U_T/2}]\right). \tag{31}$$

The second expectation equals $e^{u_0/2+s^2/8}$. For the first, integrate e^m against $\partial_m H_{-1/2}(m)$ from u_0 to infinity. Setting $x = m - u_0$ in (29) gives

$$\mathbb{E}[e^{M_T-U_T/2}] = e^{u_0/2+s^2/8}(I_1 + I_2),$$

where

$$I_1 = 2\int_0^\infty e^x\frac{1}{s}\mathfrak{n}\left(\frac{x+s^2/2}{s}\right)\mathrm{d}x = 2N\left(\frac{s}{2}\right),$$

and

$$I_2 = \int_0^\infty N\left(\frac{-x+s^2/2}{s}\right)\mathrm{d}x = s\left[\frac{s}{2}N\left(\frac{s}{2}\right) + \mathfrak{n}\left(\frac{s}{2}\right)\right].$$

Substituting into (31), together with $P_0 e^{u_0/2} = S_{1,0}$, yields (30). □

Theorem 4 is useful because it tests more than a terminal relative-price distribution: the running maximum enters explicitly. As $T \downarrow 0$, the bracket in (30) tends to zero, and the price increases with the dispersion scale σ_U. More elaborate seasoned or fixed-strike lookbacks can be derived from Lemma 3, but we do not pursue them here.

4.3. Ratio Barriers Versus Genuinely Two-Dimensional First Passage

The ratio formulation matters for positioning the result correctly. In the constant-parameter Black–Scholes–Margrabe model, $\log(S_1/S_2)$ is already an arithmetic Brownian motion for every constant correlation. As a result, a barrier written only on S_1/S_2 is a one-dimensional problem; this underlies the knock-in/out Margrabe formulas of Haug and Haug [8]. Special barrier exchange options can be even more robust; for instance, Lindset

and Persson [9] give a model-independent formula when the option knocks out as the two asset values meet. The contribution of the BD construction is therefore not the discovery of ratio-barrier reduction but its *preservation*: an exactly Brownian relative coordinate is kept for a state-dependent local-volatility/local-correlation family and exploited for a wider path-dependent class. For one-homogeneous claims, this preservation extends beyond BD (Remark 8); the gains specific to the BD design are the cash-settled ratio book of Corollaries 1 and 5 and the independence structure of Section 5. In foreign exchange, designs of this kind are usually aimed at the one-time law of the cross rate [19,20]; here, the target is its path.

This should be distinguished from the classical two-dimensional passage problems of Iyengar [1], Metzler [3] and from the double-extrema problem of He et al. [2]. Separate barriers on the two asset coordinates, or joint maxima and minima of the two individual assets, do not collapse to the ratio. After decorrelation, they lead to a wedge or quadrant problem whose angular structure depends on the correlation. The BD model does not claim to solve these individual-stock problems. Instead, Section 5 identifies a different two-dimensional region in the original price plane whose two rotated coordinates can be made independent so that its joint exit law separates.

5. Independent Coordinates and Joint Ratio–Index First Passage

We now specialise to the index-autonomous case, in which the index volatility, and equivalently the local correlation, depends on the state only through P.

Theorem 5 (Exact independence)**.** *Suppose $\nu_t = \nu(t, P_t)$ with ν measurable, bounded above and below away from zero, and such that $\mathrm{d}P = P(r - \alpha^*)\mathrm{d}t + P\nu(t,P)\mathrm{d}\beta^V$ has a unique strong solution. Equivalently, in the original parameterisation, suppose that the local correlation depends on the state only through the index,*

$$\rho_t = \rho\big(t, \sqrt{S_{1,t}S_{2,t}}\big) = \frac{4\nu^2(t,P_t) - \sigma_U^2}{4\nu^2(t,P_t) + \sigma_U^2}.$$

Then, U and P are independent*, U is a driftless Brownian motion with variance rate σ_U^2, and P is an autonomous one-dimensional diffusion.*

Proof. By Theorem 1, U solves $\mathrm{d}U = \sigma_U \mathrm{d}\beta^U$ and P solves $\mathrm{d}P = P(r - \alpha^*)\mathrm{d}t + P\nu(t,P)\mathrm{d}\beta^V$, with $\beta^U \perp \beta^V$ independent Brownian motions. Neither equation involves the other state variable, so U is a measurable functional of β^U alone and P (by strong uniqueness) is a measurable functional of β^V alone. The independence of β^U and β^V therefore transfers to U and P. □

Why orthogonality is insufficient, and how far autonomy can be weakened.

Independent Brownian drivers do not make the state processes independent when ν depends on U: observing the index quadratic variation can then convey information about the ratio path. Within the regular positive Markov class $\nu = \nu(t, U, P)$, process independence forces $\nu^2(t, u, p)$ to be independent of u on the reachable state support, up to null sets. Indeed, independence of the entire P and U paths requires the infinitesimal conditional variance of $\log P$, given the joint past, to be unchanged by the U past; that variance is $\nu^2(t, U_t, P_t)\mathrm{d}t$. With continuous coefficients and full state support, this is exactly index autonomy. Positivity of ν removes a possible sign ambiguity. Thus, there is no non-trivial weaker condition on a positive state-only $\rho(t, S_1, S_2)$ that retains process independence under these regularity conditions. One can widen the model by letting ν depend on an auxiliary factor driven independently of β^U; then P may not be autonomous as a one-dimensional process, while the enlarged index system is still independent of U. This is an extension beyond the stated two-price Markov class.

