Market-making in spot precious metals

The primary challenge of market-making in spot precious metals is navigating the liquidity provided by futures contracts. The exchange-for-physical (EFP) spread, which is the price difference between the futures and the spot, plays a pivotal role and exhibits multiple modes of relaxation corresponding to the diverse trading horizons of market participants. Alexander Barzykin, Philippe Bergault and Olivier Guéant model the EFP spread using a nested Ornstein-Uhlenbeck process, in the spirit of the two-factor Hull-White model for interest rates, and provide a computationally efficient stochastic optimal control framework that capitalises on the co-integration properties intrinsic to the precious metals sector

Recent decades have seen a profound transformation in financial markets, driven largely by the widespread adoption of electronification. This digital evolution has been paralleled by the development of new decision-making tools for market-makers and the emergence of systematic market-making for most asset classes. The case of stocks traded through central limit order books has been discussed extensively (see Guilbaud & Pham (2013) for a good example using stochastic optimal control), but over-the-counter (OTC) markets have also been addressed, using the modelling framework introduced by Avellaneda & Stoikov (2008) (see, for example, Barzykin *et al* (2023) for the foreign exchange (FX) market). These market-making models have helped dealers manage inventory risk through strategic hedging and quote skewing. In single-asset models, optimal quotes typically vary monotonically with current inventories, influenced by factors such as risk aversion, clients' pricing sensitivity and the liquidity on external platforms. Dealers aim to maximise client flow internalisation in order to mitigate the cost of external execution and reduce market impact, hedging only when inventories exceed franchise-dependent thresholds (on the internalisation-externalisation dilemma, see, for example, Barzykin *et al* (2022) and Butz & Oomen (2019)).

Multi-asset extensions of market-making models have been built to help dealers manage their risk at the portfolio level (see, for example, Barzykin *et al* 2023; Bergault *et al* 2021). In practice, market-makers often deal with large portfolios of assets with very different levels of liquidity. Illiquid assets may be difficult to internalise and costly to execute in the market. However, the risk associated with illiquid assets may sometimes be partially offset by positions in other instruments that are more liquid, leaving time to unwind the illiquid asset with less cost. In FX markets, illiquid instruments can also be traded via more liquid legs, introducing both complexities and opportunities for the dealer (see Barzykin *et al* 2023; Cartea *et al* 2020). Most existing models consider the case of asset price dynamics driven by correlated Brownian motions.

Diverging from conventional approaches that focus solely on correlated asset price dynamics, our model capitalises on the unique benefits of using co-integrated assets for hedging, regardless of whether the market-maker proposes quotes in those assets. Co-integration offers a special opportunity for market-makers to tap into enhanced liquidity pools and benefit from mean reversion. This is particularly relevant for spot dealers hedging their position through futures, as in the precious metals market. While an interbank spot market does exist for precious metals, precious metals futures are considerably more liquid with tighter spreads and higher volumes. Since the price difference between futures and the spot – the so-called exchange-for-physical (EFP) spread – is primarily driven by swap rates that are relatively stable over the market-maker's intraday risk horizon, an assumption of co-integration between spot and futures prices is very natural. There are other asset classes where the same mathematical ideas could be used (eg, market making in non-deliverable forwards using onshore spot liquidity).

Our modelling framework builds on the classical foundation laid by Avellaneda & Stoikov (2008) and Cartea *et al* (2014) (see Guéant (2016) for a detailed discussion). In this framework, the price of a single underlying instrument (the spot price) is modelled by an arithmetic Brownian motion. The arrival rates and size distribution of client trades are modelled with predetermined intensity kernels. Going beyond the initial academic literature and building on recent advances, our model considers a dealer quoting in the spot who can hedge (with costs) in both the spot and futures. The novelty comes from modelling the EFP spread by a nested Ornstein-Uhlenbeck (OU) process, as inspired by the observation of multiple relaxation times in the market, ranging from hours to days, possibly relating to the different trading horizons of different types of traders. This model is in the spirit of the two-factor Hull-White model for interest rates (Hull & White 1994). However, the observed EFP volatility significantly exceeds the underlying interest rate volatility, so the mechanism is not driven purely by interest rates but is also influenced by the cost of physical delivery, apart from speculation. This article therefore expands the now classical stochastic optimal control framework for market-making by incorporating the existence of co-integrated and liquid assets. Not only does it address a gap in the existing academic literature on market-making models, but it also provides a new class of price dynamics that is compatible with the approximation techniques developed by Bergault *et al* (2021).

In the following sections, we begin by introducing our model alongside the key equations that define the optimal strategies for market-makers in the spot precious metals market. We also discuss the merits and limitations of using nested OU processes. Subsequently, we demonstrate the applicability of the method of Bergault *et al* (2021) to our general framework. We conclude with an in-depth numerical analysis focused on the gold market, illustrating the practical implications of our model.

The model

State variables and the optimal control problem. We consider a spot market-maker in a given precious metal. We denote by $(S_t)_{t\geqslant 0}$ a reference spot price for that metal and assume that the market-maker continuously streams to clients a pricing ladder at the bid, $S^{\mathrm{b}}(t,z) = S_t - \delta^{\mathrm{b}}(t,z)$, and at the ask, $S^{\mathrm{a}}(t,z) = S_t + \delta^{\mathrm{a}}(t,z)$, where z denotes the size in ounces (oz). Transaction likelihood is assumed to depend solely on the distance between the proposed prices and the reference spot price. Subsequently, we introduce two intensity functions for the bid and ask: $(z, \delta) \mapsto \Lambda^{\mathrm{b}}(z, \delta)$ and $(z, \delta) \mapsto \Lambda^{\mathrm{a}}(z, \delta)$. We assume that the functions Λ^{b} and Λ^{a} take the form:

$$\Lambda^{\mathrm{b}}(z, \delta) = \Lambda^{\mathrm{a}}(z, \delta) = \lambda(z) f(\delta)$$

with $f(\delta) = (1 + \mathrm{e}^{\alpha + \beta\delta})^{-1}$, $\beta > 0$ and $\alpha \in \mathbb{R}$.

