

Algorithmic Market Making in Spot Precious Metals

Alexander BARZYKIN*

Philippe BERGAULT†

Olivier GUÉANT‡

Abstract

The primary challenge of market making in spot precious metals is navigating the liquidity that is mainly provided by futures contracts. The Exchange for Physical (EFP) spread, which is the price difference between futures and spot, plays a pivotal role and exhibits multiple modes of relaxation corresponding to the diverse trading horizons of market participants. In this paper, we introduce a novel framework utilizing a nested Ornstein-Uhlenbeck process to model the EFP spread. We demonstrate the suitability of the framework for maximizing the expected P&L of a market maker while minimizing inventory risk across both spot and futures. Using a computationally efficient technique to approximate the solution of the Hamilton-Jacobi-Bellman equation associated with the corresponding stochastic optimal control problem, our methodology facilitates strategy optimization on demand in near real-time, paving the way for advanced algorithmic market making that capitalizes on the co-integration properties intrinsic to the precious metals sector.

Key words: Market making, precious metals, internalization-externalization dilemma, hedging, stochastic optimal control, Riccati equations, closed-form approximations.

Introduction

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 [9] for a good example using stochastic optimal control), but OTC markets have also been addressed using the modelling framework introduced by Avellaneda and Stoikov in [1] (see for instance [3] for the FX market). These market making models have helped dealers to 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 maximize client flow internalization to mitigate the cost of external execution and reduce market impact, hedging only when inventories exceed franchise-dependent thresholds (on the internalization-externalization dilemma, see for instance [2] and [5]).

Multi-asset extensions of market making models have been built to help dealers manage their risk at the portfolio level (see for instance [3] and [4]). Market makers often deal indeed with large portfolios of assets of very different liquidity. Illiquid assets may be difficult to internalize 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. 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 capitalizes on the unique benefits of using co-integrated assets for hedging, regardless of whether the market

*HSBC, 8 Canada Square, Canary Wharf, London E14 5HQ, United Kingdom, alexander.barzykin@hsbc.com.

†Université Paris Dauphine-PSL, Ceremade, Place du Maréchal de Lattre de Tassigny, 75116 Paris, France, bergault@ceremade.dauphine.fr.

‡Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, olivier.gueant@univ-paris1.fr.

maker proposes quotes in those assets. Co-integration offers indeed a special opportunity for market makers to tap into enhanced liquidity pools and benefit from mean reversion. This is particularly relevant in the case of a spot dealer hedging their position through futures, as in the case of precious metals market. While the interbank spot market does exist, precious metals futures are considerably more liquid, with tighter spreads and higher volumes. Since the price difference between futures and spot – so-called Exchange for Physical (EFP) spread – is primarily driven by swap rates which are relatively stable over the market maker’s intraday risk horizon, the co-integration assumption between spot and futures prices is very natural.¹

Our modelling framework builds on the classical foundation laid by [1] and [7] (see the two reference books [6] and [8] for a detailed discussion). In this framework, a single underlying instrument (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 spot who can hedge (with costs) in both spot and futures. The novelty comes from the modelling of EFP by a nested Ornstein-Uhlenbeck (hereafter nested OU) process, as inspired by the observation of multiple relaxation times in the market ranging from hours to days, possibly related to different trading horizons of different types of traders, as illustrated in Figure 1. This article expands, therefore, 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 also provides a new class of price dynamics compatible with the approximation techniques developed in [4].




Figure 1: EFP spreads in basis points implied from four active futures contracts and spot mid prices of gold against USD in 2023 (green). Zoom in on 21 July also shows OTC forward rate – OTC FWD – (red) and demonstrates intraday mean reversion. Daily median difference in basis points between implied EFP and OTC forward rate (blue) illustrates mean reversion on a weekly scale.

¹There are other asset classes where the same mathematical ideas could be used. In FX, one can think of market making involving Non-Deliverable Forwards and onshore spot liquidity for instance.