The tractability below uses both process independence and the choice of boundaries. In $(U, \log P)$, the region is a rectangle; its killed generator is a sum of two one-dimensional generators and its semigroup is their tensor product. Separate stock-price barriers do not have this product geometry after the transformation, even when ν is index-autonomous. Now, define the two-dimensional region

$$\begin{aligned} \mathcal{D} := \Big\{(s_1, s_2) \in (0,\infty)^2 : \; L < \frac{s_1}{s_2} < H, \quad A < \sqrt{s_1 s_2} < B\Big\}, \\ 0 < L < R_0 < H, \qquad 0 < A < P_0 < B. \end{aligned} \tag{32}$$

a curvilinear quadrilateral bounded by the two rays $s_1 = Ls_2$, $s_1 = Hs_2$ and the two hyperbolas $s_1 s_2 = A^2$, $s_1 s_2 = B^2$. Set

$$\tau_U := \inf\{t : U_t \notin (u_-, u_+)\}, \qquad \tau_P := \inf\{t : P_t \notin (A, B)\}, \qquad \tau_{\mathcal{D}} := \tau_U \wedge \tau_P.$$

Corollary 6 (Factorisation of the joint exit law)**.** *Under the assumptions of Theorem 5, τ_U and τ_P are independent. Consequently, for all $t_1, t_2 \geq 0$,*

$$\mathbb{Q}\big(\tau_U > t_1,\ \tau_P > t_2\big) = \mathbb{Q}\big(\tau_U > t_1\big)\,\mathbb{Q}\big(\tau_P > t_2\big), \qquad \mathbb{Q}\big(\tau_{\mathcal{D}} > t\big) = \mathbb{Q}(\tau_U > t)\,\mathbb{Q}(\tau_P > t),$$

and the identity of the first boundary piece reached factorises as well: for any of the four pieces,

$$\mathbb{Q}\big(\tau_U < \tau_P,\ U_{\tau_U} = u_+,\ \tau_U \leq t\big) = \int_0^t \mathbb{Q}\big(\tau_P > v\big)\,\mathbb{Q}\big(\tau_U \in \mathrm{d}v,\ U_{\tau_U} = u_+\big),$$

with analogous expressions for the other three. The same factorisation holds for the joint law of the running extrema $\big(\max_{t\leq T} R_t,\ \max_{t\leq T} P_t\big)$ and, more generally, for any pair of functionals, one of the ratio path and one of the index path.

Proof. τ_U is a measurable functional of the path of U, and τ_P of the path of P; the two paths are independent by Theorem 5. The displayed identities are then standard consequences of independence, the last one obtained by conditioning on the exit time of U and using that $\{\tau_P > v\}$ is independent of $\mathcal{F}^U$. □

Corollary 6 provides a tractable two-dimensional exit problem in the original price plane: the survival law of a region with two ray boundaries and two hyperbolic boundaries reduces to two one-dimensional survival laws. The result holds for a broad class of index-dependent local correlations. It is a ratio–index first-passage result and differs from the joint passage times of the two individual stock coordinates.

Remark 12 (The U-leg)**.** *The U-leg is fully explicit. In particular*

$$\mathbb{Q}(\tau_U > t) = \int_{u_-}^{u_+} p_t^{\mathrm{DB}}(u_0, u)\mathrm{d}u = \sum_{n\in\mathbb{Z}} \Big[\mathcal{E}_0\big(u_-, u_+; u_0 + 2n\ell\big) - \mathcal{E}_0\big(u_-, u_+; 2u_- - u_0 + 2n\ell\big)\Big],$$

with $\mathcal{E}_0$ as in Lemma 2 evaluated at horizon t.

Remark 13 (The P-leg)**.** *The P-leg is an ordinary one-dimensional first-passage problem. With constant ν, it reduces to the classical drifted-Brownian double-barrier problem; with state-dependent $\nu(P)$, it can be handled with one-dimensional spectral methods such as those of Linetsky [32]. The contribution lies in the factorisation of the original two-dimensional region, not in a new solution theory for the remaining one-dimensional index problem.*

The factorisation also delivers prices. Write Π_0^{DKO} for the ratio double-knock-out value (22) of a continuous one-homogeneous payoff π.

