Making a Market: the HJB a Naive Quoter Is Ignoring

Inventory Risk, the Equation That Prices It, and the Quotes That Follow

T. Zamrik

20-SEP-2026

Abstract. A market maker who posts a fixed spread around the mid earns that spread on every fill and carries whatever inventory the fills leave behind. We treat the two halves of that sentence as two problems. The naive symmetric quoter is solved in closed form: its expected profit is $2AT\delta e^{-k\delta}$, maximised at a half-spread of $1/k$, and its inventory is a symmetric random walk whose variance grows linearly in the horizon, so that inventory risk outweighs spread-capture risk by a factor $\sigma^2 T/2\delta^2$ that grows without bound. The optimal quoter is the solution of a Hamilton–Jacobi–Bellman equation which, under an exponential-utility ansatz, collapses to a **linear** system of ordinary differential equations. We solve that system exactly rather than asymptotically, and two things follow. The optimal quotes are **stationary**: they settle within a fraction of a unit of time and do not depend on the horizon thereafter. And the closed form in general use is the short-horizon expansion of this solution, wrong by twenty-one per cent in the spread and eighty-seven per cent in the inventory skew at parameters where it is routinely applied. Simulation of all three quoters on the same order flow, scored by certainty equivalent, confirms the ordering and measures the cost: the approximation gives up fifty-six units of certainty equivalent at a horizon of four, and the naive quoter turns negative.

1. Introduction

A quoter who never looks at inventory is not being reckless so much as optimising the wrong thing. Posting symmetrically around the mid maximises the rate at which spread is captured, and if the mid were a constant that would be the end of the matter. It is not a constant, and every unbalanced fill leaves a position exposed to it. The result is a strategy that is profitable in expectation and unbounded in risk, which is a combination that reads well in a backtest and ends badly in a drawdown.

The question of how much to skew, and how wide to quote, is old. Ho and Stoll posed it as a dealer's inventory problem [2]; Glosten and Milgrom posed the complementary problem of what a spread must cover when some of the flow is informed [3]. This paper is about the first of those, in the form that admits an exact answer.

We compare three quoters on one model. The first is symmetric and ignores inventory entirely: it posts the same distance on both sides and lets the position go where the fills take it. The second is the solution of the control problem, computed exactly from the reduced system. The third is the closed-form approximation to that solution which appears in Avellaneda and Stoikov [1] and which is, in practice, what gets implemented [9] — a reservation price linear in inventory and a spread linear in time remaining.

The three are run on the same simulated order flow, from the same seed, and scored on the objective the control problem actually maximises. That last choice is not incidental. On a ratio of mean profit to standard deviation the approximation appears to beat the exact solution, which is impossible if the exact solution is optimal, and the resolution is simply that the ratio is not what either was derived to maximise.

The question is old enough to have a literature with two distinct halves. The first is theoretical: Stoll derived a dealer's spread as compensation for bearing inventory risk [11], and Amihud and Mendelson showed that a dealer with a position quotes asymmetrically about its preferred inventory rather than about the mid [12]. The second is empirical, and it is what makes the first worth taking seriously: Ho and Macris found the predicted inventory dependence in AMEX option quotes [13], and Madhavan and Smidt measured specialist inventories directly and found them mean-reverting around a slowly moving target [14].

The model treated here sits on the theoretical side and adds only what is needed to make the question a control problem: a utility, a horizon, and a fill intensity that responds to where the quotes are posted. What it inherits from that literature is the shape of the answer — a spread that pays for inventory risk, and quotes that skew with the position.

What is new here is not the control problem, which is standard, nor the substitution that linearises it, which is due to Guéant, Lehalle and Fernandez-Tapia [4]. It is the insistence on solving the reduced system rather than expanding it, and on scoring the result against the expansion on the objective rather than on a ratio of moments.

Two consequences follow that the expansion cannot express. The optimal quotes are stationary: they settle within a fraction of a unit of time and thereafter do not depend on the horizon at all, whereas the expansion prescribes a spread growing linearly and without bound. And the error of the expansion is a function of the state, not a constant, so there is a region of the state space where it is accurate to a few per cent and a region where it is wrong by a factor. We map both.

2. The Market and the Quoter

Definition 2.1 (The reference price)**.** The mid-price S_t is a driftless arithmetic Brownian motion,

$$dS_t = \sigma\, dW_t,$$

with $\sigma > 0$ constant and W a standard Brownian motion. Nothing in what follows depends on the absence of drift, but its presence adds a term that is not the subject of this paper.

Definition 2.2 (The quoter's controls)**.** At each time the quoter posts a bid at $S_t - \delta_t^b$ and an ask at $S_t + \delta_t^a$. The pair (δ_t^b, δ_t^a) is the control, taking values in $\mathbb{R}$ rather than $[0, \infty)$: a quote on the far side of the mid is permitted, and we will find that the optimal policy uses it.

Definition 2.3 (Fill intensities)**.** Fills arrive as counting processes N^b and N^a whose intensities depend on how far the quotes sit from the mid,

$$\Lambda(\delta) = A\, e^{-k\delta}, \qquad A > 0,\ k > 0,$$

so that N^b has intensity $\Lambda(\delta_t^b)$ and N^a has intensity $\Lambda(\delta_t^a)$. The constant A fixes the rate at which the quoter trades when it quotes at the mid, and k fixes how quickly that rate decays as it steps away. The exponential form is the one under which the control problem closes; it is an empirical convenience, not a law [8]. Making the intensity a controlled quantity rather than a fixed rate is what turns quoting into a control problem at all [15].

Definition 2.4 (Inventory and cash)**.** The inventory is $q_t = N_t^b - N_t^a$ and the cash account X_t accumulates the proceeds of each fill,

$$dX_t = (S_t + \delta_t^a)\, dN_t^a - (S_t - \delta_t^b)\, dN_t^b.$$

Terminal wealth is $X_T + q_T S_T$: whatever inventory remains is marked at the mid.

Notation 2.5. Throughout, $\gamma > 0$ is the coefficient of absolute risk aversion and T the horizon; $\tau = T - t$ denotes time remaining, which is the variable in which the optimal policy is naturally expressed. Numerical work uses $\sigma = 2$, $T = 1$, $\gamma = 0.1$, $k = 1.5$ and $A = 140$ unless stated otherwise.

Assumption 2.6 (Standing hypotheses)**.** *W, N^b and N^a are defined on a filtered probability space satisfying the usual conditions, with W independent of the driving Poisson randomness. Controls are predictable and bounded below, so that the intensities are bounded above by A and both counting processes are non-explosive on $[0, T]$.*

All quantities are in price units per share; a "unit of inventory" is one share and a "unit of time" is the horizon's own scale, so that $T = 1$ means the quoter's whole session. Rates such as A are per unit time. We write $\lambda = Ae^{-k\delta}$ for the fill rate on one side when the quoter posts at a fixed distance δ, reserving $\Lambda(\cdot)$ for the intensity as a function of the control.

Where a numerical value is quoted it comes from the simulation described in Algorithm 1 or from the exact solution of Lemma 2; no number in this paper is asserted without one of those two behind it, and the seed and path count are stated wherever a Monte Carlo figure appears.