The market-maker is also capable of hedging their position by trading on various platforms, using either the spot or futures or both. The execution rates associated with the market-maker's activities are modelled by the processes $(v_t^S)_{t \geqslant 0}$ and $(v_t^F)_{t \geqslant 0}$ for the spot and futures, respectively.

We denote by $(F_t)_{t \geqslant 0}$ the price process of the futures contract and by $(E_t)_{t \geqslant 0}$ the EFP process; that is, for all $t \geqslant 0$, $F_t = S_t + E_t$.[1] In what follows, we consider Brownian dynamics for $(S_t)_{t \geqslant 0}$ and nested OU dynamics for $(E_t)_{t \geqslant 0}$:

$$\begin{aligned}
\mathrm{d}S_t &= \sigma_S \,\mathrm{d}W_t^S, \quad \sigma_S > 0 \\
\mathrm{d}E_t &= -k_E(E_t - D_t)\,\mathrm{d}t + \sigma_E \,\mathrm{d}W_t^E, \quad k_E, \sigma_E > 0 \\
\mathrm{d}D_t &= -k_D(D_t - \bar{D})\,\mathrm{d}t + \sigma_D \,\mathrm{d}W_t^D, \quad k_D, \sigma_D \geqslant 0, \ \bar{D} \in \mathbb{R}
\end{aligned}$$

where $(W_t^S, W_t^E, W_t^D)_{t \geqslant 0}$ is a three-dimensional Brownian motion with correlation matrix R. In what follows, we denote by Σ the variance-covariance matrix $\operatorname{diag}(\sigma_S, \sigma_E, \sigma_D) R \operatorname{diag}(\sigma_S, \sigma_E, \sigma_D)$.

We denote the inventory processes of the market-maker by $(q_t^S)_{t \geqslant 0}$ and $(q_t^F)_{t \geqslant 0}$ for spot and futures contracts, respectively. Mathematically, the dynamics of $(q_t^S)_{t \geqslant 0}$ are formalised by considering the random measures $J^{\mathrm{b}}(\mathrm{d}t, \mathrm{d}z)$ and $J^{\mathrm{a}}(\mathrm{d}t, \mathrm{d}z)$ modelling the times and sizes of OTC trades on the bid and ask sides, respectively. Subsequently, the inventory dynamics for the spot can be expressed as:

$$\mathrm{d}q_t^S = \int_{z=0}^{\infty} z J^{\mathrm{b}}(\mathrm{d}t, \mathrm{d}z) - \int_{z=0}^{\infty} z J^{\mathrm{a}}(\mathrm{d}t, \mathrm{d}z) + v_t^S \,\mathrm{d}t$$

while the inventory dynamics for the futures contract are modelled by:

$$\mathrm{d}q_t^F = v_t^F \,\mathrm{d}t$$

The resulting cash process $(X_t)_{t \geqslant 0}$ of the market-maker can be written as:

$$\begin{aligned}
\mathrm{d}X_t = {} & \int_{z=0}^{\infty} S^{\mathrm{a}}(t, z) z J^{\mathrm{a}}(\mathrm{d}t, \mathrm{d}z) - \int_{z=0}^{\infty} S^{\mathrm{b}}(t, z) z J^{\mathrm{b}}(\mathrm{d}t, \mathrm{d}z) \\
& - v_t^S S_t \,\mathrm{d}t - L^S(v_t^S)\,\mathrm{d}t - v_t^F F_t \,\mathrm{d}t - L^F(v_t^F)\,\mathrm{d}t
\end{aligned}$$

where the terms $L^S(v_t^S)$ and $L^F(v_t^F)$ account for spread costs along with the temporary price impact of the market-maker upon externalising.[2] The functions L^S and L^F we consider here are:

$$L^S(v) = \psi^S |v| + \eta^S v^2 \quad \text{and} \quad L^F(v) = \psi^F |v| + \eta^F v^2$$

[1] *We consider de-seasonalised futures prices with the dependence of the interest rate on time-to-maturity factored out, because interest rates typically exhibit low intraday volatility.*

[2] *For simplicity, we have not considered permanent market impact here. We might assume that both spot and futures external transactions (ie, the processes $(v_t^S)_{t \geqslant 0}$ and $(v_t^F)_{t \geqslant 0}$) affect the spot price equally, without altering the EFP spread. Under the assumption of linear permanent market impact, our model remains effective. For more details on market impact, see the paper by Hey* et al *(2024) on the importance of having a good market impact model.*

The market-maker wants to maximise the expected utility of the mark-to-market value of their portfolio at the end of the period $[0, T]$ minus a penalty corresponding to remaining inventories. We assume that the market-maker has a constant absolute risk aversion (CARA) utility function and maximises:

$$\mathbb{E}[-\exp(-\gamma(X_T + q_T^S S_T + q_T^F F_T - K^S (q_T^S)^2 - K^F (q_T^F)^2))]$$

by selecting two $\mathcal{P} \otimes \mathcal{B}(\mathbb{R}_+^*)$-measurable processes[3] δ^{b} and δ^{a} bounded from below and two $\mathcal{P}$-measurable processes v^S and v^F, with:

$$\mathbb{E}\biggl[\int_0^T (v_t^S)^2 \,\mathrm{d}t\biggr] < +\infty \quad \text{and} \quad \mathbb{E}\biggl[\int_0^T (v_t^F)^2 \,\mathrm{d}t\biggr] < +\infty$$

and with γ indicating the market-maker's risk aversion and $K^S, K^F \geqslant 0$ being the penalty coefficients.