Corollary 7 (Knock-out exchange option on $\mathcal{D}$)**.** *Under the assumptions of Theorem 5,*

$$\mathbb{E}[e^{-rT}\pi(S_{1,T},S_{2,T})\mathbf{1}_{\{\tau_{\mathcal{D}}>T\}}]=\Pi_0^{\mathrm{DKO}}\,\widehat{\mathbb{Q}}(\tau_P>T),\qquad \frac{\mathrm{d}\widehat{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}\Big|_{\mathcal{F}_T}=\frac{e^{-(r-\alpha^*)T}P_T}{P_0},$$

where, under $\widehat{\mathbb{Q}}$*, the index solves*

$$\mathrm{d}P_t=P_t\big(r-\alpha^*+\nu^2(t,P_t)\big)\mathrm{d}t+P_t\nu(t,P_t)\,\mathrm{d}\hat{\beta}_t^V$$

for a $\widehat{\mathbb{Q}}$*-Brownian motion* $\hat{\beta}^V$*, and the law of* U *is unchanged. For constant* ν*, set* $y_0=\ln P_0$*,* $a_P=\ln A$*,* $b_P=\ln B$*,* $\ell_P=b_P-a_P$*,* $s_P=\nu\sqrt{T}$*,* $\hat{\theta}=(r-\alpha^*+\nu^2/2)/\nu^2$*,* $m_n=y_0+2n\ell_P$ *and* $\tilde{m}_n=2a_P-y_0+2n\ell_P$*; then*

$$\begin{aligned}\widehat{\mathbb{Q}}(\tau_P>T)=\sum_{n\in\mathbb{Z}}\Big\{&e^{2\hat{\theta}n\ell_P}\Big[N\Big(\tfrac{b_P-m_n-\hat{\theta}s_P^2}{s_P}\Big)-N\Big(\tfrac{a_P-m_n-\hat{\theta}s_P^2}{s_P}\Big)\Big]\\&-e^{2\hat{\theta}(a_P-y_0+n\ell_P)}\Big[N\Big(\tfrac{b_P-\tilde{m}_n-\hat{\theta}s_P^2}{s_P}\Big)-N\Big(\tfrac{a_P-\tilde{m}_n-\hat{\theta}s_P^2}{s_P}\Big)\Big]\Big\}.\end{aligned}\tag{33}$$

Proof. By (20), the payoff is $P_T\widehat{\pi}(U_T)\mathbf{1}_{\{\tau_U>T\}}\mathbf{1}_{\{\tau_P>T\}}$. The density $Z_T:=e^{-(r-\alpha^*)T}P_T/P_0=\mathcal{E}(\int\nu\,\mathrm{d}\beta^V)_T$ is a true martingale because ν is bounded, and it is a functional of β^V alone, as is $\mathbf{1}_{\{\tau_P>T\}}$. Independence of β^U and β^V gives

$$\mathbb{E}[e^{-rT}P_T\widehat{\pi}(U_T)\mathbf{1}_{\{\tau_U>T\}}\mathbf{1}_{\{\tau_P>T\}}]=P_0e^{-\alpha^*T}\,\mathbb{E}[Z_T\mathbf{1}_{\{\tau_P>T\}}]\,\mathbb{E}[\widehat{\pi}(U_T)\mathbf{1}_{\{\tau_U>T\}}],$$

and $P_0e^{-\alpha^*T}\mathbb{E}[\widehat{\pi}(U_T)\mathbf{1}_{\{\tau_U>T\}}]=\Pi_0^{\mathrm{DKO}}$ by the proof of Theorem 3. Girsanov's theorem gives the $\widehat{\mathbb{Q}}$-dynamics of P and leaves β^U a Brownian motion. For constant ν, $\ln P$ is a $\widehat{\mathbb{Q}}$-Brownian motion with drift $\hat{\theta}\nu^2$ and volatility ν; weighting the killed density on (a_P,b_P) by $\exp\{\hat{\theta}(y-y_0)-\frac{1}{2}\hat{\theta}^2s_P^2\}$ and applying Lemma 2 with (σ_U,s) replaced by (ν,s_P) gives the series. □

The denominators of all four normal CDF arguments in (33) are $s_P=\nu\sqrt{T}$; b_P denotes the upper log barrier and cannot serve as a volatility scale. For $\nu=0.22$ and the parameters of Section 6, with $A=85$, $B=115$ and $\kappa=1.05$, the series gives $\widehat{\mathbb{Q}}(\tau_P>1)=0.09163191131$ and a knock-out value 0.09568161982. Direct numerical integration of the constant-drift Black–Scholes killed log-index density gives the same survival probability to 2×10^{-16}, independently checking the CDF scaling and exponential weights. For constant ν, the correlation is constant and Corollary 7 reduces to a Black–Scholes benchmark. Its content lies in state-dependent $\nu(t,P)$, for which $\widehat{\mathbb{Q}}(\tau_P>T)$ is a one-dimensional first-passage probability that can be computed by the methods mentioned in the preceding remark.

6. Numerical Verification

All formulas above were derived through the canonical representation of Theorem 1. To confirm both the representation and the prices, we simulate the original two-dimensional system (1)–(9) directly in (S_1,S_2), never using the reduced form, and compare. The revised verification uses $N=480{,}000$ antithetic paths (240,000 independent pairs) for Specs A and B, with coupled grids $\Delta=1/252$, $1/504$ and $1/1008$. Monte Carlo standard errors are reported throughout. Coarse increments are sums of the fine increments; this allows standard errors for paired grid differences. Barrier probabilities use the exact two-sided Brownian-bridge series to numerical precision.

