Optimal Almgren–Chriss Execution with CRRA Preferences

Grant Stenger

December 2024

Abstract

We study optimal execution in a linear-impact Almgren–Chriss model with an arithmetic-Brownian unaffected price. The trader chooses a deterministic open-loop inventory path subject to temporary and permanent price impact. Rather than solving $\max \mathbb{E}[U(W_T)]$ exactly, we characterize the stationary schedules of the second-order certainty-equivalent approximation to CRRA utility, $J[Q] = M[Q] - \frac{\gamma}{2} V[Q]/M[Q]$. Under the arithmetic specification, the stationarity condition is an exact constant-coefficient boundary-value problem, $\eta \ddot{Q}_t - \alpha Q_t = -\mu/2$, where α is a global scalar pinned down by a one-dimensional fixed-point equation. With fixed endpoints, the solution is a shifted sinh-bridge. When terminal inventory is free and permanent impact is absent, the transversality condition $\dot{Q}(T) = 0$ yields a shifted cosh-bridge with stationary terminal inventory $Q_T^* = Q_0/\cosh(\kappa T)$ in the zero-drift case—a single number that captures the stationary partial liquidation. Both bridge forms are structurally invariant: they are the exact stationary families for *any* open-loop objective $\Phi(M, V)$, not just the certainty-equivalent approximation; only the scalar $\alpha = -\sigma^2 \Phi_V / \Phi_M$ changes. In the fixed-endpoint problem, permanent impact and initial wealth do not enter the interior stationarity equation but affect the schedule globally through $M[Q]$; in the free-endpoint extension, permanent impact additionally enters the transversality condition.

1 Introduction

Optimal execution studies how to move a large position over a finite horizon while balancing price impact against market risk. Bertsimas and Lo (1998) derive dynamic optimal trading strategies that minimize expected execution cost. Almgren and Chriss (2000) then provide the canonical mean–variance benchmark, using an arithmetic random-walk specification for short horizons with linear temporary and permanent impact. That framework remains the reference point for both theory and practice because it separates the main economic forces cleanly and admits closed-form schedules.

Risk-averse execution has been studied from several angles. Schied and Schöneborn (2009) analyze optimal liquidation under general risk aversion. Almgren (2003) extends the framework to nonlinear impact. Forsyth et al. (2012) develop numerical methods for mean-quadratic-variation objectives, and Guéant et al. (2012) treat the complementary problem of limit-order execution.

This paper revisits the Almgren–Chriss benchmark under CRRA preferences, but in a deliberately modest way. We do not solve the exact problem $\max \mathbb{E}[U(W_T)]$—indeed, under ABM terminal wealth is Gaussian, so standard CRRA utility is not globally well-defined (Remark 4.1). Instead, with CRRA preferences as in Merton (1971) and the standard second-order certainty-equivalent expansion, we replace expected utility by the surrogate

$$J[Q] = M[Q] - \frac{\gamma}{2} \frac{V[Q]}{M[Q]}.$$

The resulting problem is still nontrivial because the variance penalty is scaled by the endogenous mean $M[Q]$. The main mathematical point is that the model must be specified consistently with the intended closed form. If one keeps geometric Brownian motion for the unaffected price, then the exact stationary equation under this objective has time-varying coefficients. Closed-form GBM execution results do exist in the Almgren–Chriss framework, but under different risk criteria; see Gatheral and Schied (2011) and Schied (2013). We therefore take the short-horizon arithmetic route from the outset, which is the natural setting for a genuine constant-coefficient sinh solution.