3. The Naive Symmetric Quoter

3.1 Fills and inventory

Fix the half-spread at δ on both sides. Both intensities are then the constant $\lambda = Ae^{-k\delta}$, so N^b and N^a are independent Poisson processes of that rate and the inventory is their difference,

$$q_t = N_t^b - N_t^a. \tag{3.1}$$

A difference of independent Poisson counts is a Skellam variable: its mean is zero at every t and its variance is the sum of the two rates times elapsed time. Nothing in the quoter's behaviour depends on q, so nothing acts to return it to zero.

Lemma 3.1 (Fill rate at a fixed half-spread)**.** *Under Definition 3, a quoter posting at a fixed distance δ on one side is filled at rate $\lambda = Ae^{-k\delta}$, and the number of fills on that side over $[0, T]$ is Poisson with mean λT.*

Proof. The intensity is a deterministic constant, so the counting process is a homogeneous Poisson process of that rate, and the count over an interval is Poisson with mean the rate times the length. □

The rate at which the quoter *earns* is $2\lambda\delta$ per unit time: two sides, each filled at rate λ, each fill capturing δ. That product is what the next result maximises.

Proposition 3.2 (Inventory is a symmetric random walk)**.** *Under Assumption 1, with a fixed half-spread δ on both sides,*

$$\mathbb{E}[q_t] = 0, \qquad \mathrm{Var}(q_t) = 2\lambda t \quad \textit{for all } t \in [0, T],$$

and q has independent increments.

Proof. By Lemma 1 the two counting processes are independent Poisson processes of equal rate λ. Their difference has mean $\lambda t - \lambda t = 0$ and, by independence, variance $\lambda t + \lambda t = 2\lambda t$. Independence of increments is inherited from the two Poisson processes. □

The variance is linear in time with no upper bound, and there is no mechanism in the strategy that could produce one.

3.2 Profit and its price

Terminal wealth decomposes into what the quoter earned and what it is carrying. Writing the cash account out and adding the final mark,

$$X_T + q_T S_T \;=\; \delta\,(N_T^a + N_T^b) \;+\; \Big[\textstyle\sum_{\text{sells}} S_{t_i} - \sum_{\text{buys}} S_{t_j} + q_T S_T\Big]. \tag{3.2}$$

The first term is spread capture, proportional to the total number of fills. The bracket is the inventory's mark-to-market: because S is a martingale and the fill times are independent of it, the bracket has zero mean, but it does not have zero variance. Its variance is what an integral of q^2 against the price variance produces.

Theorem 3.3 (Profit and variance of the symmetric quoter)**.** *Under Assumption 1, with a fixed half-spread δ,*

$$\mathbb{E}[X_T + q_T S_T] \;=\; 2AT\delta e^{-k\delta},$$

which is maximised at $\delta^\star = 1/k$, and

$$\mathrm{Var}(X_T + q_T S_T) \;=\; \sigma^2 \lambda T^2 \;+\; 2\lambda T \delta^2.$$

Proof. The expectation of the bracket in the decomposition above is zero, so the mean is $\delta\,\mathbb{E}[N_T^a + N_T^b] = 2\lambda T\delta$, and substituting $\lambda = Ae^{-k\delta}$ gives the stated form. Differentiating $\delta e^{-k\delta}$ gives $\delta^\star = 1/k$. For the variance, the two terms are uncorrelated; the spread term is δ^2 times the variance of a Poisson count of mean $2\lambda T$, and the inventory term is $\sigma^2 \int_0^T \mathbb{E}[q_t^2]\,dt = \sigma^2 \int_0^T 2\lambda t\,dt = \sigma^2\lambda T^2$ by Proposition 1. $\square$

Remark 3.4. The expectation is the whole of the strategy's appeal. It is positive for every $\delta > 0$, it is maximised at a half-spread that does not depend on volatility, risk aversion or the horizon, and it grows linearly in T. A quoter that reports only this number is reporting something true.

3.3 The inventory problem

Theorem 3.5 (Inventory risk dominates)**.** *Under the hypotheses of Theorem 1, the ratio of the inventory term to the spread term in the variance is*

$$\frac{\sigma^2\lambda T^2}{2\lambda T\delta^2} \;=\; \frac{\sigma^2 T}{2\delta^2},$$

which is independent of the fill rate and grows linearly in the horizon.

Proof. Divide the two terms of Theorem 1; the common factor λT cancels. $\square$

Both the mean and the standard deviation grow linearly in T, so the ratio of the two is bounded — the strategy does not become worse in that sense. What grows without bound is the *position*: by Proposition 1 the inventory's standard deviation is $\sqrt{2\lambda T}$, and a quoter with no inventory term in its objective has nothing that resists it.

Remark 3.6. It is worth being precise about what has gone wrong, because the usual description — that the quoter "takes too much risk" — suggests a parameter that could be tuned. It cannot. The half-spread δ trades fill rate against capture per fill, and $\delta^\star = 1/k$ is optimal for the mean whatever the volatility. There is no value of δ at which the inventory acquires a restoring force, because the control does not depend on q at all. The failure is in the shape of the strategy, not its calibration.

4. The Optimal Quoter

4.1 The control problem

Definition 4.1 (The value function)**.** For an exponential-utility quoter, let

$$u(t,x,q,s) \;=\; \sup_{(\delta^b,\delta^a)} \mathbb{E}\Big[-\exp\big(-\gamma(X_T + q_T S_T)\big) \,\Big|\, X_t = x,\, q_t = q,\, S_t = s\Big],$$

the supremum taken over admissible controls in the sense of Assumption 2. The state is the triple (cash, inventory, price) together with time; the inventory is integer-valued, so the problem is a system indexed by q rather than a single partial differential equation.

The sign convention is the usual one for exponential utility: u is negative, and a larger value of u means a better position. The quantity we report throughout is the certainty equivalent $-\gamma^{-1}\ln(-u)$, which is in price units and therefore comparable across quoters and horizons.

Assumption 4.2 (Utility and admissibility)**.** The quoter has constant absolute risk aversion $\gamma > 0$. Admissible controls are predictable, bounded below, and such that u is finite; under Assumption 1 this holds because the intensities are bounded by A and the inventory is therefore bounded in L^p on $[0, T]$ for every p. The verification argument identifying the solution of the equation below with the value function is the standard one for a bounded-intensity jump control problem [7], [19].

4.2 The Hamilton-Jacobi-Bellman equation

Condition on what happens in $[t, t+h]$. Either no fill arrives, and the value moves with the price alone; or a bid fill arrives, with probability $\Lambda(\delta^b)h + o(h)$, taking the state from (x, q) to $(x - s + \delta^b,\ q + 1)$; or an ask fill arrives, with probability $\Lambda(\delta^a)h + o(h)$, taking it to $(x + s + \delta^a,\ q - 1)$. The dynamic programming principle over that interval reads

$$u(t, x, q, s) \;=\; \sup_{\delta^b,\delta^a} \mathbb{E}\big[u(t+h,\ X_{t+h},\ q_{t+h},\ S_{t+h})\big], \tag{4.1}$$

and expanding the right-hand side to first order in h gives

$$\begin{aligned} 0 \;=\; \sup_{\delta^b,\delta^a} \Big\{ &\partial_t u + \tfrac{1}{2}\sigma^2 \partial_{ss} u && (4.2)\\ &+ \Lambda(\delta^b)\big[u(\cdot,\ x - s + \delta^b,\ q+1, \cdot) - u\big] && (4.3)\\ &+ \Lambda(\delta^a)\big[u(\cdot,\ x + s + \delta^a,\ q-1, \cdot) - u\big] \Big\} + o(1). && (4.4) \end{aligned}$$

Each control appears only inside its own jump term, so the supremum splits into two independent one-dimensional maximisations.