■ **Solution.** We denote by $u \colon [0, T] \times \mathbb{R}^6 \to \mathbb{R}$ the value function of this stochastic control problem. The Hamilton-Jacobi-Bellman equation associated with it is:

$$\begin{aligned}
0 = {} & \partial_t u - k_E(E - D)\partial_E u - k_D(D - \bar{D})\partial_D u + \tfrac{1}{2} \operatorname{Tr}(\Sigma \nabla_{SED}^2 u) \\
& + \mathcal{L}^{\mathrm{b}} u + \mathcal{L}^{\mathrm{a}} u + \sup_{v^S} (v^S \partial_{q^S} u - (L^S(v^S) + v^S S)\partial_x u) \\
& + \sup_{v^F} (v^F \partial_{q^F} u - (L^F(v^F) + v^F (S + E))\partial_x u)
\end{aligned} \tag{1}$$

with terminal condition:

$$\begin{aligned}
& u(T, x, q^S, q^F, S, E, D) \\
& \quad = -\exp(-\gamma(x + q^S S + q^F (S + E) - K^S (q^S)^2 - K^F (q^F)^2))
\end{aligned}$$

where:

$$\begin{aligned}
& \mathcal{L}^{\mathrm{b/a}} u(t, x, q^S, q^F, S, E, D) \\
& \quad = \int_0^{\infty} \sup_{\delta^{\mathrm{b/a}}} f(\delta^{\mathrm{b/a}}) \bigl(u(t, x \mp z(S \mp \delta^{\mathrm{b/a}}), q^S \pm z, q^F, S, E, D) \\
& \qquad\qquad - u(t, x, q^S, q^F, S, E, D)\bigr) \lambda(z)\,\mathrm{d}z
\end{aligned}$$

This equation can be simplified through the use of the ansatz:

$$\begin{aligned}
& u(t, x, q^S, q^F, S, E, D) \\
& \quad = -\exp(-\gamma(x + q^S S + q^F (S + E) + \theta(t, q^S, q^F, E, D)))
\end{aligned}$$

where $\theta \colon [0, T] \times \mathbb{R}^4 \to \mathbb{R}$ is a differentiable function such that:

$$\theta(T, q^S, q^F, E, D) = -K^S (q^S)^2 - K^F (q^F)^2$$

The partial differential equation associated with θ is:

$$\begin{aligned}
0 = {} & \partial_t \theta - k_E(E - D)(q^F + \partial_E \theta) - k_D(D - \bar{D})\partial_D \theta \\
& + \tfrac{1}{2} \operatorname{Tr}(\tilde{\Sigma} \nabla_{ED}^2 \theta) - \frac{\gamma}{2} \begin{pmatrix} q^S + q^F \\ q^F + \partial_E \theta \\ \partial_D \theta \end{pmatrix}^{\mathrm{T}} \Sigma \begin{pmatrix} q^S + q^F \\ q^F + \partial_E \theta \\ \partial_D \theta \end{pmatrix} \\
& + \mathcal{J}_H \theta + \mathcal{H}^S(\partial_{q^S} \theta) + \mathcal{H}^F(\partial_{q^F} \theta)
\end{aligned} \tag{2}$$

[3] *Here, $\mathcal{P}$ denotes the σ-algebra of $\mathbb{F}$-predictable subsets of $\Omega \times [0, T]$ and $\mathcal{B}(\mathbb{R}_+^*)$ denotes the Borelian sets of $\mathbb{R}_+^*$.*

where $\tilde{\Sigma}$ is the submatrix of Σ obtained by removing the first row and the first column, and where:

$$\mathcal{H}^{S/F}: p \in \mathbb{R} \mapsto \sup_{v}(vp - L^{S/F}(v))$$

and:

$$\begin{aligned}\mathcal{I}_H\theta = {} & \int_0^{\infty} zH(z, \mathcal{J}^{+}\theta(t, q^S, q^F, E, D, z))\lambda(z)\,\mathrm{d}z \\ & + \int_0^{\infty} zH(z, \mathcal{J}^{-}\theta(t, q^S, q^F, E, D, z))\lambda(z)\,\mathrm{d}z\end{aligned}$$

with:

$$H: (z, p) \in (0, +\infty) \times \mathbb{R} \mapsto \sup_{\delta} \frac{f(\delta)}{\gamma z}(1 - \mathrm{e}^{-\gamma z(\delta - p)})$$

and:

$$\mathcal{J}^{\pm}\theta = \frac{\theta(t, q^S, q^F, E, D) - \theta(t, q^S \pm z, q^F, E, D)}{z}$$

Under classical assumptions on the intensities (here on the function f; see, for example, Guéant (2016)), it can be proved that, given a smooth solution to (2), the optimal controls are given by:[4]

$$\left.\begin{aligned}\delta^{\mathrm{b/a}*}(t, z) &= \bar{\delta}(z, \mathcal{J}^{\pm}\theta(t, q^S_{t-}, q^F_t, E_t, D_t, z)) \\ v_t^{S/F*} &= (\mathcal{H}^{S/F})'(\partial_{q^{S/F}}\theta(t, q^S_{t-}, q^F_t, E_t, D_t))\end{aligned}\right\} \qquad (3)$$

where $\bar{\delta}(z, p) = f^{-1}(\gamma z H(z, p) - \partial_p H(z, p))$.