The main contributions are as follows. We derive terminal wealth, mean, and variance consistently with the arithmetic specification and deterministic open-loop controls, compute the first variation of the certainty-equivalent functional, and show that every stationary path satisfies the constant-coefficient ODE $\eta\ddot{Q}_t - \alpha Q_t = -\mu/2$ with a global scalar coefficient α. We solve this boundary-value problem exactly and obtain a shifted sinh-bridge (fixed endpoints) and, in the zero-permanent-impact free-endpoint case, a shifted cosh-bridge. The transversality condition pins down the stationary terminal inventory as $Q_T^* = q_\mu + (Q_0 - q_\mu)/\cosh(\kappa T)$; with zero drift this simplifies to $Q_T^* = Q_0/\cosh(\kappa T)$. We then prove a structural invariance result: both bridge forms are the exact stationary families for *any* open-loop objective $\Phi(M, V)$, not just the certainty-equivalent approximation; only the scalar $\alpha = -\sigma^2\Phi_V/\Phi_M$ changes. Invariance identifies the universal trajectory family; the CE-CRRA fixed point then determines *where on* that family a trader with specific risk preferences and wealth operates. Finally, we make precise how initial total wealth W_0 and linear permanent impact θ affect the schedule: in the fixed-endpoint problem, neither appears in the interior stationarity equation, but both change the fixed point for α through the mean term $M[Q]$. In the free-endpoint extension, θ additionally enters the transversality condition.

The paper characterizes stationary schedules of the certainty-equivalent approximation. When multiple positive roots of the fixed-point equation arise, we select the one with the largest objective value. We do not claim uniqueness of the positive root or global optimality of the CE functional under arbitrary parameters.

Our viewpoint is intentionally narrow. The schedule is deterministic and chosen at time 0; it is a pre-committed open-loop benchmark, not a dynamically adaptive trading rule. The arithmetic specification is a short-horizon approximation, not a long-horizon positive-price model. The CRRA objective is handled through a certainty-equivalent approximation, not exact expected-utility maximization. Within those limits, the paper gives a clean and internally consistent closed form.

The rest of the paper is organized as follows. Section 2 sets out the model. Section 3 derives terminal wealth, mean, and variance. Section 4 introduces the CRRA certainty-equivalent approximation. Section 5 gives the first variation and stationary equation. Section 6 solves the constant-coefficient boundary-value problem and derives the scalar fixed point. Section 7 extends the analysis to free terminal inventory and derives the cosh-bridge. Section 8 proves structural invariance of both bridges and illustrates it with an exact CARA example. Section 9 discusses comparative statics. Section 10 provides numerical illustrations. Section 11 concludes. Appendix A collects the integral formulas and Appendix B gives a short numerical recipe.

2 Model

2.1 Unaffected price and impact

We work over a fixed finite horizon $[0, T]$.

Assumption 2.1 (Short-horizon arithmetic benchmark)**.** The unaffected price process $\bar{S}$ follows the arithmetic Brownian motion

$$d\bar{S}_t = \mu\,dt + \sigma\,dB_t, \qquad \mu \in \mathbb{R},\ \sigma > 0, \tag{2.1}$$

where B is a standard Brownian motion.

Assumption 2.1 is the standard short-horizon Almgren–Chriss specification; see Almgren and Chriss (2000). On longer horizons one may prefer a multiplicative model, but then, under the present certainty-equivalent CRRA objective, the stationary equation no longer has constant coefficients (see Remark 5.3).

Following Almgren–Chriss, we distinguish the impacted mid-price and the execution price:

$$S_t^{\text{mid}} = \bar{S}_t + \theta(Q_t - Q_0), \qquad S_t^{\text{exec}} = S_t^{\text{mid}} + \eta v_t, \qquad v_t := \dot{Q}_t, \tag{2.2}$$

with temporary-impact coefficient $\eta > 0$ and permanent-impact coefficient $\theta \geq 0$.

The permanent term $\theta(Q_t - Q_0)$ implies no initial scar: at $t = 0$, $S_0^{\text{mid}} = \bar{S}_0$. It also implies a lasting footprint: under complete liquidation ($Q_T = 0$), the mid-price at the horizon is $S_T^{\text{mid}} = \bar{S}_T - \theta Q_0$. This is the classical Almgren–Chriss cumulative-flow specification. It differs from dynamic liquidity/resilience models such as Obizhaeva and Wang (2013), where the state variable and economic mechanism are different.