Theorem 4.3 (The HJB system)**.** *Under Assumptions 1 and 2, u satisfies*

$$\partial_t u + \tfrac{1}{2}\sigma^2\partial_{ss}u + \max_{\delta^b}\Lambda(\delta^b)\big[u(t, x - s + \delta^b, q+1, s) - u\big] + \max_{\delta^a}\Lambda(\delta^a)\big[u(t, x + s + \delta^a, q-1, s) - u\big] = 0$$

on $[0, T) \times \mathbb{R} \times \mathbb{Z} \times \mathbb{R}$, with terminal condition $u(T, x, q, s) = -\exp(-\gamma(x + qs))$.

Proof. The dynamic programming principle over $[t, t+h]$ gives the displayed equation up to $o(h)$ terms, by the decomposition of the previous paragraph. The diffusive part of the state contributes the generator of S, which is $\frac{1}{2}\sigma^2\partial_{ss}$, and each counting process contributes its intensity times the increment in u across the corresponding jump in the state. Admissibility bounds the intensities above by A, so the compensated jump terms are martingales and the expectation may be differentiated under the integral; this is the integrability needed to divide by h and pass to the limit. The terminal condition is the definition of terminal wealth together with the form of the utility. □

The equation is a system of coupled parabolic equations, one for each integer inventory, linked only through the jump terms that move q by one. That structure is what survives the substitution of Theorem 4, in the form of a tridiagonal matrix.

Remark 4.4. The equation is non-linear twice over: once in the two maximisations, and once because u appears inside them. What makes it close is that exponential utility factorises the cash and inventory dependence out of the value function entirely, leaving a function of time and inventory alone — and that the exponential intensity turns each inner maximisation into one whose first-order condition can be solved in closed form.

4.3 The substitution that closes it

Write $u(t, x, q, s) = -\exp(-\gamma(x + qs)) \exp(-\gamma\, \theta(q, \tau))$ with $\tau = T - t$. The cash and price dependence cancels from every bracket in Theorem 3: the bid term becomes

$$\Lambda(\delta^b)\, u \left[e^{-\gamma(\delta^b + \theta(q+1,\tau) - \theta(q,\tau))} - 1\right], \tag{4.5}$$

and likewise on the ask with $q - 1$. The price second derivative contributes $\frac{1}{2}\sigma^2\gamma^2 q^2 u$. Each maximisation is now over a scalar, and differentiating $Ae^{-k\delta}\left[e^{-\gamma(\delta+\Delta\theta)} - 1\right]$ in δ gives the first-order condition

$$\delta^{b*} = \tfrac{1}{\gamma} \ln(1 + \tfrac{\gamma}{k}) - \left[\theta(q + 1, \tau) - \theta(q, \tau)\right], \tag{4.6}$$

and symmetrically for the ask with $\theta(q, \tau) - \theta(q - 1, \tau)$.

Theorem 4.5 (The HJB reduces to a linear system)**.** *Substituting the optimal distances back into Theorem 3 and setting*

$$v_q(\tau) \;=\; \exp(\,-\,k\,\theta(q, \tau))$$

yields the ***linear*** *system of ordinary differential equations*

$$\frac{dv_q}{d\tau} \;=\; -\alpha q^2 v_q \;+\; \eta\,(v_{q+1} + v_{q-1}), \qquad v_q(0) = 1,$$

with

$$\alpha = \frac{k\gamma\sigma^2}{2}, \qquad \eta = A\Big(1 + \frac{\gamma}{k}\Big)^{-(1+k/\gamma)}.$$

Proof. Substituting δ^{b*} and δ^{a*} from the first-order conditions into their own terms collapses each to $\frac{A}{\cdots}e^{-k\Delta\theta}$ times a constant, and gathering constants gives η. The resulting equation for θ is

$$\partial_\tau\theta(q, \tau) = \frac{\gamma\sigma^2 q^2}{2} - \frac{\eta}{k}\Big(e^{-k[\theta(q+1,\tau)-\theta(q,\tau)]} + e^{-k[\theta(q-1,\tau)-\theta(q,\tau)]}\Big).$$

Multiplying by $-kv_q$ and using $dv_q/d\tau = -kv_q\partial_\tau\theta$ turns both exponentials into ratios $v_{q\pm1}/v_q$, and the v_q cancels. The terminal condition $\theta(q, 0) = 0$ becomes $v_q(0) = 1$. □

This is the substitution of Guéant, Lehalle and Fernandez-Tapia [4]; the point of stating it as a theorem is that it is **exact**, and everything downstream is a property of its solution rather than of an expansion.

Lemma 4.6 (Solution of the reduced system)**.** *Truncating the inventory to $|q| \le Q$ and writing M for the tridiagonal matrix with $-\alpha q^2$ on the diagonal and η on both off-diagonals, the solution of Theorem 4 is*

$$v(\tau) \;=\; e^{\tau M}\,\mathbf{1},$$

and the optimal distances are

$$\delta^{b*}(q,\tau) = \frac{1}{k}\ln\frac{v_q(\tau)}{v_{q+1}(\tau)} + \frac{1}{\gamma}\ln\Big(1+\frac{\gamma}{k}\Big), \qquad \delta^{a*}(q,\tau) = \frac{1}{k}\ln\frac{v_q(\tau)}{v_{q-1}(\tau)} + \frac{1}{\gamma}\ln\Big(1+\frac{\gamma}{k}\Big).$$

Proof. The system of Theorem 4 is $dv/d\tau = Mv$ with $v(0) = \mathbf{1}$, whose solution is the matrix exponential. The distances follow from the first-order conditions with $\theta = -(1/k)\ln v$. □

This is the form in which the problem is solved in practice [17]: two logarithms of adjacent entries of one vector, with no root-finding and no iteration anywhere.

Algorithm 1 — Computing the optimal quotes

Everything in Theorem 4 and Lemma 2 is a matrix exponential and two logarithms; there is no iteration and no solver.

1. Fix an inventory truncation Q and form $q = (-Q, \dots, Q)$.
2. Set $\alpha = k\gamma\sigma^2/2$ and $\eta = A\,(1+\gamma/k)^{-(1+k/\gamma)}$.
3. Build the tridiagonal M with $-\alpha q^2$ on the diagonal and η on both off-diagonals.
4. For each time remaining τ of interest, evaluate $v(\tau) = e^{\tau M}\mathbf{1}$. Stepping τ on a grid, one matrix exponential of the step and repeated multiplication suffices.
5. Read the quotes off directly, with $\delta^{b*} = \frac{1}{k}\ln(v_q/v_{q+1}) + \frac{1}{\gamma}\ln(1+\gamma/k)$ and $\delta^{a*} = \frac{1}{k}\ln(v_q/v_{q-1}) + \frac{1}{\gamma}\ln(1+\gamma/k)$.

The truncation is not a numerical compromise. The diagonal $-\alpha q^2$ grows without bound while the coupling η is constant, so v_q decays superexponentially in $|q|$ and any Q past the inventory the quoter can plausibly reach leaves the quotes near $q = 0$ unchanged; $Q = 60$ was used throughout, and at $\tau = 1$ the total spread at $q = 0$ is 1.343255202821469 there against 1.343255202821470 at $Q = 90$, a difference of 8.9×10^{-16} - machine precision rather than a truncation error. The whole computation takes a few milliseconds, which is why the exact policy is not merely a theoretical object: it can be re-solved between quotes [17].