In the forthcoming 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 utilizing nested OU processes. Subsequently, we demonstrate the applicability of the method from [4] 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 optimal control problem

We consider a spot market maker in a given precious metal. We denote by $(S_t)_t$ a reference spot price for that metal and assume that the market maker continuously streams to clients a pricing ladder at the bid, $S^b(t, z) = S_t - \delta^b(t, z)$, and at the ask, $S^a(t, z) = S_t + \delta^a(t, z)$, where z denotes the sizes 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 bid and ask: $(z, \delta) \mapsto \Lambda^b(z, \delta)$ and $(z, \delta) \mapsto \Lambda^a(z, \delta)$. We assume that the functions Λ^b and Λ^a take the form²

$$\Lambda^b(z, \delta) = \Lambda^a(z, \delta) = \Lambda(z, \delta) = \lambda(z)f(\delta) \quad \text{with} \quad f(\delta) = \frac{1}{1 + e^{\alpha + \beta\delta}}.$$

The market maker is also capable of hedging their position by trading on various platforms, using either or both spot and futures (the most liquid contract is typically chosen). The execution rates associated with the market maker's activities are modeled by the processes $(v_t^S)_t$ and $(v_t^F)_t$ for spot and futures, respectively.

We denote by $(F_t)_t$ the price process of the futures contract and by $(E_t)_t$ the EFP process, i.e. $F_t = S_t + E_t$.³ In what follows, we consider a Brownian dynamics for $(S_t)_t$ and a nested OU dynamics⁴ for $(E_t)_t$, i.e.,

$$dS_t = \sigma_S dW_t^S, \quad \sigma_S > 0$$

and

$$dE_t = -k_E (E_t - D_t) dt + \sigma_E dW_t^E, \quad k_E, \sigma_E > 0,$$

with

$$dD_t = -k_D (D_t - \bar{D}) dt + \sigma_D dW_t^D, \quad k_D, \sigma_D \geq 0, \quad \bar{D} \in \mathbb{R},$$

where $(W_t^S, W_t^E, W_t^D)_t$ is a three-dimensional Brownian motion with correlation matrix R . In what follows,

we denote by Σ the variance-covariance matrix $\begin{pmatrix} \sigma_S & 0 & 0 \\ 0 & \sigma_E & 0 \\ 0 & 0 & \sigma_D \end{pmatrix} R \begin{pmatrix} \sigma_S & 0 & 0 \\ 0 & \sigma_E & 0 \\ 0 & 0 & \sigma_D \end{pmatrix}$.

We denote the inventory process of the market maker in spot by $(q_t^S)_t$ and the inventory in the futures contract by $(q_t^F)_t$. Mathematically, the dynamics of $(q_t^S)_t$ is formalized by considering the random measures $J^b(dt, dz)$ and $J^a(dt, dz)$ modeling the times and sizes of OTC trades on the bid and ask sides, respectively. Subsequently, the inventory dynamics for spot can be expressed as

$$dq_t^S = \int_{z=0}^{\infty} z J^b(dt, dz) - \int_{z=0}^{\infty} z J^a(dt, dz) + v_t^S dt,$$

while the dynamics for the inventory in the futures contract is modelled by

$$dq_t^F = v_t^F dt.$$

The resulting cash process $(X_t)_t$ of the market maker writes

$$dX_t = \int_{z=0}^{\infty} S^a(t, z) z J^a(dt, dz) - \int_{z=0}^{\infty} S^b(t, z) z J^b(dt, dz) - v_t^S S_t dt - L^S(v_t^S) dt - v_t^F F_t dt - L^F(v_t^F) dt,$$

²Generalizations do not raise difficulties.

³We consider de-seasonalized futures prices with interest rate dependence on time to maturity factored out.

⁴It is possible to expand the nested dynamics to include more than one level if there is statistical or practical need.

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 externalizing.⁵ The functions L^S and L^F are typically nonnegative, strictly convex, and asymptotically super-linear. Here, we consider

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

The market maker wants to maximize 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 CARA utility function, and maximizes

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

by selecting δ^b , δ^a , v^S and v^F , where γ indicates the market maker's risk aversion and $K \geq 0$ is a penalty constant.