Appendix A describes the simulation scheme.

We take

$$r=0.03,\quad \sigma_U=0.25,\quad S_{1,0}=100,\quad S_{2,0}=95,\quad L=0.80,\quad H=1.30,\quad T=1,$$

so that $u_0 = 0.05129$, $u_- = -0.22314$, $u_+ = 0.26236$, $\ell = 0.48551$ and $P_0 = 97.4679$. Two structurally different correlation specifications are used, both of the form $\rho = \rho(u)$:

$$\text{Spec A (U-shaped):} \quad \rho_A(u) = 0.60 - 0.40\ \mathrm{sech}^2(4u), \quad \rho_A(0) = 0.20,\ \rho_A(\pm\infty) = 0.60,$$
$$\text{Spec B (hump-shaped):} \quad \rho_B(u) = 0.10 + 0.70\, e^{-2u^2}, \quad \rho_B(0) = 0.80,\ \rho_B(\pm\infty) = 0.10.$$

The induced common volatilities $\sigma = \sigma_U/\sqrt{2(1-\rho)}$ differ substantially: as $|u|$ varies, Spec A ranges over $[0.1976, 0.2795]$ and Spec B over $[0.1863, 0.3953]$, and the two are monotone in opposite directions. Any dependence of prices on $\rho(\cdot)$ would be clearly visible.

Table 2 confirms that the log-ratio is a driftless Brownian motion with variance rate σ_U^2 and that both discounted assets are martingales, as asserted by Theorem 1. Monte Carlo standard errors are given in parentheses.

Table 2. Structural checks from direct simulation of (1)–(9); $N = 480{,}000$ and $\Delta = 1/252$. Errors use independent antithetic pair averages. The ratio paths are shared across specifications; antithetic pairing makes their sample mean exactly u_0, up to roundoff.

Quantity	Spec A	Spec B	Theoretical
$\mathrm{Var}(U_T - u_0)/T$	0.06252 (0.00018)	0.06252 (0.00018)	0.06250
$\mathbb{E}[U_T] - u_0$	0.00000 (0.00000)	0.00000 (0.00000)	0
$e^{-rT}\mathbb{E}[S_{1,T}]$	99.999 (0.008)	99.991 (0.019)	100
$e^{-rT}\mathbb{E}[S_{2,T}]$	95.003 (0.006)	95.000 (0.019)	95

Table 3 reports two fixed-maturity contracts that provide a genuine test of Theorem 2.

Table 3. Invariant prices from the revised direct simulation, $N = 480{,}000$ antithetic paths, $\Delta = 1/252$, with exact two-sided bridge survival and exact Brownian-bridge maximum sampling. Standard errors appear in parentheses; the two specifications share the ratio innovations.

Contract	MC, Spec A	MC, Spec B	Closed Form
Double knock-out exchange, Corollary 3			
$\kappa = 0.90$	4.0710 (0.0124)	4.0724 (0.0128)	4.0692
$\kappa = 1.05$	1.0443 (0.0043)	1.0446 (0.0046)	1.0442
$\kappa = 1.20$	0.0636 (0.0007)	0.0636 (0.0007)	0.0640
Relative-performance lookback, Theorem 4			
$S_{2,T}\mathcal{M}_T - S_{1,T}$	21.5696 (0.0146)	21.5693 (0.0173)	21.5615

To show that the result is not a reparametrised Black–Scholes model, Table 4 prices a fixed-strike basket call, which lies outside the invariant class because of its currency strike.

The same direct simulation and paths are used; sharing ratio and index innovations across specifications sharpens the comparison.

Table 4. Revised discounted values of $(S_{1,T} + S_{2,T} - K)^+$, which lies outside the invariant class; $N = 480{,}000$, $\Delta = 1/252$. Standard errors appear in parentheses.

Basket Strike K	MC, Spec A	MC, Spec B
180	25.506 (0.023)	36.754 (0.059)
195	16.481 (0.026)	29.291 (0.059)
210	9.972 (0.024)	23.114 (0.057)

The comparison isolates the content of the theory. Specs A and B agree on both the DKO and the lookback within simulation error, as the invariance theorem predicts, yet they produce very different basket prices. Fixing σ_U therefore does not erase the local-correlation surface; it renders that surface invisible only to the precisely identified invariant class.

Table 5 reports the three grids for a representative DKO, the lookback and a basket, together with paired fine-minus-daily differences. All strikes in Tables 3 and 4 are included in the replication results. The ratio increments already have their exact Brownian law on every grid; refinement tests the index-volatility freezing error and its interaction with the payoff, not a convergence approximation to the ratio variance. The fine-grid basket changes are small relative to their prices and much smaller than the differences between Specs A and B. Paired errors can detect small grid effects that are hidden by the standard errors of separate price estimates, so stability is not asserted to mean zero discretisation error.

Table 5. Coupled-grid convergence, $N = 480{,}000$. Parentheses give Monte Carlo standard errors; in the last row for each specification, they are standard errors of the paired pathwise grid differences.