Remarks on nested OU processes. Optimal quoting and hedging strategies in the above equations are contingent on current inventories and the current values of the processes $(E_t)_{t\geqslant 0}$ and $(D_t)_{t\geqslant 0}$. While the EFP and inventories are directly observable, $(D_t)_{t\geqslant 0}$ is not, introducing a challenge for practical implementation.

Assuming $(D_t)_{t\geqslant 0}$ is fixed, as in a simple OU model where $k_D = \sigma_D = 0$, leads to market-makers engaging in overly confident statistical arbitrage. This could lead to strategies that are strongly reliant on the EFP mean-reverting to the constant value of the process $(D_t)_{t\geqslant 0}$. To encourage caution, our model introduces variability in $(D_t)_{t\geqslant 0}$ by setting k_D and σ_D to positive values. These parameters can then be viewed as hyperparameters, allowing market-makers to modulate their confidence. Under that approach, the current value of the process $(D_t)_{t\geqslant 0}$ might be estimated or chosen pragmatically.

Another (more rigorous) point of view consists in statistically estimating all the parameters and filtering the signal to get an approximation of the process $(D_t)_{t\geqslant 0}$ at all times. For that purpose, we can assume a correlation structure:

$$R = \begin{pmatrix} 1 & \rho & 0 \\ \rho & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

with $\rho \in [-1, 1]$. Then, all the parameters can be estimated using classical maximum likelihood methods (without observing $(D_t)_{t\geqslant 0}$) because $(E_t)_{t\geqslant 0}$ is a Gaussian process whose covariance function is known in closed form. Once the parameters are estimated, we use filtering techniques to write:

$$\begin{aligned}\mathrm{d}S_t &= \sigma_S\,\mathrm{d}\hat{W}_t^S \\ \mathrm{d}E_t &= -k_E(E_t - \hat{D}_t)\,\mathrm{d}t + \sigma_E\,\mathrm{d}\hat{W}_t^E \\ \mathrm{d}\hat{D}_t &= -k_D(\hat{D}_t - \bar{D})\,\mathrm{d}t + \frac{1}{\sqrt{1-\rho^2}}\frac{k_E}{\sigma_E}v_t^2\,\mathrm{d}\hat{W}_t^D\end{aligned}$$

where:

$$\hat{D}_t = \mathbb{E}[D_t \mid (S_s)_{s\leqslant t}, (E_s)_{s\leqslant t}], \quad v_t^2 = \mathbb{V}(D_t \mid (S_s)_{s\leqslant t}, (E_s)_{s\leqslant t})$$

and:

$$\hat{W}_t^S = W_t^S, \quad \hat{W}_t^E = W_t^E + \frac{k_E}{\sigma_E}\int_0^t (D_s - \hat{D}_s)\,\mathrm{d}s$$
$$\hat{W}_t^D = \frac{\hat{W}_t^E - \rho\hat{W}_t^S}{\sqrt{1-\rho^2}}$$

define a three-dimensional Brownian motion adapted to the natural filtration of the processes $(S_t)_{t\geqslant 0}$ and $(E_t)_{t\geqslant 0}$ with rank 2 correlation structure:

$$\hat{R} = \begin{pmatrix} 1 & \rho & 0 \\ \rho & 1 & \sqrt{1-\rho^2} \\ 0 & \sqrt{1-\rho^2} & 1 \end{pmatrix}$$

Using standard Bayesian filtering techniques, we have that $(v_t^2)_{t\geqslant 0}$ is in fact deterministic and satisfies:

$$\frac{\mathrm{d}v_t^2}{\mathrm{d}t} = -\frac{1}{1-\rho^2}\frac{k_E^2}{\sigma_E^2}v_t^4 - 2k_D v_t^2 + \sigma_D^2$$

In particular, assuming we have observed the spot and futures prices for a long time, we can replace v_t^2 with its asymptotic value:

$$v_\infty^2 = \frac{\sigma_D^2}{k_D + \sqrt{k_D^2 + \dfrac{1}{1-\rho^2}\dfrac{k_E^2}{\sigma_E^2}\sigma_D^2}}$$

Our market-making problem can then be solved as if $(D_t)_{t\geqslant 0}$ was observed by replacing the unobservable process $(D_t)_{t\geqslant 0}$ with its observable counterpart $(\hat{D}_t)_{t\geqslant 0}$, up to the replacement[5] in the partial differential equation (2) of R by $\hat{R}$ and σ_D by:

$$\hat{\sigma}_D = \frac{1}{\sqrt{1-\rho^2}}\frac{k_E}{\sigma_E}v_\infty^2 = \sigma_D\frac{\xi}{k_D + \sqrt{k_D^2 + \xi^2}}$$

where:

$$\xi = \frac{1}{\sqrt{1-\rho^2}}k_E\frac{\sigma_D}{\sigma_E}$$

Approximation technique

Approximating numerically the solution θ of (2) using grid methods poses challenges due to the high dimensionality of the state space and the complex geometry of the frequently visited inventory states at optimality. To avoid

[4] *Because q^S is not a continuous process, we consider in (3) the left limit of q^S at time t, denoted by q^S_{t-}.*

[5] *When we filter, we remain in the same family of processes, and therefore the Hamilton-Jacobi-Bellman equation remains the same up to the value of those coefficients.*

grids, we adopt a methodology similar to that described by Bergault *et al* (2021), which involves approximating (2) with a closely related equation. The solution to that equation, denoted by $\check{\theta}$, will be a quadratic polynomial and serve as an approximation to the original value function θ. Consequently, $\check{\theta}$ will be used in lieu of θ in (3) to derive strategies that are nearly optimal.