2.2 Inventory and admissible controls

Inventory evolves according to

$$dQ_t = v_t\,dt, \qquad Q(0) = Q_0, \quad Q(T) = Q_T. \tag{2.3}$$

We take $v_t > 0$ to mean net buying and $v_t < 0$ net selling. The classical liquidation problem corresponds to $Q_0 > 0$ and $Q_T = 0$, but we keep Q_T arbitrary.

Assumption 2.2 (Deterministic open-loop control)**.** The trading schedule is deterministic and chosen at time 0. The admissible class is

$$\mathcal{A} := \big\{Q \in C^2([0,T]) : Q(0) = Q_0,\ Q(T) = Q_T\big\}.$$

Under Assumption 2.2, all randomness in terminal wealth comes from the unaffected price process (2.1). The resulting schedule is a deterministic benchmark, not a feedback policy.

3 Terminal wealth, mean, and variance

We define total marked-to-market wealth by

$$W_t := C_t + Q_t S_t^{\text{mid}},$$

where C_t is cash. Because $S_0^{\text{mid}} = \bar{S}_0$, the initial total wealth is $W_0 = C_0 + Q_0\bar{S}_0$, where $\bar{S}_0 > 0$ is the initial unaffected price.

3.1 Self-financing identity

Because trades execute at S_t^{exec}, cash evolves as

$$dC_t = -S_t^{\text{exec}} v_t \, dt.$$

Therefore

$$\begin{aligned} dW_t &= dC_t + Q_t \, dS_t^{\text{mid}} + S_t^{\text{mid}} \, dQ_t \\ &= -S_t^{\text{exec}} v_t \, dt + Q_t \big(d\bar{S}_t + \theta \, dQ_t \big) + S_t^{\text{mid}} v_t \, dt \\ &= -\eta v_t^2 \, dt + \theta Q_t v_t \, dt + Q_t \, d\bar{S}_t. \end{aligned}$$

Using (2.1) gives the terminal-wealth representation below.

Proposition 3.1. *For any deterministic admissible inventory path Q,*

$$W_T = W_0 - \eta \int_0^T \dot{Q}_t^2 \, dt + \mu \int_0^T Q_t \, dt + \theta \int_0^T Q_t \dot{Q}_t \, dt + \sigma \int_0^T Q_t \, dB_t. \tag{3.1}$$

3.2 Mean and variance

Since endpoints are fixed,

$$\int_0^T Q_t \dot{Q}_t \, dt = \frac{1}{2} \big(Q_T^2 - Q_0^2 \big).$$

Taking expectations in (3.1) and using Itô isometry for the stochastic integral yields the next result.

Proposition 3.2. *For any deterministic admissible inventory path Q,*

$$M[Q] := \mathbb{E}[W_T] = W_0 - \eta \int_0^T \dot{Q}_t^2 \, dt + \mu \int_0^T Q_t \, dt + \frac{\theta}{2} \big(Q_T^2 - Q_0^2 \big), \tag{3.2}$$

$$V[Q] := \mathrm{Var}(W_T) = \sigma^2 \int_0^T Q_t^2 \, dt. \tag{3.3}$$

Remark 3.3 (Two structural points). First, linear permanent impact is path-independent in the mean once endpoints are fixed: it contributes only the boundary term $\frac{\theta}{2}(Q_T^2 - Q_0^2)$. Second, W_0 is initial total marked-to-market wealth, not initial cash alone, because wealth is defined by $W_t = C_t + Q_t S_t^{\text{mid}}$.

Remark 3.4 (Gaussian terminal wealth). Under Assumption 2.2, the only random term in (3.1) is the stochastic integral $\sigma \int_0^T Q_t \, dB_t$, so $W_T \sim \mathcal{N}(M[Q], V[Q])$. In particular, W_T has full real-line support whenever $V[Q] > 0$.