Spec	Δ	DKO, $\kappa = 1.05$	Lookback	Basket, $K = 195$
A	1/252	1.0443 (0.0043)	21.570 (0.015)	16.481 (0.026)
A	1/504	1.0448 (0.0044)	21.570 (0.015)	16.483 (0.026)
A	1/1008	1.0447 (0.0044)	21.570 (0.015)	16.484 (0.026)
Fine minus daily, A		0.00046 (0.00076)	−0.00005 (0.00006)	0.00340 (0.00039)
B	1/252	1.0446 (0.0046)	21.569 (0.017)	29.291 (0.059)
B	1/504	1.0451 (0.0047)	21.569 (0.017)	29.287 (0.059)
B	1/1008	1.0451 (0.0047)	21.569 (0.017)	29.286 (0.059)
Fine minus daily, B		0.00049 (0.00080)	0.00013 (0.00010)	−0.00497 (0.00067)

The revised runs replace the submitted $\max\{1 - p_\uparrow - p_\downarrow, 0\}$ correction by the exact two-sided Brownian-bridge series. The paired corrections coincide at double precision in these daily-grid DKO runs. This is explained by a bound rather than by rounding: for a driftless Brownian step with variance $v = \sigma_U^2 \Delta$, the probability of touching both barriers is at most $4N(-\ell/\sqrt{v})$. Inclusion–exclusion shows that this is the only event omitted by the two one-sided probabilities. Using a union bound over T/Δ steps and the bounded exchange payoff on the corridor gives, at the largest payoff ($\kappa = 0.90$),

$$0 \le \Pi_{\text{exact bridge}} - \Pi_{\text{one-sided sum}} \le P_0 e^{-\alpha^* T} M \frac{T}{\Delta} 4N\left(-\frac{\ell}{\sigma_U \sqrt{\Delta}}\right) < 1.85 \times 10^{-204}. \tag{34}$$

The normal-tail bound follows by conditioning on the first boundary hit: the two successive Brownian hitting distances add to at least ℓ, and their first-passage-time convolution is the hitting law for the sum of the distances. The price bound uses the conditional mean-one index exponential; it holds for the ratio-only Euler schemes here, as well as for the reduced ratio claim. It is far below the reported Monte Carlo errors. Structural checks and unbarriered baskets do not use a barrier correction, while the lookback maximum was already sampled exactly; hence, this particular approximation introduces no error into those rows. It can matter for a much narrower corridor or a coarser grid, for which the exact series is preferable. To test Corollary 7, we retain $\nu(p) = \min\{\max\{0.22(P_0/p)^{1.5}, 0.05\}, 0.8\}$, $A = 85$, $B = 115$ and $\kappa = 1.05$. Correlation ranges over approximately $[0.31, 0.65]$ inside the index corridor. The original two-dimensional system is simulated with $N = 400{,}000$ on the same three coupled grids. The U bridge is exact; the $\log P$ bridge is exact for each frozen constant-coefficient Euler step and remains a discretisation of the state-dependent index diffusion.

As an independent check, we solve the one-dimensional backward survival equation under $\widehat{\mathbb{Q}}$ in $y = \log p$:

$$\partial_t h = \tfrac{1}{2}\nu^2(e^y)\partial_{yy} h + \left(r - \alpha^* + \tfrac{1}{2}\nu^2(e^y)\right)\partial_y h, \quad h(0, y) = 1, \quad h(t, \log A) = h(t, \log B) = 0.$$

Central spatial differences and an implicit-midpoint scheme with two initial half-step backward-Euler steps give survival probabilities 0.10454995, 0.10454950 and 0.10454944

with, respectively, 400, 800 and 1600 spatial intervals and 4 times as many time steps. The last spatial refinement changes the product price by less than 7×10^{-8}, well below the Monte Carlo errors. Table 6 compares that benchmark with direct pricing; the three direct estimates agree with the product within their reported simulation uncertainty.

Table 6. Knock-out exchange option on $\mathcal{D}$ under the index-dependent correlation, with the original contract parameters and $N = 400{,}000$. Parentheses are Monte Carlo standard errors; the one-dimensional PDE values are deterministic benchmarks.

Quantity	Value
Direct two-dimensional MC, $\Delta = 1/252$	0.11075 (0.00164)
Direct two-dimensional MC, $\Delta = 1/504$	0.11066 (0.00167)
Direct two-dimensional MC, $\Delta = 1/1008$	0.11094 (0.00170)
Paired 1/1008 minus 1/252	0.00018 (0.00045)
Π_0^{DKO}, closed form	1.04419539
$\widehat{\mathbb{Q}}(\tau_P > T)$, one-dimensional PDE	0.10454944
Product, Corollary 7	0.10917004

For this frozen-step index simulation, the omitted two-barrier intersection in the old one-sided correction also has a conservative price bound below 8.3×10^{-15} at the daily grid. One may use the Brownian range bound $4N\{-(\ell_P - d\Delta)/(2\bar{\nu}\sqrt{\Delta})\}$ per step, where $\bar{\nu} = 0.27014$ and $d = \sup_{p\in[A,B]} |r - \alpha^* - \nu^2(p)/2| = 0.01430$; on survival, the discounted payoff is bounded by $e^{-rT} B \sup \widehat{\pi}$. This is a bound for the bridge-intersection approximation, not for coefficient freezing. The latter is assessed by the coupled grids and the PDE comparison above.