To obtain our new equation, instead of L^S and L^F we employ the quadratic approximations $\check{L}^S(v) = \eta^S v^2$ and $\check{L}^F(v) = \eta^F v^2$. Subsequently, we replace $\mathscr{H}^S$ and $\mathscr{H}^F$ in (2) with $\check{\mathscr{H}}^{S/F}(p) = p^2/(4\eta^{S/F})$. Also, following Bergault *et al* (2021), we approximate the function H by a quadratic function:

$$\check{H}(p) = \alpha_0 + \alpha_1 p + \tfrac{1}{2}\alpha_2 p^2$$

We then obtain the following partial differential equation:

$$\begin{aligned} 0 &= \partial_t \check{\theta} - k_E(E-D)(q^F + \partial_E \check{\theta}) - k_D(D - \bar{D})\partial_D \check{\theta} \\ &\quad + \tfrac{1}{2}\operatorname{Tr}(\tilde{\Sigma}\nabla^2_{ED}\check{\theta}) - \frac{\gamma}{2}\begin{pmatrix} q^S + q^F \\ q^F + \partial_E \check{\theta} \\ \partial_D \check{\theta} \end{pmatrix}^{\mathrm{T}} \Sigma \begin{pmatrix} q^S + q^F \\ q^F + \partial_E \check{\theta} \\ \partial_D \check{\theta} \end{pmatrix} \\ &\quad + \mathscr{J}_{\check{H}}\check{\theta} + \frac{1}{4\eta^S}(\partial_{q^S}\check{\theta})^2 + \frac{1}{4\eta^F}(\partial_{q^F}\check{\theta})^2 \end{aligned} \tag{4}$$

with terminal condition $\check{\theta}(T, q^S, q^F, E, D) = -K^S(q^S)^2 - K^F(q^F)^2$.

As mentioned above, the interest of the above approximation is that the solution to (4) is a polynomial of degree 2 in $y = (q^S, q^F, E, D)^{\mathrm{T}}$. Let us apply the following ansatz:

$$\check{\theta}(t, q^S, q^F, E, D) = -y^{\mathrm{T}} A(t) y - y^{\mathrm{T}} B(t) - C(t)$$

where $A\colon [0,T] \mapsto \mathcal{S}_4(\mathbb{R})$, $B\colon [0,T] \mapsto \mathbb{R}^4$ and $C\colon [0,T] \mapsto \mathbb{R}$ are differentiable functions such that:

$$A(T) = \begin{pmatrix} -K^S & 0 & 0 & 0 \\ 0 & -K^F & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad B(T) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \text{and} \quad C(T) = 0$$

Plugging this polynomial ansatz into (4) yields a system of ODEs for A, B and C:[6]

$$\begin{aligned} A'(t) &= A(t) M^A A(t) + A(t) U^A + {U^A}^{\mathrm{T}} A(t) + R^A \\ B'(t) &= A(t) M^A B(t) + A(t) V^B + {U^A}^{\mathrm{T}} B(t) \end{aligned}$$

where:

$$M^A = \left(\begin{array}{cc|c} 4\alpha_2 \displaystyle\int_0^{+\infty} z\lambda(z)\,\mathrm{d}z + \frac{1}{\eta^S} & 0 & \\ 0 & \dfrac{1}{\eta^F} & 0_{2\times 2} \\ \hline \multicolumn{2}{c|}{0_{2\times 2}} & -2\gamma\tilde{\Sigma} \end{array}\right)$$

$$U^A = \left(\begin{array}{c|c} 0_{2\times 2} & 0_{2\times 2} \\ \hline \gamma \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \Sigma \begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 0 & 0 \end{pmatrix} & \begin{matrix} k_E & -k_E \\ 0 & k_D \end{matrix} \end{array}\right)$$

$$R^A = -\tfrac{1}{2}\gamma \left(\begin{array}{c|c} I_2 & 0_{2\times 1} \\ \hline 0_{2\times 2} & 0_{2\times 1} \end{array}\right) \Sigma \left(\begin{array}{c|c} I_2 & 0_{2\times 2} \\ \hline 0_{1\times 2} & 0_{1\times 2} \end{array}\right) - \tfrac{1}{2}k_E \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}$$

$$V^B = (0, 0, 0, -2k_D \bar{D})^{\mathrm{T}}$$

This system of ODEs can be easily solved numerically to finally get $\check{\theta}$. Then, replacing θ by $\check{\theta}$ in (3) leads to the following approximations of the optimal controls:

$$\left.\begin{aligned} \check{\delta}^{\mathrm{b/a}*}(t,z) &= \bar{\delta}(z(e^S)^{\mathrm{T}} A(t) e^S \pm 2y_{t-}^{\mathrm{T}} A(t) e^S \pm (e^S)^{\mathrm{T}} B(t)) \\ \check{v}_t^{S/F*} &= (\mathscr{H}^{S/F})'(-2(e^{S/F})^{\mathrm{T}} A(t) y_{t-} - (e^{S/F})^{\mathrm{T}} B(t)) \end{aligned}\right\} \tag{5}$$

where $e^S = (1,0,0,0)^{\mathrm{T}}$ and $e^F = (0,1,0,0)^{\mathrm{T}}$.