Solution

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

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

with terminal condition $u(T, x, q^S, q^F, S, E, D) = -\exp \left(-\gamma \left(x + q^S S + q^F (S + E) - K \left((q^S)^2 + (q^F)^2 \right) \right) \right)$, where

$$\begin{aligned} \mathcal{L}^b u(t, x, q^S, q^F, S, E, D) &= \int_0^\infty \sup_{\delta^b} f(\delta^b) \left(u(t, x - z(S - \delta^b), q^S + z, q^F, S, E, D) - u(t, x, q^S, q^F, S, E, D) \right) \lambda(z) dz \\ \mathcal{L}^a u(t, x, q^S, q^F, S, E, D) &= \int_0^\infty \sup_{\delta^a} f(\delta^a) \left(u(t, x + z(S + \delta^a), q^S - z, q^F, S, E, D) - u(t, x, q^S, q^F, S, E, D) \right) \lambda(z) dz. \end{aligned}$$

This equation can be simplified through the use of the ansatz

$$u(t, x, q^S, q^F, S, E, D) = -\exp \left(-\gamma \left(x + q^S S + q^F (S + E) + \theta(t, q^S, q^F, E, D) \right) \right),$$

where $\theta : [0, T] \times \mathbb{R}^4 \rightarrow \mathbb{R}$ is a differentiable function such that $\theta(T, q^S, q^F, E, D) = -K \left((q^S)^2 + (q^F)^2 \right)$.

The partial differential equation associated with θ is indeed the following:

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

where $\tilde{\Sigma}$ is the submatrix of Σ obtained by removing the first row and the first column, $\mathcal{H}^S$ and $\mathcal{H}^F$ are the Hamiltonian functions defined by

$$\mathcal{H}^S : p \in \mathbb{R} \mapsto \sup_{v^S} (v^S p - L^S(v^S)) \quad \text{and} \quad \mathcal{H}^F : p \in \mathbb{R} \mapsto \sup_{v^F} (v^F p - L^F(v^F)),$$

⁵We have not considered permanent market impact here for the sake of simplicity. We might assume that both spot and futures external transactions (i.e. the processes $(v_t^S)_t$ and $(v_t^F)_t$) 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 reference book [12] and the recent paper [10] on the importance to have a good market impact model.

and

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

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

Under classical assumptions on the intensities (here on the function f – see for instance [8]), it can be proved that, given a smooth solution to the equation (2), the optimal controls are given by⁶

$$\begin{cases} \delta^{b*}(t, z) = \bar{\delta} \left(z, \frac{\theta(t, q_{t-}^S, q_t^F, E_t, D_t) - \theta(t, q_{t-}^S + z, q_t^F, E_t, D_t)}{z} \right) \\ \delta^{a*}(t, z) = \bar{\delta} \left(z, \frac{\theta(t, q_{t-}^S, q_t^F, E_t, D_t) - \theta(t, q_{t-}^S - z, q_t^F, E_t, D_t)}{z} \right) \\ v_t^{S*} = \mathcal{H}^{S'}(\partial_{q^S}\theta(t, q_{t-}^S, q_t^F, E_t, D_t)) \\ v_t^{F*} = \mathcal{H}^{F'}(\partial_{q^F}\theta(t, q_{t-}^S, q_t^F, E_t, D_t)) \end{cases} \quad (3)$$

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 our model are contingent upon current inventories and the current values of the processes $(E_t)_t$ and $(D_t)_t$. While the EFP and inventories are directly observable, $(D_t)_t$ is not, introducing a challenge for practical implementation. However, assuming $(D_t)_t$ 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 indeed to strategies that are strongly reliant on the EFP mean reverting to the constant value of the process $(D_t)_t$.

To encourage caution, our model introduces variability in $(D_t)_t$ 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. With that vision, the current value of the process $(D_t)_t$ might be estimated or pragmatically chosen.

Another point of view consists in statistically estimating all the parameters and filtering the signal to get an approximation of $(D_t)_t$ 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 (without observing $(D_t)_t$) methods because $(E_t)_t$ is a Gaussian process whose covariance function is known in closed form. Once parameters are estimated, we use filtering techniques to write

$$\begin{cases} dS_t &= \sigma_S d\widehat{W}_t^S \\ dE_t &= -k_E (E_t - \widehat{D}_t) dt + \sigma_E d\widehat{W}_t^E \\ d\widehat{D}_t &= -k_D (\widehat{D}_t - \bar{D}) dt + \frac{1}{\sqrt{1-\rho^2}} \frac{k_E}{\sigma_E} \nu_t^2 d\widehat{W}_t^D \end{cases}$$

where

$$\widehat{D}_t = \mathbb{E}[D_t | (S_s)_{s \leq t}, (E_s)_{s \leq t}], \quad \nu_t^2 = \mathbb{V}(D_t | (S_s)_{s \leq t}, (E_s)_{s \leq t}),$$

and

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

⁶Because q^S is not a continuous process, we consider in (5) the left limit of q^S at time t denoted by q_{t-}^S .

define a three-dimensional Brownian motion adapted to the natural filtration of the processes $(S_t)_t$ and $(E_t)_t$ with correlation structure $\widehat{R} = \begin{pmatrix} 1 & \rho & 0 \\ \rho & 1 & \sqrt{1-\rho^2} \\ 0 & \sqrt{1-\rho^2} & 1 \end{pmatrix}$ of rank 2.