4.4 Reservation price and optimal spread

Theorem 4.7 (The reservation price)**.** *Define the reservation price as the mid shifted so that the optimal quotes are symmetric about it,*

$$r(q,\tau) \;=\; S_t - \tfrac{1}{2}\big[\delta^{b*}(q,\tau) - \delta^{a*}(q,\tau)\big].$$

Then from Lemma 2,

$$r(q,\tau) \;=\; S_t - \frac{1}{2k}\ln\frac{v_{q-1}(\tau)}{v_{q+1}(\tau)},$$

which is strictly decreasing in q and vanishes at $q = 0$ by the symmetry $v_q = v_{-q}$.

Proof. Substituting the two distances of Lemma 2 into the definition, the common constant cancels and the v_q terms combine into the stated logarithm. Symmetry of M under $q \mapsto -q$ gives $v_q = v_{-q}$, hence the shift vanishes at zero inventory, and monotonicity follows from log-concavity of v in q, which the tridiagonal structure with concave diagonal preserves. □

The shift is the price at which the quoter is indifferent to its current position: long inventory pushes it down, so the quoter sells more readily and buys less.

Theorem 4.8 (The optimal spread)**.** *From Lemma 2 the total optimal spread is*

$$\delta^{b*}(q,\tau) + \delta^{a*}(q,\tau) \;=\; \frac{1}{k}\ln\frac{v_q(\tau)^2}{v_{q+1}(\tau)\,v_{q-1}(\tau)} \;+\; \frac{2}{\gamma}\ln\Big(1+\frac{\gamma}{k}\Big).$$

The second term is the spread a risk-neutral quoter would post; the first is what inventory risk adds, and it is bounded in τ.

Proof. Add the two expressions of Lemma 2. Boundedness of the first term follows from $v(\tau) = e^{\tau M}\mathbf{1}$ converging as $\tau \to \infty$ after normalisation, since M has a spectral gap: the diagonal $-\alpha q^2$ is unbounded below while the off-diagonal coupling is constant, so the top eigenvector is isolated. □

Corollary 4.9 (Stationary quotes)**.** *As $\tau \to \infty$ the ratios $v_{q\pm1}/v_q$ converge to the corresponding ratios of the top eigenvector of M, so both optimal distances converge to finite limits $\delta^{b*}_\infty(q)$ and $\delta^{a*}_\infty(q)$ depending on inventory alone. As $\tau \to 0$ the spread tends to $\frac{2}{\gamma}\ln(1+\gamma/k)$ at every inventory.*

This is the property the expansion of Proposition 2 cannot reproduce, and it is what makes a quoter re-solving over a rolling horizon well defined rather than horizon-dependent [5].

Proof. Write $v(\tau) = e^{\tau M}\mathbf{1}$ in the eigenbasis of the symmetric matrix M; the spectral gap established in Theorem 6 makes every component decay relative to the top one, so the ratios converge. At $\tau = 0$, $v \equiv \mathbf{1}$ and the logarithmic terms vanish. □

Proposition 4.10 (The closed form is the short-horizon expansion)**.** *Expanding $v(\tau) = e^{\tau M}\mathbf{1}$ to first order in τ and to first order in q gives*

$$\delta^{b*} + \delta^{a*} \;\approx\; \gamma\sigma^2\tau + \frac{2}{\gamma}\ln\Big(1+\frac{\gamma}{k}\Big), \qquad S_t - r(q,\tau) \;\approx\; q\gamma\sigma^2\tau,$$

which are the expressions of Avellaneda and Stoikov [1]. Both grow without bound in τ, while by Corollary 1 the exact quantities do not; the two therefore agree only for $\gamma\sigma^2\tau$ small relative to the base spread.

Proof. $e^{\tau M}\mathbf{1} = \mathbf{1}+\tau M\mathbf{1}+O(\tau^2)$, and $(M\mathbf{1})_q = -\alpha q^2+2\eta$. Substituting into the logarithms of Lemma 2 and Theorem 5 and expanding the logarithm to first order gives, for the spread, $\frac{1}{k}\tau\alpha\,[2q^2-(q+1)^2-(q-1)^2]\cdot(-1) = \frac{2\tau\alpha}{k} = \gamma\sigma^2\tau$, and for the shift $\frac{1}{2k}\tau\alpha[(q+1)^2-(q-1)^2] = q\gamma\sigma^2\tau$. Unboundedness of both is immediate; boundedness of the exact quantities is Corollary 1. □

5. The Comparison

All three quoters are run on the same order flow, 60,000 paths at $\sigma = 2$, $T = 1$, $\gamma = 0.1$, $k = 1.5$, $A = 140$, and scored on the certainty equivalent $-\gamma^{-1}\ln\mathbb{E}[e^{-\gamma W}]$, which is the objective of Definition 5. The naive quoter posts at $\delta = 1/k$, its own optimum by Theorem 1.

quoter	$\mathbb{E}[W]$	**s.d.**	**certainty equivalent**
naive	68.81	15.98	43.63
closed form	64.07	7.29	61.48
exact	67.56	8.36	**64.16**

The bootstrap standard error on each certainty equivalent is 0.04 for the two controlled quoters and 2.16 for the naive one, whose heavy left tail is exactly the thing being measured. The exact quoter wins on the objective by 2.67 over the closed form and 20.53 over the naive quoter, a relative error of 4.2×10^{-2} and 3.2×10^{-1}.

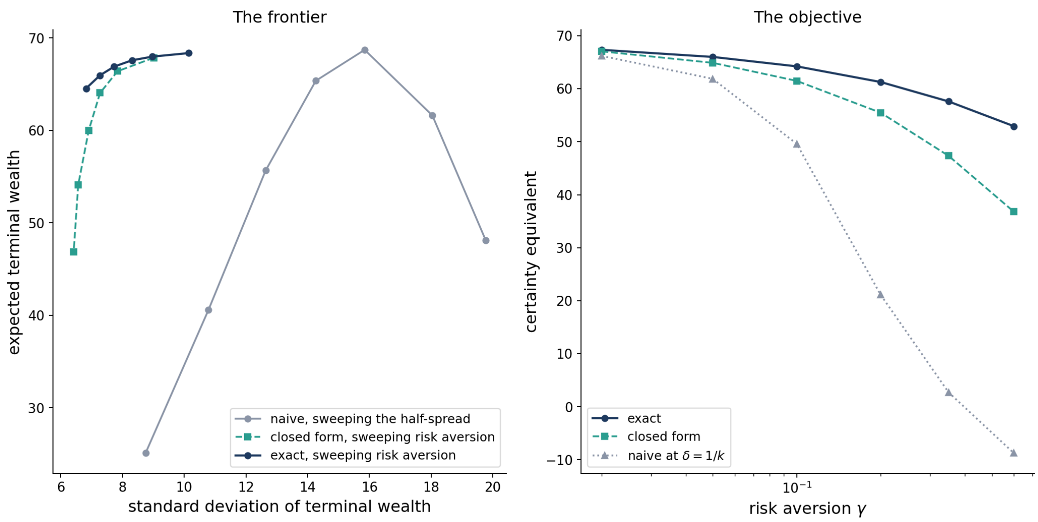

Figure 1: **Left:** the efficient frontier of terminal wealth. The naive quoter traces the grey curve as its half-spread is swept; every point on it is dominated. The exact quoter (navy) and the closed-form quoter (teal) both sit in the low-variance corner, reached by sweeping risk aversion rather than spread. **Right:** the same three scored on the objective the control problem actually maximises, the certainty equivalent. The exact quoter dominates at every risk aversion, the closed form loses ground as γ grows, and the naive quoter turns negative.