Remark 1 It is important to note that in (5) we employ the original functions $\bar{\delta}$, $\mathscr{H}^S$ and $\mathscr{H}^F$ without resorting to quadratic approximations. The only function approximated is θ, which is replaced by $\check{\theta}$. In the terminology of reinforcement learning, this is equivalent to approximating the true value function and then selecting greedy actions based on this approximation. In particular, the spread effects introduced by the functions L^S and L^F are duly considered.

Numerical results and discussion

For illustration, we consider market-making in spot gold (XAU/USD) with access to futures liquidity. A typical size discretisation for the pricing ladder is 100, 200, 500, 1000, 2000 and 5000 oz. We assume a standard client trade intensity represented by $\lambda(z) = (1600, 600, 1000, 600, 120, 80)$, respectively, and price sensitivity parameters are taken to be $\alpha = -0.8$ and $\beta = 5\text{bp}^{-1}$ (where 'bp' denotes basis points). The gold spot and EFP volatility parameters are set to $\sigma_S = 140\text{bp} \times \text{day}^{-1/2}$ and $\sigma_E = 5\text{bp} \times \text{day}^{-1/2}$, respectively. Spot-EFP correlation is typically small and assumed to be zero in this example. The EFP price relaxation rate is $k_E = 8\ \text{day}^{-1}$. We assume first, as a benchmark, a simple OU framework with $k_D = \sigma_D = 0$ and with $\bar{D}$ fixed at zero. The standard functional dependence of the instantaneous market impact on the execution rate is assumed: $L^S(v) = \psi^S|v| + \eta^S v^2$ and $L^F(v) = \psi^F|v| + \eta^F v^2$ with $\psi^S = 0.4\text{bp}$, $\psi^F = 0.2\text{bp}$, $\eta^S = 7 \times 10^{-8}\text{bp} \times \text{day} \times \text{oz}^{-1}$ and $\eta^F = 3 \times 10^{-8}\text{bp} \times \text{day} \times \text{oz}^{-1}$.[7] We consider a time horizon of $T = 1$ hour, which ensures convergence towards stationary quotes and hedging rates at time $t = 0$. Terminal penalty coefficients are set to zero.

Figure 1 demonstrates how the optimal market-making controls vary with the spot inventory. When there are no futures in the book, the equilibrium

[6] *We report here only the equations for A and B, as C is irrelevant for the computation of the optimal strategy.*

[7] *The parameters were selected by analysing a subset of trades by HSBC's market-making franchise. However, they should be considered as representative not of HSBC but rather of a typical institutional spot gold dealer. Daily turnover in this setup is approximately US$1 billion. The basis point convention is used for price changes, which is an approximation insignificant in practical terms given the timeframe of the problem (see Barzykin* et al *2022).*

1 Optimal gold spot pricing ladder, spot execution rate and futures execution rate as functions of the spot inventory with no futures inventory

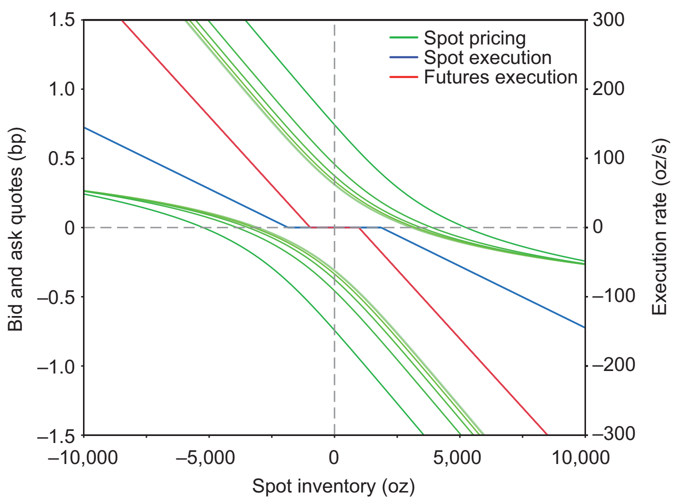

EFP price deviation is zero; risk aversion $\gamma = 3 \times 10^{-4}\ \mathrm{bp}^{-1}\ \mathrm{oz}^{-1}$; other parameters are given in the text. Different shades of green correspond to different sizes in the price ladder (lighter colours indicate a smaller size and a tighter spread)

2 Gold spot and futures no-execution zones as functions of the spot inventory and the volatility-normalised EFP price deviation