Using standard Bayesian filtering techniques, $(\nu_t^2)_t$ is in fact deterministic and satisfies

$$\frac{d\nu_t^2}{dt} = -\frac{1}{1-\rho^2} \frac{k_E^2}{\sigma_E^2} \nu_t^4 - 2k_D \nu_t^2 + \sigma_D^2.$$

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

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

Our market making problem can then be solved by replacing the unobservable process $(D_t)_t$ by the observable process $(\widehat{D}_t)_t$ up to the replacement in equation (2) of R by $\widehat{R}$ and σ_D by

$$\widehat{\sigma}_D = \frac{1}{\sqrt{1-\rho^2}} \frac{k_E}{\sigma_E} \nu_\infty^2 = \sigma_D \frac{\xi}{k_D + \sqrt{k_D^2 + \xi^2}}, \quad \text{where} \quad \xi = \frac{1}{\sqrt{1-\rho^2}} k_E \frac{\sigma_D}{\sigma_E}.$$

Approximation technique

Approximating numerically the solution θ of equation (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 grids, we adopt a methodology similar to that described in [4], which involves approximating equation (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 utilized in lieu of θ in (5) to derive strategies that are nearly optimal.

To obtain our new equation, we employ instead of L^S and L^F the quadratic approximations $\check{L}^S(v) = \eta^S v^2$ and $\check{L}^F(v) = \eta^F v^2$. Subsequently, we substitute $\mathcal{H}^S$ and $\mathcal{H}^F$ in (2) by

$$\check{\mathcal{H}}^S : p \in \mathbb{R} \mapsto \sup_{v^S} (v^S p - \check{L}^S(v^S)) = \frac{p^2}{4\eta^S} \quad \text{and} \quad \check{\mathcal{H}}^F : p \in \mathbb{R} \mapsto \sup_{v^F} (v^F p - \check{L}^F(v^F)) = \frac{p^2}{4\eta^F}.$$

Also, following [4], we approximate the function H by a quadratic function

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

We then obtain the following partial differential equation:

$$\begin{aligned} 0 = & \partial_t \check{\theta} - k_E (E - D) (q^F + \partial_E \check{\theta}) - k_D (D - \bar{D}) \partial_D \check{\theta} + \frac{1}{2} \text{Tr}(\widetilde{\Sigma} \nabla_{ED}^2 \check{\theta}) \\ & - \frac{\gamma}{2} \begin{pmatrix} q^S + q^F \\ q^F + \partial_E \check{\theta} \\ \partial_D \check{\theta} \end{pmatrix}^\top \Sigma \begin{pmatrix} q^S + q^F \\ q^F + \partial_E \check{\theta} \\ \partial_D \check{\theta} \end{pmatrix} + \mathcal{J}_{\check{H}} \check{\theta} + \frac{1}{4\eta^S} (\partial_{q^S} \check{\theta})^2 + \frac{1}{4\eta^F} (\partial_{q^F} \check{\theta})^2 \end{aligned} \quad (4)$$

with terminal condition $\check{\theta}(T, q^S, q^F, E, D) = -K((q^S)^2 + (q^F)^2)$, whose solution $\check{\theta}$ will be our approximation of θ .