4 CRRA certainty-equivalent approximation

Let terminal utility be CRRA:

$$U(w) = \begin{cases} \dfrac{w^{1-\gamma} - 1}{1-\gamma}, & \gamma > 0, \ \gamma \neq 1, \\ \log w, & \gamma = 1, \end{cases} \tag{4.1}$$

with relative risk-aversion parameter $\gamma > 0$; see Merton (1971).

As in the standard local approximation behind certainty-equivalent formulas, expand $U(W_T)$ around $M[Q] = \mathbb{E}[W_T]$:

$$\mathbb{E}[U(W_T)] \approx U(M) + \frac{1}{2}U''(M)V.$$

For CRRA utility, $U'(w) = w^{-\gamma}$ and $U''(w) = -\gamma w^{-\gamma-1}$. Converting the approximation to money-metric form gives

$$\mathrm{CE}(Q) \approx M[Q] - \frac{\gamma}{2}\frac{V[Q]}{M[Q]}.$$

This motivates the working objective

$$J[Q] := M[Q] - \frac{\gamma}{2}\frac{V[Q]}{M[Q]}. \tag{4.2}$$

Remark 4.1 (Scope of the approximation). The functional (4.2) is a second-order certainty-equivalent approximation, not exact expected-utility maximization. The paper therefore does not claim an exact CRRA solution; it gives an exact stationary analysis of (4.2) in the arithmetic Almgren–Chriss model. The approximation is meaningful when $M[Q] > 0$ and $V[Q] \ll M[Q]^2$ along the schedules of interest. Under ABM with nonzero volatility, terminal wealth is Gaussian and therefore has positive probability of being nonpositive; exact CRRA expected utility is thus not well-defined without an additional wealth-flooring or utility-extension device. The CE approximation avoids this domain problem by working locally around positive mean wealth.

Remark 4.2 (Connection with mean–variance). If the relative variance is small, $V[Q] \ll M[Q]^2$, then the effective risk-weight in (4.2) is approximately $\gamma/(2M[Q])$. In this sense the certainty-equivalent objective behaves like a mean–variance problem with an endogenous wealth-dependent penalty.

5 First variation and stationary equation

Let $Q \in \mathcal{A}$ and let $h \in C^2([0,T])$ satisfy $h(0) = h(T) = 0$. Define the variation $Q^\varepsilon := Q + \varepsilon h$.

5.1 Variation of the mean and variance

Lemma 5.1. *The first variations of M and V at Q in the direction h are*

$$\delta M[Q;h] := \frac{\mathrm{d}}{\mathrm{d}\varepsilon}M[Q^\varepsilon]\Big|_{\varepsilon=0} = \int_0^T \big(2\eta\ddot{Q}_t + \mu\big)h_t\,\mathrm{d}t, \tag{5.1}$$

$$\delta V[Q;h] := \frac{\mathrm{d}}{\mathrm{d}\varepsilon}V[Q^\varepsilon]\Big|_{\varepsilon=0} = 2\sigma^2\int_0^T Q_t h_t\,\mathrm{d}t. \tag{5.2}$$

Proof. Differentiate (3.2)–(3.3) under the integral sign. For the temporary-impact term,

$$\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\left(-\eta\int_0^T (\dot{Q}_t + \varepsilon\dot{h}_t)^2\,\mathrm{d}t\right)\Big|_{\varepsilon=0} = -2\eta\int_0^T \dot{Q}_t\dot{h}_t\,\mathrm{d}t = 2\eta\int_0^T \ddot{Q}_t h_t\,\mathrm{d}t,$$

where the boundary term vanishes because $h(0) = h(T) = 0$. The drift term contributes $\mu\int_0^T h_t\,\mathrm{d}t$, and the permanent-impact boundary term does not vary because endpoints are fixed. This proves (5.1). Formula (5.2) is immediate from (3.3). $\square$