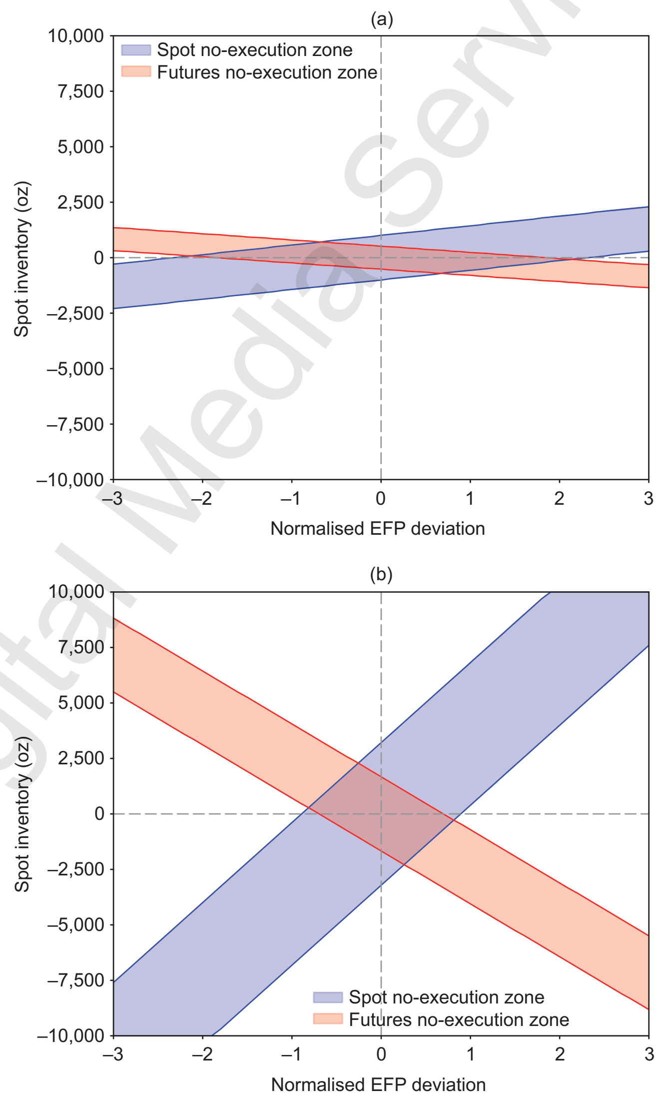

(a) $\gamma = 10^{-3}\ \mathrm{bp}^{-1}\ \mathrm{oz}^{-1}$. (b) $\gamma = 10^{-4}\ \mathrm{bp}^{-1}\ \mathrm{oz}^{-1}$. There is zero futures inventory. No execution happens in the shaded areas, with the corresponding instrument being sold above the upper boundary and bought below the lower boundary

spot position is zero, and the dealer will skew quotes to attract client flow in the direction towards this equilibrium. As noticed in other papers dealing with the internalisation-externalisation dilemma (see, for example, Barzykin *et al* 2022), there exists a pure internalisation area where skewing will be considered as the only risk-reducing option. Larger positions will, however, also involve hedging in the market. The onset for hedging with futures is earlier than for the spot, and the corresponding rate of execution is also higher for futures than for the spot. This is totally understandable given the difference in liquidity (and thus in the cost of trading). A nonzero position in futures would lead to a shifted quasi-equilibrium in the spot, corresponding to an EFP position (where each futures position is paired with another in the spot in the opposite direction).

EFP mean reversion influences optimal controls. A deviation of the EFP spread from the expected mean will shift the equilibrium inventory along with the corresponding execution onsets and quote skew, as shown in figures 2 and 3. Understandably, when the deviation is negative the dealer will tend to accumulate a long EFP position, and vice versa. Figure 2 demonstrates that EFP mean reversion will lead to rare direct arbitrage opportunities only under strong risk aversion. This is related to the cost of opportunistically entering into an EFP position (two spreads have to be crossed). The dealer will instead skew quotes in the required direction (see figure 3) and wait for the opportunity to materialise while making a spread. With a lower risk aversion, the appetite to capitalise on EFP mean reversion increases, leading to opportunistic execution at extreme deviations (where the upper/lower futures execution onset is below/above the lower/upper spot execution onset. The skew will be repurposed from EFP risk management to opportunistic skew in response to EFP deviation.

We introduced nested OU processes in this paper to mitigate this repurposing, which is mainly due to overconfidence when using the OU model. Uncertainty over the mean EFP deviation, through the introduction of the stochastic process $(D_t)_{t \geq 0}$, creates additional risk that is taken into account through risk aversion, resulting in the dealer being less inclined to keep an EFP position. Figure 4 demonstrates the effect of σ_D on the propensity to capitalise on EFP mean reversion. Here, $k_D = 0.2\ \mathrm{day}^{-1}$, $\gamma = 3 \times 10^{-4}\ \mathrm{bp}^{-1}\ \mathrm{oz}^{-1}$ and other parameters are as defined earlier. As expected, the larger σ_D is, the smaller the skew and the higher the arbitrage onset.

Concluding remarks

We have extended a stochastic optimal control framework for OTC market-making by incorporating co-integrated liquidity for hedging. Using a computationally efficient approximation technique, our methodology facilitates strategy optimisation on demand in near real-time, potentially benefiting both electronic and voice traders. A specific example of the spot gold market is analysed in detail, demonstrating efficient risk management that benefits