As announced above, the interest of the above approximation is that the solution to (4) is a polynomial of degree 2 in q^S , q^F , E , and D . Let us make indeed the following ansatz:

$$\check{\theta}(t, q^S, q^F, E, D) = - \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix}^\top A(t) \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix} - \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix}^\top B(t) - C(t),$$

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

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

Plugging this polynomial ansatz into (4) yields a system of ODEs for A , B , and C :⁷

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

where

$$M^A = \left(\begin{array}{c|c} 4\alpha_2 \int_0^{+\infty} z\lambda(z)dz + \frac{1}{\eta^S} & 0 \\ \hline 0 & \frac{1}{\eta^F} \end{array} \middle| \begin{array}{c} 0_{2 \times 2} \\ -2\gamma\tilde{\Sigma} \end{array} \right), \quad U^A = \left(\begin{array}{c|c} 0_{2 \times 2} & 0_{2 \times 2} \\ \hline \gamma \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \Sigma \begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 0 & 0 \end{pmatrix} & \begin{pmatrix} k_E & -k_E \\ 0 & k_D \end{pmatrix} \end{array} \right),$$

$$R^A = -\frac{1}{2}\gamma \left(\begin{array}{c|c} I_2 & 0_{2 \times 1} \\ \hline 0_{2 \times 2} & 0_{2 \times 1} \end{array} \right) \Sigma \left(\begin{array}{c|c} I_2 & 0_{2 \times 2} \\ \hline 0_{1 \times 2} & 0_{1 \times 2} \end{array} \right) - \frac{1}{2}k_E \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad V^B = \begin{pmatrix} 0 \\ 0 \\ 0 \\ -2k_D \bar{D} \end{pmatrix}.$$

This system of ODEs can easily be solved numerically. It leads to the following approximations of the optimal controls:

$$\left\{ \begin{array}{l} \check{\delta}^{b*}(t, z) = \bar{\delta} \left(z(e^S)^\top A(t)e^S + 2 \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix}^\top A(t)e^S + (e^S)^\top B(t) \right) \\ \check{\delta}^{a*}(t, z) = \bar{\delta} \left(z(e^S)^\top A(t)e^S - 2 \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix}^\top A(t)e^S - (e^S)^\top B(t) \right) \\ \check{v}_t^{S*} = \mathcal{H}^{S'} \left(-2(e^S)^\top A(t) \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix} - (e^S)^\top B(t) \right) \\ \check{v}_t^{F*} = \mathcal{H}^{F'} \left(-2(e^F)^\top A(t) \begin{pmatrix} q^S \\ q^F \\ E \\ D \end{pmatrix} - (e^F)^\top B(t) \right), \end{array} \right. \quad (5)$$

where $e^S = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$ and $e^F = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}$.

Remark 1. It is important to note that in (5), we employ the original functions $\bar{\delta}$, $\mathcal{H}^S$ and $\mathcal{H}^F$ 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.

⁷We only report here the equations for A and B , as C is irrelevant for the computation of the optimal strategy.

Numerical results and discussion

For illustration, we consider market making in spot gold (XAUUSD) with access to futures liquidity.

A typical pricing ladder size discretization is set as 100, 200, 500, 1000, 2000 and 5000 oz. We assume a standard client trade intensity represented by

$$\lambda(z) = \begin{cases} 1600 & \text{if } z = 100 \\ 600 & \text{if } z = 200 \\ 1000 & \text{if } z = 500 \\ 600 & \text{if } z = 1000 \\ 120 & \text{if } z = 2000 \\ 80 & \text{if } z = 5000 \end{cases} \quad \text{and} \quad f(\delta) = \frac{1}{1 + e^{\alpha + \beta\delta}}$$

where price sensitivity parameters are taken to be $\alpha = -0.8$ and $\beta = 5 \text{ bp}^{-1}$. Gold spot and EFP volatility parameters are set to $\sigma_S = 140 \text{ bp} \cdot \text{day}^{-1/2}$ and $\sigma_E = 5 \text{ bp} \cdot \text{day}^{-1/2}$, respectively. Spot-EFP correlation is typically small and assumed to be zero in this example. 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 $\bar{D}$ is fixed at zero. Standard functional dependence of instantaneous market impact on execution rate is assumed: $L^S(v) = \phi^S |v| + \eta^S v^2$ and $L^F(v) = \phi^F |v| + \eta^F v^2$ with $\phi^S = 0.4 \text{ bp}$, $\phi^F = 0.2 \text{ bp}$, $\eta^S = 7 \cdot 10^{-8} \text{ bp} \cdot \text{day} \cdot \text{oz}^{-1}$ and $\eta^F = 3 \cdot 10^{-8} \text{ bp} \cdot \text{day} \cdot \text{oz}^{-1}$. We consider a time horizon of $T = 1$ hour that ensures convergence towards stationary quotes and hedging rates at time $t = 0$. Terminal penalty is set to zero unless specified otherwise.⁸




Figure 2: Optimal gold spot pricing ladder (green), spot execution rate (blue) and futures execution rate (red) as functions of spot inventory for zero (left) and 1000 oz (right) futures inventory. EFP price deviation is zero, risk aversion $\gamma = 3 \cdot 10^{-4}$, other parameters in the text.