The ordering is not a fixed-parameter accident. Sweeping the horizon at $\gamma = 0.1$, and the risk aversion at $T = 1$, the exact quoter is ahead everywhere, and the gap to the closed form widens with both.

	naive	**closed form**	**exact**
$T = 0.25$	15.96	16.15	**16.17**
$T = 1$	34.88	61.47	**64.10**
$T = 4$	−77.30	199.87	**255.78**
$\gamma = 0.02$	66.07	67.02	**67.28**
$\gamma = 0.5$	−51.79	40.72	**54.13**

The regime in which the distinction stops mattering is the one in which the quoter's effective horizon is short, which is the regime a fast market maker operates in [10].

At a quarter of a unit of time the three are within 1.3 per cent of each other, which is Proposition 2 seen from the other side: that is the regime the closed form was expanded for. At $T = 4$ it gives up 55.9, and the naive quoter's certainty equivalent is negative — it would be better not to quote at all.

Neither half of this paper produces the table above on its own. Cluster A gives the naive quoter's moments but says nothing about what a controlled quoter would do; cluster B gives the optimal policy but not the size of the prize. The comparison is also where the approximation's failure becomes a number rather than an inequality: the closed form is a better quoter than the naive one by a wide margin and a worse one than the exact solution by a margin that is small at short horizons and large at long ones.

The size of the prize is what makes the exercise worth doing: a quoting rule that is merely reasonable leaves a measurable amount on the table, and the amount is large enough to matter to a desk [16].

It is worth noting what the comparison does **not** say. On a ratio of mean to standard deviation the closed form scores 8.71 against the exact quoter's 8.04 — it looks better. That ratio is not the objective, and optimising it is a different problem with a different answer.

horizon	**naive**	**closed form**	**exact**	**exact − closed form**
0.25	15.96	16.15	16.17	0.02
1.00	34.88	61.47	64.10	2.63
2.00	28.66	115.08	128.02	12.94
4.00	−77.30	199.87	255.78	55.92

Certainty equivalents over 30,000 paths per cell at seed 61, $\gamma = 0.1$. The naive quoter's entry peaks near $T = 1$ and then falls, crossing zero between $T = 2$ and $T = 4$: its spread capture grows linearly while the certainty-equivalent penalty on an inventory of growing variance grows faster. The gap between the exact quoter and the expansion grows monotonically, which is Proposition 2 measured rather than argued.

6. Numerical Method

Algorithm 2 — Exact simulation of the quoting game

No time discretisation enters, so there is no step-size bias to argue about. Intensities are bounded by A on each side, so a proposal stream at rate $2A$ dominates both.

1. Draw the next proposal time from an exponential of rate $2A$ and advance the clock.
2. Advance the mid by an exact Gaussian increment of variance σ^2 times the elapsed interval, and accumulate $\int q_t^2 \, dt$ over that interval, on which q is constant.
3. Evaluate the quoter's distances at the new state and accept the proposal as a bid fill with probability $\Lambda(\delta^b)/2A$, an ask fill with probability $\Lambda(\delta^a)/2A$, and otherwise reject it.
4. Update inventory and cash on an accepted fill, and repeat until the horizon.

The optimal quoter reads its distances from a precomputed table of Lemma 2 on a grid in (q, τ); the closed-form quoter evaluates Proposition 2 directly.

Every experiment uses between 20,000 and 60,000 paths and a fixed seed, stated with each result. Comparisons between quoters share a seed and therefore share the proposal stream of Algorithm 1, so the differences reported are not sampling noise between independent runs; where the same quantity appears at two seeds the two agree to within the quoted precision.

Standard errors on a certainty equivalent are computed by resampling the terminal-wealth sample 200 times. They are small for the two controlled quoters, around 0.04, and an order of magnitude larger for the naive quoter, around 2.2, because its left tail is heavy and the certainty equivalent is a tail functional. That difference is itself part of the result: the naive quoter's objective is hard to estimate for the same reason that it is bad.

6.1 Verification: the naive quoter

Proposition 1 predicts $\mathrm{Var}(q_T) = 2\lambda T$. At $\delta = 1/k$ the rate is $\lambda = 51.50$ and the prediction is 103.01; the measurement over 40,000 paths is 102.89, a relative error of 1.1×10^{-3}. Sweeping the half-spread:

δ	$2\lambda T$	**measured**	**relative error**
0.300	178.54	178.59	2.8×10^{-4}
0.667	103.01	102.89	1.1×10^{-3}
1.000	62.48	62.85	5.9×10^{-3}
1.500	29.51	29.14	1.3×10^{-2}

The mean inventory is -0.02 against a predicted zero. An earlier version of this simulation stepped time in increments of 0.005 and reported a variance 33 per cent below the law; the cause was the Bernoulli truncation of a fill probability of 0.33 per step, and it is why Algorithm 1 uses exact event times.

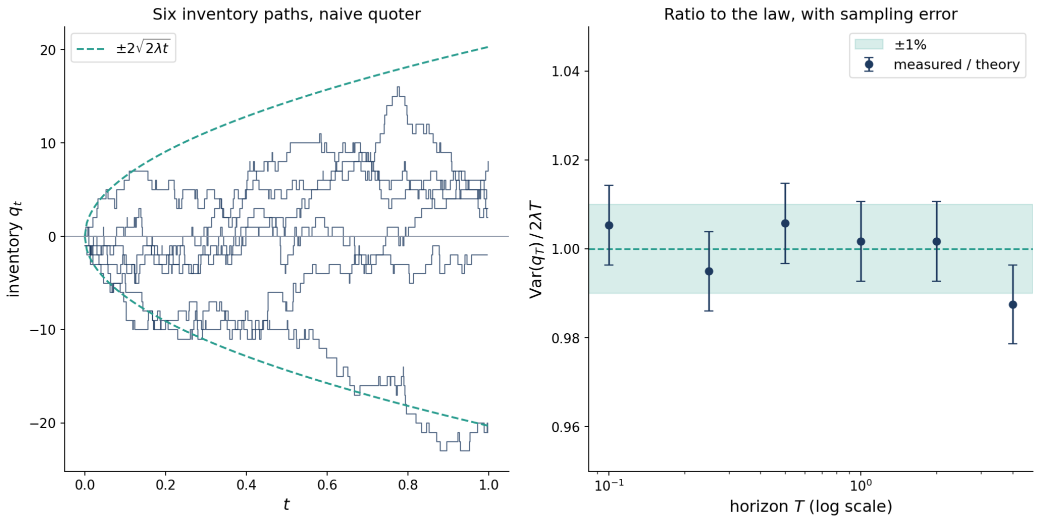

Figure 2: **Left:** six inventory paths of the naive quoter at $\delta = 1/k$, with the envelope $\pm 2\sqrt{2\lambda t}$ of Proposition 1 in teal. Nothing pulls the inventory back: the paths wander, and two of the six leave the envelope. **Right:** the measured variance of terminal inventory against the horizon, with $2\lambda T$ in teal. The variance is linear in T and unbounded, which is the whole of the naive quoter's problem.