7. Discussion and Conclusions

The BD construction makes the correlation constraint part of the parameterisation. Equal instantaneous volatilities render the log-ratio and log-product orthogonal; fixing the log-ratio variance then leaves the orthogonal index variance free. The canonical representation maps this remaining variance into an automatically admissible local correlation.

The pricing consequence is intentionally narrow. Scale invariance already explains why relative prices are natural for homogeneous claims, and for one-homogeneous claims the invariance holds whenever the relative variance is constant. The BD design adds invariance of the cash-settled ratio book, because the ratio is a Brownian motion under the pricing measure itself, and the construction extends to a state-dependent dispersion volatility (Remark 9). Once σ_U is fixed, the invariant classes are observationally equivalent to a constant-parameter Black–Scholes–Margrabe model with equal volatilities and the same relative variance, whereas fixed-strike basket claims can remain highly sensitive to the local-correlation specification.

The common instantaneous volatility is also a genuine restriction: two arbitrary single-name local-volatility surfaces cannot be prescribed independently. Because the common volatility is cross-dependent, the resulting marginal smiles need not coincide, but they are jointly constrained. The model is therefore most natural for pairs for which a common instantaneous variance is an acceptable structural approximation.

The second main result is the ratio–index decomposition of a non-rectangular two-dimensional exit region. When the index is autonomous, the rotated coordinates are independent and the joint exit law factorises; the price of an exchange option knocked out on that region is then a ratio double-knock-out value times a one-dimensional index survival probability under an index-tilted measure. This does not solve the classical problem of two fixed individual-stock barriers; rather, it shows how dependence design can create

non-trivial regions for which multi-dimensional first passage is analytically separable. Extensions are left for separate work.

Calibration, Pricing and Correlation-Risk Management

The numerical example is a model-identification experiment, not a fit to market data. An exchange-option or relative-performance book constrains σ_U, but its common price under Specs A and B cannot select between their dependence surfaces. At fixed σ_U, index options, fixed-currency-strike baskets and single-name options provide information about ν; the correlation follows from (12). A practical calibration would first estimate σ_U from available exchange or other relative-performance quotes then fit a bounded finite-dimensional $\nu(t, u, p; \theta)$ jointly to basket/index and single-name prices, checking residuals and the equal-variance restriction. Sparse quotes do not uniquely identify an unrestricted local-correlation surface. Historical covariance can help choose or regularise a specification but is not itself a risk-neutral option calibration. In the index-autonomous subclass, both calibration and index-barrier pricing use one-dimensional index solvers. Outside that subclass, the non-invariant calibration generally remains two-dimensional or simulation based.

Single-name quotes are a compatibility check, not two freely prescribable inputs. Where Markovian projection applies, their projected local variances satisfy

$$\sigma^2_{i,\mathrm{loc}}(t,s) = \mathbb{E}_{\mathbb{Q}}\left[\, \nu_t^2 + \frac{\sigma_U^2}{4} \middle| S_{i,t} = s \right], \qquad i = 1, 2.$$

Different conditioning can produce different marginal smiles, despite the common pathwise variance; nevertheless, those smiles must be jointly compatible with the BD dynamics. Persistent misfit is evidence to relax the model, not to force calibration.

For investors, a ratio-barrier exchange option provides contingent participation in one asset's outperformance of the other while the relative price stays in a prescribed range. The lookback protects the gap between the best observed relative performance and its terminal level. These are relative-performance exposures; fixed strike baskets additionally expose the investor to aggregate price and index-volatility risk. Their very different prices under Specs A and B show why a relative-performance calibration alone is insufficient for a basket book.

The separation also yields useful hedge formulas. Between updates of any additional ratio-path state, let a smooth degree-one claim have value $C = Pf(t, u)$, with that state suppressed. Its stock deltas are

$$\Delta_1 = \frac{P}{S_1}\left(\frac{f}{2} + f_u\right), \qquad \Delta_2 = \frac{P}{S_2}\left(\frac{f}{2} - f_u\right). \tag{35}$$

For a cash-settled ratio value $g(t, u)$, they are $\Delta_1 = g_u/S_1$ and $\Delta_2 = -g_u/S_2$. In the two-Brownian market with $\nu > 0$, the stock diffusion matrix is invertible, so these deltas match both diffusion exposures; the residual value is financed in the money-market account under the usual smoothness and self-financing conditions. Ratio barriers are hedged up to knock-out, with their prescribed boundary settlement. The reduced equations are $f_t + \sigma_U^2 f_{uu}/2 - \alpha^* f = 0$ and $g_t + \sigma_U^2 g_{uu}/2 - rg = 0$, with the contract's boundary and terminal conditions, and neither contains ν.