Figure 2 demonstrates the variation of optimal market making controls with spot inventory. When there are no futures in the book, the equilibrium spot position is zero, and the dealer will skew quotes to attract client flow in the direction towards this equilibrium. As noticed previously (see [2]), there exists a pure internalization area where skewing will be considered as the only risk reducing option. Larger positions will also involve hedging in the market. We can see that the onset for hedging with futures is earlier and the corresponding rate of execution is higher for futures than for spot. This is totally understandable given

⁸These parameters have been selected by analyzing a subset of HSBC market making franchise. However, they should not be considered as representative of HSBC but rather of a typical institutional spot gold dealer. Daily turnover in this setup is approximately \$1B. Basis point convention is used for price changes, which is an approximation insignificant in practical terms given the time frame of the problem [2].

the liquidity (and thus the cost of trading) difference. A non-zero position in futures leads to a shifted quasi-equilibrium in spot, corresponding to an EFP position (where each futures position is paired with another in spot in the opposite direction). It is optimal for the dealer to unwind the underlying risk towards this point when there is an unbalanced futures or spot position.




Figure 3: Average gold spot (blue) and futures (red) inventory relaxation following a 1000 oz client spot trade with a termination condition in 1 hour obtained by numerically averaging $2 \cdot 10^4$ Monte Carlo trajectories. The dashed line (yellow) shows the dynamics of $q^S + q^F$.




Figure 4: Inventory probability distribution of a spot gold market maker with access to futures hedging.

Figure 3 illustrates the expected inventory relaxation following a relatively large 1000 oz OTC trade for $\gamma = 3 \cdot 10^{-4}$. It was calculated via Monte Carlo simulation ($2 \cdot 10^4$ trajectories) of the model given the initial condition and optimal controls with terminal penalty $K = 10^{-3}$. We see that part of the risk is hedged with futures very quickly. The remaining excess spot risk is unwound via skew (attracting the offsetting client flow) and spot execution. Figure 2 shows that skew can be quite aggressive for this level of risk aversion, with offers better than mid price. Effectively, during this fast relaxation stage, the dealer creates an EFP position which is later unwound very slowly. Figure 4 shows an inventory probability distribution extracted from a sufficiently long Monte Carlo trajectory (10^7 seconds). As expected, the system spends most of the time along the low-risk EFP diagonal.




Figure 5: Gold spot and futures no-execution zones as functions of spot inventory and volatility-normalised EFP price deviation. Zero futures inventory, $\gamma = 10^{-3}$ (left) and 10^{-4} (right). No execution in shaded areas, selling the corresponding instrument above the upper boundary and buying below the lower boundary.




Figure 6: Gold top of book spot skew as a function of EFP position and volatility-normalised EFP price deviation. $\gamma = 10^{-3}$ (left) and 10^{-4} (right).

EFP mean reversion influences optimal controls. EFP deviation from the expected mean shifts the equilibrium inventory along with the corresponding execution onsets and quote skew, as shown in Figures 5 and 6. Understandably, when the deviation is negative the dealer will tend to accumulate long EFP position, and vice versa. Figure 5 demonstrates that EFP mean reversion will only lead to rare direct arbitrage opportunities under strong risk aversion. This is related to the cost of opportunistically entering into an EFP position (one would have to cross two spreads). The dealer will instead skew quotes in the required direction (see Figure 6) and wait for the opportunity to materialize while making spread. With a lower risk aversion, the appetite to capitalize on EFP mean reversion increases leading to opportunistic execution at extreme deviations (where the upper futures execution onset is below the lower spot execution onset and, similarly, where the lower futures execution onset is above the upper spot execution onset). The skew will be repurposed from EFP risk management to opportunistic skew on EFP deviation, as shown in Figure 6.