5.2 First variation of the certainty equivalent

Write $M := M[Q]$ and $V := V[Q]$ for brevity. The quotient rule gives

$$\delta J[Q;h] = \left(1 + \frac{\gamma V}{2M^2}\right) \delta M[Q;h] - \frac{\gamma}{2M} \delta V[Q;h].$$

Substituting Lemma 5.1 yields

$$\delta J[Q;h] = \int_0^T \left[\left(1 + \frac{\gamma V}{2M^2}\right) (2\eta \ddot{Q}_t + \mu) - \frac{\gamma \sigma^2}{M} Q_t \right] h_t \, \mathrm{d}t. \tag{5.3}$$

Hence every stationary path must satisfy the following first-order stationarity condition.

Proposition 5.2 (Stationary equation)**.** *If Q is a stationary point of* (4.2) *in the class $\mathcal{A}$ with $M[Q] > 0$, then*

$$\eta \ddot{Q}_t - \alpha Q_t = -\frac{\mu}{2}, \qquad 0 \le t \le T, \tag{5.4}$$

where

$$\alpha := \frac{\gamma \sigma^2 M}{2M^2 + \gamma V} = \frac{\gamma \sigma^2}{2M\left(1 + \frac{\gamma V}{2M^2}\right)}. \tag{5.5}$$

Proof. By (5.3), $\delta J[Q;h] = 0$ for every admissible variation h if and only if the bracketed term vanishes pointwise. Dividing by $2(1 + \gamma V/(2M^2))$ gives (5.4)–(5.5). □

Remark 5.3 (Why the arithmetic specification matters). Equation (5.4) is a constant-coefficient ODE because under the arithmetic model the variance functional is $\sigma^2 \int_0^T Q_t^2 \, \mathrm{d}t$. Under a GBM specification ($\mathrm{d}\bar{S}_t = \mu \bar{S}_t \, \mathrm{d}t + \sigma \bar{S}_t \, \mathrm{d}B_t$) with deterministic controls, the stochastic-integral contribution to the variance becomes

$$\sigma^2 \bar{S}_0^2 \int_0^T Q_t^2 \, e^{(2\mu + \sigma^2)t} \, \mathrm{d}t,$$

which already acquires time-varying coefficients. (Under GBM the drift integral $\mu \int Q_t \bar{S}_t \, \mathrm{d}t$ is also stochastic, so W_T is no longer Gaussian, (M,V)-sufficiency breaks down, and the full $\mathrm{Var}(W_T)$ has additional terms; the qualitative point is unchanged.) The exact sinh-bridge is therefore specific to the arithmetic model.

Remark 5.4 (Role of W_0 and θ). Neither W_0 nor θ appears in the local ODE (5.4). This does not mean that they are irrelevant. Both enter the mean term $M[Q]$ in (3.2), and therefore they change the global scalar α. In particular, permanent impact is path-independent in the mean but still affects the stationary schedule indirectly through the denominator of the certainty-equivalent penalty.

6 Closed-form solution and scalar fixed point

Section 5 shows that the entire local problem reduces to the constant-coefficient ODE (5.4). The remaining issue is global consistency, because α depends on $M[Q]$ and $V[Q]$.

6.1 Solution for fixed α

Fix $\alpha > 0$ and define

$$\kappa := \sqrt{\alpha/\eta}, \qquad q_\mu := \frac{\mu}{2\alpha}.$$

Also define the boundary basis functions

$$\phi_0(t;\kappa) := \frac{\sinh(\kappa(T-t))}{\sinh(\kappa T)}, \qquad \phi_T(t;\kappa) := \frac{\sinh(\kappa t)}{\sinh(\kappa T)}. \tag{6.1}$$