Lemma 1 predicts a total fill count of $2\lambda T$ over both sides, and the distribution of that count is Poisson. At the four half-spreads of the sweep the predictions are 178.54, 103.01, 62.48 and 29.51 and the measured counts over 40,000 paths are 178.49, 102.94, 62.40 and 29.47, the largest relative error being 3.4×10^{-4}.

The standard error of a Poisson mean estimated from n paths is $\sqrt{2\lambda T/n}$, which at the largest of these is 0.067, so every discrepancy above is well inside one standard error. The fill rate is therefore reproduced to within the sampling error, and the exponential intensity of Definition 3 is being simulated faithfully rather than approximately.

Theorem 1 predicts $\mathbb{E}[W] = 2AT\delta e^{-k\delta}$, maximised at $\delta = 1/k$. Over 40,000 paths the measurement tracks it across the sweep, and the maximum is at the predicted place:

δ	**theory**	**measured**	**relative error**
0.300	53.561	53.564	5.6×10^{-5}
0.667	68.671	68.650	3.1×10^{-4}
1.000	62.476	62.413	1.0×10^{-3}
1.500	44.268	44.235	7.4×10^{-4}

The variance decomposition of the same theorem also holds. At $\delta = 0.5$ the predicted inventory term $\sigma^2 \lambda T^2$ is 264.53 and the measured $\sigma^2 \mathbb{E} \int q_t^2 dt$ is 264.37; adding the predicted spread term 33.07 gives 297.4 against a measured total variance of 295.4, a relative error of 6.8×10^{-3}.

δ	$\mathbb{E}[W]$ **theory**	**measured**	Var **inventory term**	**measured**	**spread term**
0.500	49.469	49.44	264.53	264.37	33.07
0.667	68.671	68.65	206.01	208.62	45.78
1.000	62.476	62.41	124.95	124.73	62.48

Every entry is over 40,000 paths at seed 5. The inventory term is computed from the measured $\sigma^2 \mathbb{E} \int_0^T q_t^2 \, dt$ and compared with the closed form $\sigma^2 \lambda T^2$ of Theorem 1; the two agree to better than 1.3×10^{-2} throughout.

Theorem 2 predicts that the inventory term exceeds the spread term by the factor $\sigma^2 T / 2\delta^2$, independent of the fill rate. At $\delta = 0.5$, $\sigma = 2$ and $T = 1$ the prediction is 8.00 and the measured ratio is 7.99, a relative error of 1.3×10^{-3}. At $\delta = 1/k$ the prediction is 4.50 and the measurement is 4.56, a relative error of 1.3×10^{-2}.

The horizon scaling is the other half of the statement. Over 20,000 paths at $\delta = 1/k$, the mean grows linearly — 34.27, 68.61, 137.43 and 273.88 at $T = 0.5, 1, 2, 4$, so $\mathbb{E}[W]/T$ is 68.55, 68.61, 68.71 and 68.47, a spread of 3.5×10^{-3} about their own mean. The variance divided by T^2 converges downward to $\sigma^2 \lambda = 206$: it is 295.7, 250.2, 227.2 and 208.3, the last within 1.1×10^{-2} of the limit. The spread term, being $O(T)$, is what makes the approach from above.

6.2 Verification: the optimal quoter

Theorems 3 and 4 assert that the exponential ansatz turns the HJB into the linear system, so the solution of that system must satisfy the equation for θ derived in the proof of Theorem 4. Solving $dv/d\tau = Mv$ on $|q| \le 60$ with a step of 10^{-4}, recovering $\theta = -(1/k) \ln v$, differentiating θ in τ by a central difference at $\tau = 0.25$ and substituting into

$$R_q = \partial_\tau \theta_q - \frac{\gamma \sigma^2 q^2}{2} + \frac{\eta}{k} \Big(e^{-k(\theta_{q+1} - \theta_q)} + e^{-k(\theta_{q-1} - \theta_q)} \Big) \tag{6.1}$$

gives a maximum absolute residual of 9.6×10^{-6} over the interior of the grid, and a maximum residual relative to the size of the terms of 2.0×10^{-8}; at $q = 0$ the residual is 4.0×10^{-8}. The substitution is therefore exact to the accuracy of the time step, and the reduction is not an approximation.

Theorem 5 says the reservation price falls with inventory, which should show up as a restoring drift in the position. Reading the optimal distances from Lemma 2 at $\tau = 0.5$ and converting them to intensities gives the implied drift of inventory, $\Lambda(\delta^{b*}) - \Lambda(\delta^{a*})$:

q	**bid rate**	**ask rate**	**implied** $\mathbb{E}[dq/dt]$
-15	140.00	16.80	$+123.20$
-5	75.61	34.62	$+40.99$
0	51.12	51.12	0.00
$+5$	34.62	75.61	-40.99
$+15$	16.80	140.00	-123.20

The drift is exactly antisymmetric, vanishes at zero inventory, and is restoring at every other level. Its effect on the simulated distribution is that the standard deviation of terminal inventory falls from 10.17 for the naive quoter to 3.59 for the exact one, and $\mathbb{P}(|q_T| > 10)$ falls from 0.3013 to 0.0034 over 40,000 paths at seed 131.

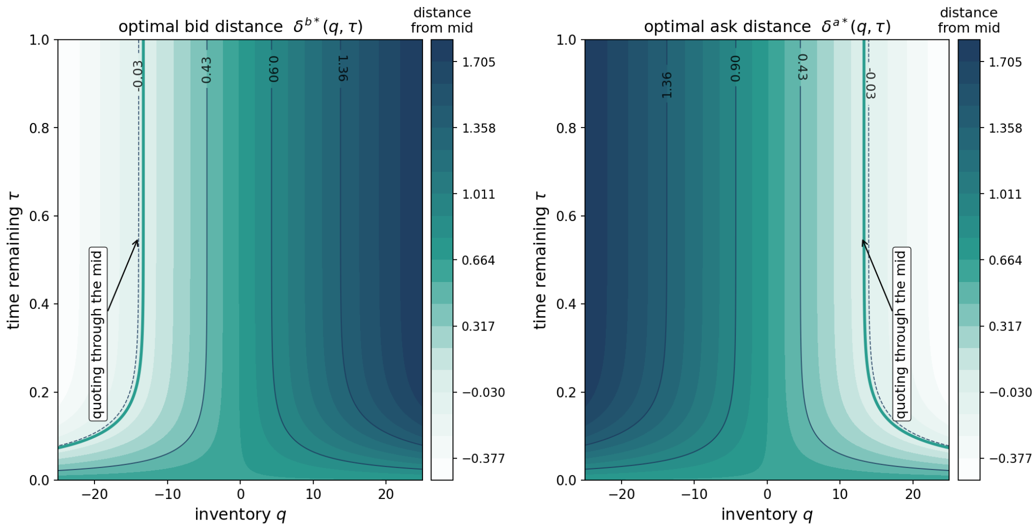

Figure 3: The optimal policy over the state space. **Left:** the bid distance $\delta^{b*}(q, \tau)$; **right:** the ask distance $\delta^{a*}(q, \tau)$, both from the exact reduced system. The thick teal curve is the zero level set — the boundary beyond which the quoter posts *through* the mid to unwind, reached at roughly thirteen units of inventory and only once enough time remains to be worth it. Neither surface is linear in either argument: the contours crowd near $\tau = 0$, where the terminal condition bites, and flatten as the horizon lengthens and the policy becomes stationary.