We introduced nested OU processes in this paper to mitigate the latter repurposing which is mainly due to overconfidence in the OU model. Uncertainty on mean EFP deviation through the introduction of the stochastic process $(D_t)_t$ creates additional risk that is taken into account by risk aversion, resulting in the dealer being less inclined to keep EFP position. Figure 7 demonstrates the effect of σ_D on the propensity to capitalize on EFP mean reversion. Here $k_D = 0.2 \text{ day}^{-1}$, $\gamma = 3 \cdot 10^{-4}$, and other parameters are as defined earlier. As expected, the larger is σ_D , the smaller is the skew and the higher is the arbitrage onset.




Figure 7: Top of book skew for the volatility-normalized EFP deviation of $\epsilon \equiv E/\sigma_E = 1$ and $D = 0$ along with the value of ϵ corresponding to spot execution onset as functions of normalized volatility of the mean, σ_D/σ_E , for zero spot and futures inventory and $D = 0$.

The choice of risk aversion is ultimately in the hands of the dealer. The rule of thumb is to decrease γ while the expected benefit in P&L exceeds the expected increase in risk. Figure 8 demonstrates the effect of risk aversion on volume share of hedging, P&L and risk. In this specific example, it is clear that decreasing γ below 10^{-4} is questionable as the expected risk increases much faster than P&L. Similarly, increasing γ above 10^{-3} is not efficient.




Figure 8: Volume share of client trades and hedging, hourly P&L and standard deviation of hourly P&L as functions of risk aversion.




Figure 9: Optimal market making strategy performance on real market data (XAUUSD, 12 January 2024) with simulated client flow and risk aversion of $\gamma = 3 \cdot 10^{-4}$. Top chart displays implied EFP spread and dealer's EFP position during the day, bottom chart shows realised cumulative P&L with open positions marked to market mid. Insert illustrates volume share of OTC and hedging trades.

Figure 9 illustrates the performance of the optimal strategy over a single trading day on real market data but with simulated client flow (uniform intensity throughout the day). EFP position clearly echoes the mean-reverting nature of EFP – the dealer aims to keep EFP risk against the direction of EFP deviation. Volume share of futures execution significantly exceeds that of spot due to lower cost. Introducing futures hedging into risk management of OTC spot improves Sharpe ratio by at least 30% (not shown).

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 optimization on demand in near real-time. A specific example of spot gold market is analyzed in detail demonstrating efficient risk management benefiting from access to the significantly more liquid gold futures with reduced transaction costs while capitalizing on EFP mean reversion. EFP spread is modeled by a nested OU 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 generalizing most of the optimization problems involving classical OU processes (see for instance the trading problem of [11]).

Statement and acknowledgement

The results presented in this paper are part of the research works 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.

References

- [1] Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [2] Alexander Barzykin, Philippe Bergault, and Olivier Guéant. Market-making by a foreign exchange dealer. *Risk*, September 2022.
- [3] Alexander Barzykin, Philippe Bergault, and Olivier Guéant. Dealing with multi-currency inventory risk in foreign exchange cash markets. *Risk*, March 2023.
- [4] Philippe Bergault, David Evangelista, Olivier Guéant, and Douglas Vieira. Closed-form approximations in multi-asset market making. *arXiv preprint arXiv:1810.04383*, 2020.
- [5] Maximilian Butz and Roel Oomen. Internalisation by electronic fx spot dealers. *Quantitative Finance*, 19(1):35–56, 2019.
- [6] Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and high-frequency trading*. Cambridge University Press, 2015.
- [7] Álvaro Cartea, Sebastian Jaimungal, and Jason Ricci. Buy low, sell high: A high frequency trading perspective. *SIAM Journal on Financial Mathematics*, 5(1):415–444, 2014.
- [8] Olivier Guéant. *The Financial Mathematics of Market Liquidity: From optimal execution to market making*, volume 33. CRC Press, 2016.
- [9] Fabien Guilbaud and Huyen Pham. Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1):79–94, 2013.
- [10] Natascha Hey, Jean-Philippe Bouchaud, Iacopo Mastromatteo, Johannes Muhle-Karbe, and Kevin Webster. The cost of misspecifying price impact. *Risk*, January 2024.
- [11] Alexander Lipton and Marcos Lopez de Prado. A closed-form solution for optimal ornstein-uhlenbeck driven trading strategies. *International Journal of Theoretical and Applied Finance*, 23(08):2050056, 2020.
- [12] Kevin T Webster. *Handbook of Price Impact Modeling*. Chapman and Hall/CRC, 2023.