This does not mean that these contracts have zero conventional correlation sensitivity. At fixed marginal volatilities, $\partial\sigma_U^2/\partial\rho = -2\sigma_1\sigma_2$, so changing correlation changes the dispersion parameter. The invariance is along BD perturbations that hold σ_U fixed and adjust the common marginal variance together with correlation. A risk manager can therefore stress ν at fixed σ_U to isolate index/dependence model risk in baskets and index-barrier claims then stress σ_U separately for relative-performance risk. Common ratio prices are useful pricing and hedge benchmarks, but they do not establish that the underlying models have the same basket risk or market calibration.

Funding: This research received no external funding.

Data Availability Statement: No new data were created or analysed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Simulation Scheme

The revised results are obtained from the original log-asset Euler equations. At each step, evaluate ρ and the common $\sigma = \sigma_U/\sqrt{2(1-\rho)}$ at the current original asset state and set

$$\Delta W_1 = \sqrt{\frac{1+\rho}{2}}\Delta B_c + \sqrt{\frac{1-\rho}{2}}\Delta B_d, \qquad \Delta W_2 = \sqrt{\frac{1+\rho}{2}}\Delta B_c - \sqrt{\frac{1-\rho}{2}}\Delta B_d,$$

where $\Delta B_c, \Delta B_d$ are independent $N(0,\Delta)$ variables. Then

$$X_{i,t_{k+1}} = X_{i,t_k} + \left(r - \tfrac{1}{2}\sigma^2\right)\Delta + \sigma\Delta W_i.$$

This is a symmetric square root of the original correlation matrix. All pricing payoffs are calculated from the resulting $S_i = e^{X_i}$. The ratio and index are recorded as diagnostics and for monitoring. Fine increments on 1/1008 are summed in groups of two and four for the 1/504 and 1/252 grids. The same independent innovation pairs drive Specs A and B. Seeds are 4,626,192 for the ratio-dependent simulations and 4,626,193 for the index-dependent simulation; batch sizes are 10,000. The ratio experiments pair each Brownian innovation path with its negative and each bridge uniform with its complement. Their 480,000 paths form 240,000 independent antithetic pairs, and errors use the standard deviation of the pair-averaged payoffs divided by $\sqrt{240{,}000}$. The index experiment uses 400,000 independent paths and the usual sample standard deviation divided by $\sqrt{N}$. No price control variates are used. Paired grid errors apply the appropriate formula to $H_j^{\text{fine}} - H_j^{\text{daily}}$.

For endpoints $x, y \in (a,b)$ and a constant-variance Brownian step, let $v = \eta^2\Delta$, $\ell = b-a$ and $\phi_v(z) = e^{-z^2/(2v)}/\sqrt{2\pi v}$. The conditional survival probability is the killed/free transition-density ratio

$$q_\Delta(x,y) = \frac{\sum_{n\in\mathbb{Z}}\left[\phi_v(y-x-2n\ell) - \phi_v(y+x-2a-2n\ell)\right]}{\phi_v(y-x)}. \tag{A1}$$

It is zero if either endpoint is outside the interval. A constant drift cancels in this ratio. The step probabilities are multiplied as conditional survival weights, rather than drawing extra Bernoulli variables. The image series, or the eigenfunction kernel for a narrow interval, is truncated below numerical precision; only roundoff-sized excursions outside $[0,1]$ are clipped. At the grids and corridor widths used here, the implementation retains the zero and first positive images and the four closest negative images; the remaining terms underflow in double precision. This evaluates the exact series to floating-point precision, rather than retaining only the two one-sided crossing probabilities. The analytic distinction is quantified in (34).

Use $\eta = \sigma_U$ for U, which has its exact Brownian path law. For $\log P$ in the index-autonomous experiment, use the frozen step value $\eta = \nu(P_{t_k})$. The latter is exact bridge monitoring of the Euler step, not exact simulation of the continuous local-volatility index diffusion. Time-step refinement and the independent PDE benchmark test that remaining approximation. Ignoring bridge monitoring entirely instead leaves the usual barrier discretisation error [33].

Running maximum.

For a Brownian step of U, the conditional maximum is sampled by

$$M_k = \tfrac{1}{2}\Big(U_{t_k} + U_{t_{k+1}} + \sqrt{(U_{t_{k+1}} - U_{t_k})^2 - 2\sigma_U^2 \Delta \log E_k}\Big), \qquad E_k \sim \mathrm{Unif}(0,1).$$

For the coupled comparison, fine-grid endpoints and bridge maxima define the same continuous ratio maximum for all three grids. Coarser index Euler payoffs are evaluated on that shared ratio path. The maximum sampling is exact for U; the joint asset simulation still has its index-coefficient freezing error.

For applications rather than verification, the canonical representation is preferable: simulate U as a Brownian motion and evolve P under the chosen index specification. The accompanying replication script supplies the simulations, both truncation bounds, the constant-ν integral check and the refined one-dimensional PDE benchmark.