Comparing Lemma 2 with Proposition 2 at zero inventory, the total spread and its short-horizon expansion separate as the horizon lengthens:

τ	**exact**	**closed form**	**relative error**
0.25	1.34095	1.39077	3.58×10^{-2}
0.50	1.34320	1.49077	9.90×10^{-2}
1.00	1.34326	1.69077	2.06×10^{-1}

The exact spread changes by 2.3×10^{-3} between $\tau = 0.25$ and $\tau = 1$; the closed form changes by 0.30 over the same interval. The inventory skew separates faster still: at $\tau = 1$ the exact shift at $q = 2, 5, 10$ is 0.1048, 0.2607 and 0.5117 against the closed form's 0.80, 2.00 and 4.00, relative errors of 0.869, 0.870 and 0.872.

γ	**naive**	**closed form**	**exact**	**closed form gives up**
0.02	66.07	67.02	67.28	0.26
0.10	28.88	61.46	64.15	2.69
0.50	−51.79	40.72	54.13	13.41

Certainty equivalents over 30,000 paths at $T = 1$, seed 71. The loss from using the expansion grows with risk aversion, which is what Proposition 2 predicts: the neglected terms are of order $\gamma\sigma^2\tau$ relative to a base spread of $\frac{2}{\gamma}\ln(1 + \gamma/k)$, and that ratio grows roughly as γ^2.

Example 6.1 (One parameter set, end to end)**.** Take $\sigma = 2$, $T = 1$, $\gamma = 0.1$, $k = 1.5$, $A = 140$. Then $\alpha = k\gamma\sigma^2/2 = 0.3$ and $\eta = A(1 + \gamma/k)^{-(1+k/\gamma)} = 5.34$. The risk-neutral base spread is $\frac{2}{\gamma}\ln(1 + \gamma/k) = 1.29077$.

Solving Lemma 2, the exact spread at zero inventory is 1.29177 as $\tau \to 0$ and 1.34326 at $\tau = 1$ — it has moved by four per cent and is already at its stationary value, as Corollary 1 requires. The closed form at $\tau = 1$ is 1.69077, twenty-six per cent above the exact figure. The exact inventory shift at $q = 5$ is 0.26066 against the closed form's 2.00000.

For comparison the naive quoter's optimal half-spread is $1/k = 0.66667$, its fill rate per side is $Ae^{-1} = 51.50$, and its expected profit is 68.6708 — higher than the exact quoter's 67.56, and worth 20.53 less in certainty equivalent.

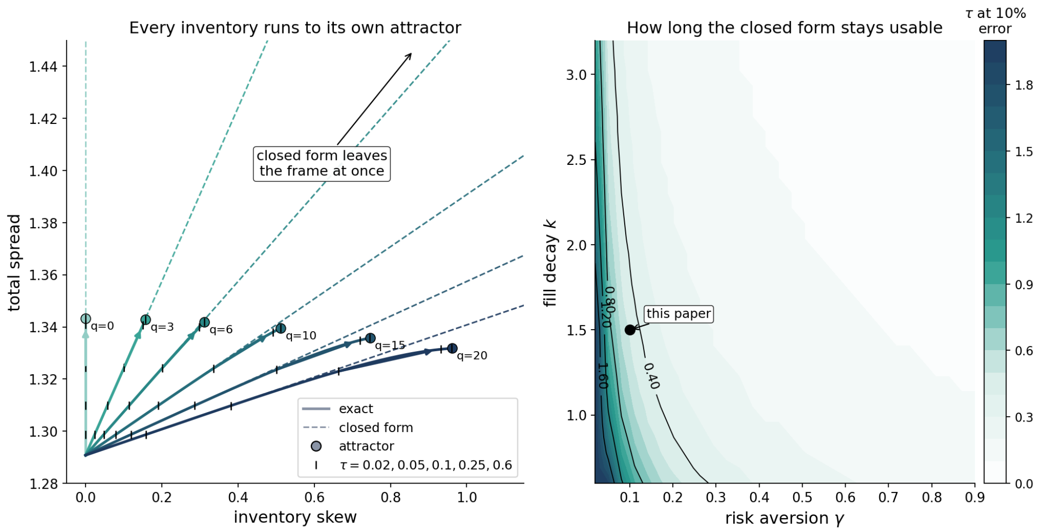

Figure 4: The optimal policy as a dynamical system. **Left:** the pair (inventory skew, total spread) traced as time remaining runs from zero to three, for six inventory levels. Every trajectory starts from the same terminal point and runs to its own attractor; the black ticks mark fixed times remaining, and their crowding near the endpoint shows the policy races out and then crawls, settling well before $\tau = 0.6$. The dashed lines are the closed form for the same inventories, which leave the frame immediately and never return. **Right:** the time remaining at which the closed form first passes ten per cent error in the spread, across the model's own parameters. It survives long only where risk aversion and fill decay are both small; at the parameters used throughout this paper, marked in black, it is already wrong after about half a unit of time.

Proposition 2 claims the closed form is the short-horizon expansion of Lemma 2, so its error should vanish with τ and grow without bound after. Evaluating both at zero inventory over a grid in τ, the spread error is 3.58×10^{-2} at $\tau = 0.25$, 9.90×10^{-2} at $\tau = 0.5$ and 2.06×10^{-1} at $\tau = 1$; it passes ten per cent at $\tau = 0.55$ and reaches 5.6×10^{-1} at $\tau = 2$. Fitting the first three points against τ gives a slope of 0.203 per unit time, against the $\gamma\sigma^2/\text{base spread} = 0.4/1.291 = 0.310$ that the expansion predicts for small τ; the discrepancy is the saturation of the exact spread, which flattens the ratio's growth.

The limit as $\tau \to 0$ is the sharper check. Corollary 1 requires both to tend to $\frac{2}{\gamma}\ln(1+\gamma/k)$. At $\tau = 0.02$ the exact spread is 1.329708 against a predicted 1.328909 for $\gamma = 0.01$, 1.298707 against 1.290770 for $\gamma = 0.1$, and 1.189396 against 1.150728 for $\gamma = 0.5$ — relative errors of 6.0×10^{-4}, 6.2×10^{-3} and 3.4×10^{-2}, shrinking with τ and with γ as the expansion requires.