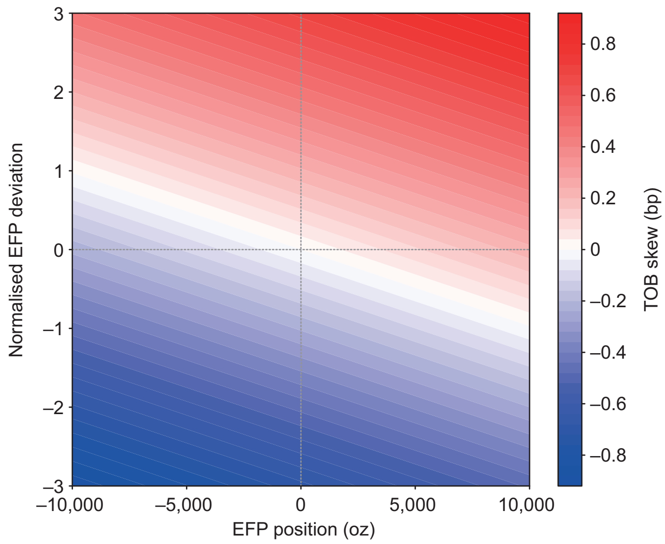

3 Gold top-of-book spot skew as a function of the EFP position and the volatility-normalised EFP price deviation

There is zero futures inventory, and $\gamma = 10^{-3}$ bp^{-1} oz^{-1}

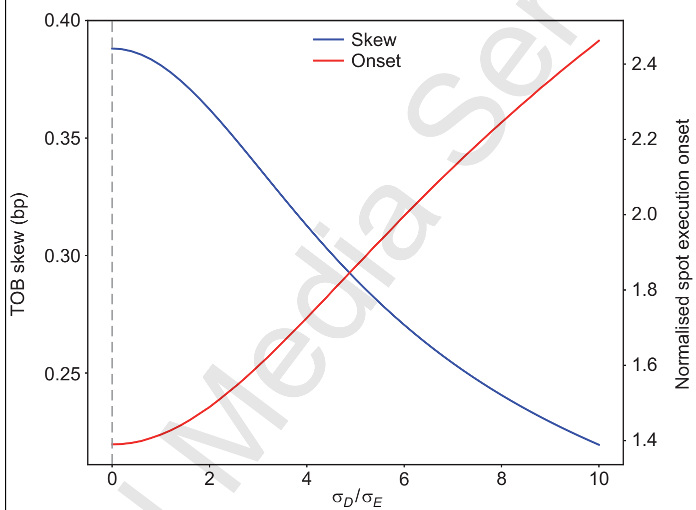

4 Top-of-book skew for the volatility-normalised EFP deviation of $\varepsilon \equiv E/\sigma_E = 1$ and $D = 0$ along with the value of ε corresponding to the spot execution onset as functions of the normalised volatility of the mean, σ_D/σ_E, for zero spot and futures inventory and $D = 0$

from access to the significantly more liquid gold futures with reduced transaction costs while capitalising on EFP mean reversion. In the spirit of the two-factor Hull-White model for interest rates, the EFP spread is modelled by a nested Ornstein-Uhlenbeck process describing the observed multiple modes of relaxation corresponding to the diverse trading horizons of market participants. Interestingly, nested OU processes also appear as a way to 'robustify' pure OU strategies with uncertainty in the mean-reversion parameters, a commonly encountered case. They could find application in generalising most of the optimisation problems involving classical OU processes (see, for instance, the trading problem of Lipton & Lopez de Prado (2020)). ■

Alexander Barzykin is a director at HSBC GFX and commodities, based in London. Philippe Bergault is an associate professor in applied mathematics at Université Paris Dauphine-PSL. Olivier Guéant is a full professor of applied mathematics at Université Paris 1 Panthéon-Sorbonne and an adjunct professor of quantitative finance at ENSAE – IP Paris. The results presented in this paper are part of the research work carried out within the HSBC Research Initiative. The views expressed are those of the authors and do not necessarily reflect the views or the practices at HSBC. The authors are grateful to Richard Anthony, James Donaldson and Manuel Abellan-Lopez (HSBC) for helpful discussions and support throughout the project.
Email: alexander.barzykin@hsbc.com, bergault.philippe@protonmail.com, olivier.gueant@univ-paris1.fr.

REFERENCES

Avellaneda M and S Stoikov, 2008
High-frequency trading in a limit order book
Quantitative Finance 8(3), pages 217–224

Barzykin A, P Bergault and O Guéant, 2022
Market-making by a foreign exchange dealer
Risk September, www.risk.net/7952481

Barzykin A, P Bergault and O Guéant, 2023
Dealing with multi-currency inventory risk in foreign exchange cash markets
Risk March, www.risk.net/7956238

Bergault P, D Evangelista, O Guéant and D Vieira, 2021
Closed-form approximations in multi-asset market making
Applied Mathematical Finance 28(2), pages 101–142

Hull JC and A White, 1994
Numerical procedures for implementing term structure models. II. Two-factor models
Journal of Derivatives 2(2), pages 37–48

Butz M and R Oomen, 2019
Internalisation by electronic FX spot dealers
Quantitative Finance 19(1), pages 35–56

Cartea Á, S Jaimungal and J Ricci, 2014
Buy low, sell high: a high frequency trading perspective
SIAM Journal on Financial Mathematics 5(1), pages 415–444

Cartea Á, S Jaimungal and T Jia, 2020
Trading foreign exchange triplets
SIAM Journal on Financial Mathematics 11(3), pages 690–719

Guéant O, 2016
The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making (volume 33)
CRC Press

Guilbaud F and H Pham, 2013
Optimal high-frequency trading with limit and market orders
Quantitative Finance 13(1), pages 79–94

Hey N, J-P Bouchaud, I Mastromatteo, J Muhle-Karbe and K Webster, 2024
The cost of misspecifying price impact
Risk January, www.risk.net/7958754

Lipton A and M Lopez de Prado, 2020
A closed-form solution for optimal Ornstein-Uhlenbeck driven trading strategies
International Journal of Theoretical and Applied Finance 23(08), paper 2050056