References

1. Iyengar, S. Hitting lines with two-dimensional Brownian motion. *Siam J. Appl. Math.* **1985**, *45*, 983–989. [CrossRef]
2. He, H.; Keirstead, W.P.; Rebholz, J. Double lookbacks. *Math. Finance* **1998**, *8*, 201–228. [CrossRef]
3. Metzler, A. On the first passage problem for correlated Brownian motion. *Stat. Probab. Lett.* **2010**, *80*, 277–284. [CrossRef]
4. Geman, H.; El Karoui, N.; Rochet, J.-C. Changes of numéraire, changes of probability measure and option pricing. *J. Appl. Probab.* **1995**, *32*, 443–458. [CrossRef]
5. Hoogl, J.K.; Neumann, C.D.D. Local scale invariance and contingent claim pricing. *Int. J. Theor. Appl. Finance* **2001**, *4*, 1–21. [CrossRef]
6. Hoogl, J.K.; Neumann, C.D.D. Local scale invariance and contingent claim pricing II: Path-dependent contingent claims. *Int. J. Theor. Appl. Finance* **2001**, 4, 23–43. [CrossRef]
7. Margrabe, W. The value of an option to exchange one asset for another. *J. Finance* **1978**, *33*, 177–186. [CrossRef]
8. Haug, E.G.; Haug, J. The collector: Knock-in/out Margrabe. *Wilmott* **2002**, *2002*, 38–41. [CrossRef]
9. Lindset, S.; Persson, S.-A. A note on a barrier exchange option: The world's simplest option formula? *Finance Res. Lett.* **2006**, *3*, 207–211. [CrossRef]
10. Dupire, B. Pricing with a smile. *Risk* **1994**, *7*, 18–20.
11. Derman, E.; Kani, I. Riding on a smile. *Risk* **1994**, *7*, 32–39.
12. Gyöngy, I. Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probab. Theory Relat. Fields* **1986**, *71*, 501–516. [CrossRef]
13. Piterbarg, V. Markovian projection for volatility calibration. *Risk* **2007**, *20*, 84–89.
14. Cont, R.; Deguest, R. Equity correlations implied by index options: Estimation and model uncertainty analysis. *Math. Finance* **2013**, *23*, 496–530. [CrossRef]
15. Guyon, J. Cross-dependent volatility. *Risk* **2016**, *29*, 61–65. [CrossRef]
16. Guyon, J. Calibration of local correlation models to basket smiles. *J. Comput. Finance* **2017**, *21*, 1–51. [CrossRef]
17. Jourdain, B.; Sbai, M. Coupling index and stocks. *Quant. Finance* **2012**, *12*, 805–818. [CrossRef]
18. Brown, C.A.; Handley, J.C.; Palmer, K. A closer look at barrier exchange options. *J. Futur. Mark.* **2013**, *33*, 29–43. [CrossRef]
19. Austing, P. Repricing the cross smile: An analytic joint density. *Risk* **2011**, 72–75.
20. De Col, A.; Gnoatto, A.; Grasselli, M. Smiles all around: FX joint calibration in a multi-Heston model. *J. Bank. Finance* **2013**, *37*, 3799–3818. [CrossRef]
21. Escobar-Anel, M. Local stochastic correlation models for derivative pricing. *Stats* **2025**, *8*, 65. [CrossRef]
22. Götz, B.; Escobar, M.; Zagst, R. Closed-form pricing of two-asset barrier options with stochastic covariance. *Appl. Math. Finance* **2014**, *21*, 363–397. [CrossRef]
23. Escobar, M.; Ferrando, S.; Wen, X. Barrier options in three dimensions. *Int. J. Finance Mark. Deriv.* **2014**, *3*, 260–292. [CrossRef]
24. Escobar, M.; Panz, S.; Zagst, R. Pricing multiple barrier derivatives under stochastic volatility. *J. Comput. Finance* **2020**, *24*, 77–101. [CrossRef]
25. Kunitomo, N.; Ikeda, M. Pricing options with curved boundaries. *Math. Finance* **1992**, *2*, 275–298. [CrossRef]
26. Goldman, M.B.; Sosin, H.B.; Gatto, M.A. Path dependent options: Buy at the low, sell at the high. *J. Finance* **1979**, *34*, 1111–1127. [CrossRef]
27. Longin, F.; Solnik, B. Is the correlation in international equity returns constant: 1960–1990? *J. Int. Money Finance* **1995**, *14*, 3–26. [CrossRef]

28. Driessen, J.; Maenhout, P.J.; Vilkov, G. The price of correlation risk: Evidence from equity options. *J. Finance* **2009**, *64*, 1377–1406. [CrossRef]
29. Karatzas, I.; Shreve, S.E. *Brownian Motion and Stochastic Calculus*, 2nd ed.; Springer: New York, NY, USA, 1991. [CrossRef]
30. Borodin, A.N.; Salminen, P. *Handbook of Brownian Motion—Facts and Formulae*, 2nd ed.; Birkhäuser: Basel, Switzerland, 2002.
31. Conze, A.; Viswanathan, R. Path dependent options: The case of lookback options. *J. Finance* **1991**, *46*, 1893–1907. [CrossRef]
32. Linetsky, V. Lookback options and diffusion hitting times: A spectral expansion approach. *Finance Stoch.* **2004**, *8*, 373–398. [CrossRef]
33. Broadie, M.; Glasserman, P.; Kou, S. A continuity correction for discrete barrier options. *Math. Finance* **1997**, *7*, 325–349. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.