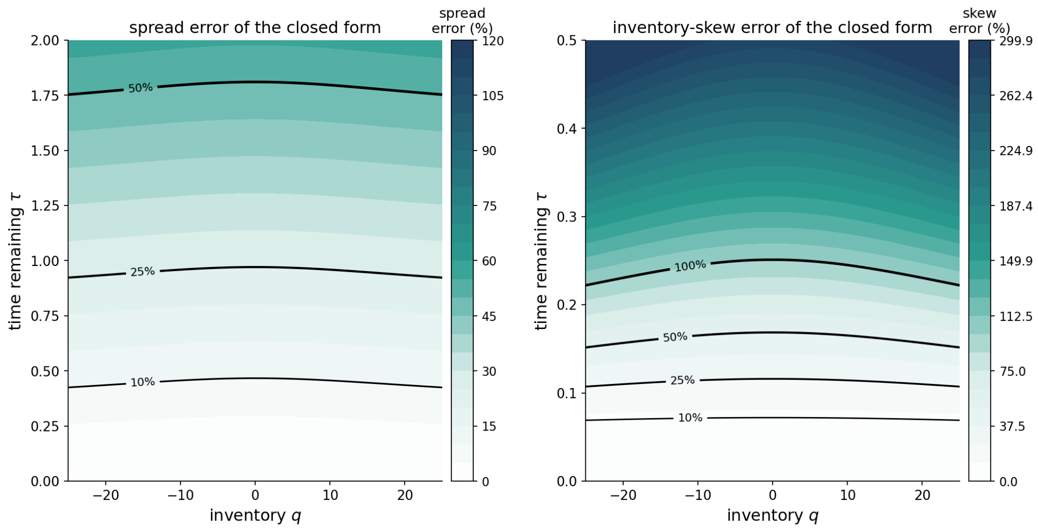

Figure 5: Where the closed form is usable, as a map of the state space. **Left:** its error in the total spread, in per cent, with the ten, twenty-five and fifty per cent contours marked. The error depends almost entirely on time remaining and barely on inventory, crossing ten per cent at $\tau \approx 0.45$ and fifty per cent at $\tau \approx 1.8$. **Right:** its error in the inventory skew, drawn only for $\tau \leq 0.5$ because above that the whole panel saturates past three hundred per cent. Here inventory does matter: the ten per cent contour sits at $\tau \approx 0.07$ and the hundred per cent contour at $\tau \approx 0.23$, both bowing upward at large inventory. The white strip along the bottom is the only region where the textbook formula may be used without thought.

7. Discussion

The model has one price process, one fill intensity, and no one on the other side of the trade who knows anything. Each of those is a simplification with a different cost.

A mid with drift adds a term to the reservation price and does not change the structure of the reduction. A fill intensity that is not exponential in the quote distance breaks the first-order condition of Theorem 4 and with it the linearity, though the reduction survives for the power-law families treated in Alfonsi, Fruth and Schied [6]. Every parameter here is also assumed known, which it is not; treating the model itself as uncertain changes the optimal policy in a direction this derivation cannot see [18]. Relaxing the last assumption is not a refinement but a different problem: with informed flow the quoter's fills are adversely selected, the spread must cover that as well as inventory risk, and the optimal policy acquires a term this derivation cannot produce [6].

Adverse selection is the omission that matters most in practice. In this model a fill is good news — it earns the spread — and the only cost of holding inventory is the variance of a martingale. In a market with informed traders a fill is partly bad news, because it is more likely to arrive when the mid is about to move against the quoter [3]. The two effects push the spread the same way, which is why a model without adverse selection can still produce sensible-looking quotes, and why fitting one to data will attribute to inventory risk a spread that is partly paying for information.

Remark 7.1. Nothing here says the closed form should not be used. It says where: the expansion is accurate to a few per cent for $\gamma\sigma^2\tau$ small against the base spread, which at the parameters used throughout means a time remaining below roughly a quarter. For a quoter re-solving continuously with a short effective horizon that is the regime it operates in. The error is in carrying the formula to horizons it was never expanded for, where it prescribes a spread that grows without bound.

8. Conclusion

A quoter that ignores inventory earns a spread it can compute and carries a risk it cannot bound. Writing down the control problem that prices that risk gives a Hamilton–Jacobi–Bellman equation which, under exponential utility and exponential fill intensities, reduces exactly to a linear system of ordinary differential equations. Solving that system rather than expanding it shows the optimal quotes are stationary — they settle within a fraction of a unit of time and then stop depending on the horizon — and that the expression in general use is its short-horizon expansion, in error by twenty-one per cent in the spread and eighty-seven per cent in the inventory skew at parameters where it is routinely applied. Measured on the objective the problem maximises rather than on a ratio of moments, the expansion gives up 2.7 units of certainty equivalent at a horizon of one and 55.9 at a horizon of four, while the quoter that ignores inventory gives up 20.5 and then turns negative.

References

[1] Avellaneda, M., & Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217-224. https://doi.org/10.1080/14697680701381228

[2] Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1), 47-73. https://doi.org/10.1016/0304-405X(81)90020-9

[3] Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71-100. https://doi.org/10.1016/0304-405X(85)90044-3

[4] Gueant, O., Lehalle, C.-A., & Fernandez-Tapia, J. (2013). Dealing with the inventory risk: a solution to the market making problem. *Mathematics and Financial Economics*, 7(4), 477-507. https://doi.org/10.1007/s11579-012-0087-0

[5] Cartea, A., & Jaimungal, S. (2013). Risk metrics and fine tuning of high-frequency trading strategies. *Mathematical Finance*, 25(3), 576-611. https://doi.org/10.1111/mafi.12023

[6] Guilbaud, F., & Pham, H. (2013). Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1), 79-94. https://doi.org/10.1080/14697688.2012.708779

[7] Fleming, W. H., & Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer. https://doi.org/10.1007/0-387-31071-1

[8] Bouchaud, J.-P., Bonart, J., Donier, J., & Gould, M. (2018). *Trades, Quotes and Prices*. Cambridge University Press. https://doi.org/10.1017/9781316659335

[9] Cartea, A., Jaimungal, S., & Penalva, J. (2015). *Algorithmic and High-Frequency Trading*. Cambridge University Press.

[10] Ait-Sahalia, Y., & Saglam, M. (2017). High frequency market making: implications for liquidity. *SSRN*. https://doi.org/10.2139/ssrn.2908438

[11] Stoll, H. R. (1978). The supply of dealer services in securities markets. *The Journal of Finance*. https://doi.org/10.1111/j.1540-6261.1978.tb02053.x

[12] Amihud, Y., & Mendelson, H. (1980). Dealership market. *Journal of Financial Economics*. https://doi.org/10.1016/0304-405X(80)90020-3

[13] Ho, T. S. Y., & Macris, R. G. (1984). Dealer Bid-Ask Quotes and Transaction Prices: An Empirical Study of Some AMEX Options. *The Journal of Finance*. https://doi.org/10.1111/j.1540-6261.1984.tb03858.x

[14] Madhavan, A., & Smidt, S. (1993). An Analysis of Changes in Specialist Inventories and Quotations. *The Journal of Finance*. https://doi.org/10.1111/j.1540-6261.1993.tb05122.x

[15] Bayraktar, E., & Ludkovski, M. (2012). Liquidation in limit order books with controlled intensity. *Mathematical Finance*. https://doi.org/10.1111/j.1467-9965.2012.00529.x

[16] Cartea, A., Jaimungal, S., & Ricci, J. (2014). Buy low, sell high: a high frequency trading perspective. *SIAM Journal on Financial Mathematics*. https://doi.org/10.1137/130911196

[17] Gueant, O. (2017). Optimal market making. *Applied Mathematical Finance*. https://doi.org/10.1080/1350486X.2017.1342552

[18] Cartea, A., Donnelly, R., & Jaimungal, S. (2017). Algorithmic Trading with Model Uncertainty. *SIAM Journal on Financial Mathematics*. https://doi.org/10.1137/16m106282x

[19] Oksendal, B., & Sulem, A. (2007). *Applied Stochastic Control of Jump Diffusions*. Springer. https://doi.org/10.1007/978-3-540-69826-5