Proposition 6.1 (Shifted sinh-bridge)**.** *For fixed $\alpha > 0$, the unique C^2 solution of*

$$\eta \ddot{Q}_t - \alpha Q_t = -\frac{\mu}{2}, \qquad Q(0) = Q_0, \quad Q(T) = Q_T,$$

is

$$Q_t = q_\mu + (Q_0 - q_\mu)\phi_0(t;\kappa) + (Q_T - q_\mu)\phi_T(t;\kappa). \tag{6.2}$$

The corresponding trading rate is

$$v_t = \dot{Q}_t = -\kappa(Q_0 - q_\mu)\frac{\cosh(\kappa(T-t))}{\sinh(\kappa T)} + \kappa(Q_T - q_\mu)\frac{\cosh(\kappa t)}{\sinh(\kappa T)}. \tag{6.3}$$

Proof. A constant particular solution is $q_\mu = \mu/(2\alpha)$. Writing $Q_t = q_\mu + R_t$ reduces the ODE to $R_t'' - \kappa^2 R_t = 0$, whose boundary-value solution is the usual sinh-bridge. □

Formula (6.2) can also be written as

$$Q_t = Q_0\phi_0(t;\kappa) + Q_T\phi_T(t;\kappa) + q_\mu\big(1 - \phi_0(t;\kappa) - \phi_T(t;\kappa)\big). \tag{6.4}$$

Using the identity

$$\phi_0(t;\kappa) + \phi_T(t;\kappa) = \frac{\cosh(\kappa(t - T/2))}{\cosh(\kappa T/2)} \le 1, \tag{6.5}$$

we obtain a clean pointwise drift effect.

Corollary 6.2 (Direct effect of drift at fixed α)**.** *Holding α fixed,*

$$\frac{\partial Q_t}{\partial \mu} = \frac{1 - \phi_0(t;\kappa) - \phi_T(t;\kappa)}{2\alpha} \ge 0, \qquad 0 \le t \le T, \tag{6.6}$$

with strict inequality for $0 < t < T$. Hence positive drift shifts the inventory path upward and delays liquidation, while negative drift shifts it downward and accelerates liquidation.

Remark 6.3 (Classical zero-drift case). When $\mu = 0$, we have $q_\mu = 0$ and (6.2) reduces to the homogeneous Almgren–Chriss bridge:

$$Q_t = Q_0\frac{\sinh(\kappa(T-t))}{\sinh(\kappa T)} + Q_T\frac{\sinh(\kappa t)}{\sinh(\kappa T)}.$$

Remark 6.4 (Risk-neutral limit). If $\alpha = 0$, the stationary equation becomes $\eta\ddot{Q}_t = -\mu/2$, whose solution is the quadratic path

$$Q_{\mathrm{RN}}(t) = Q_0 + \frac{Q_T - Q_0}{T}t + \frac{\mu}{4\eta}t(T-t).$$

The shifted sinh-bridge connects continuously to this risk-neutral quadratic schedule as $\alpha \downarrow 0$.

6.2 Moments along the closed-form family

Let

$$a := Q_0 - q_\mu, \qquad b := Q_T - q_\mu, \qquad S := \sinh(\kappa T), \qquad C := \cosh(\kappa T).$$

Then (6.2) becomes $Q_t = q_\mu + R_t$ with

$$R_t = a\,\phi_0(t;\kappa) + b\,\phi_T(t;\kappa).$$

Proposition 6.5 (Integral formulas)**.** *For the family* (6.2),

$$I_0(\alpha) := \int_0^T Q_t\,\mathrm{d}t = q_\mu T + \frac{(a+b)(C-1)}{\kappa S}, \tag{6.7}$$

$$\begin{aligned} I_2(\alpha) &:= \int_0^T Q_t^2\,\mathrm{d}t \\ &= q_\mu^2 T + 2q_\mu \frac{(a+b)(C-1)}{\kappa S} + \frac{(a^2+b^2)C - 2ab}{2\kappa S} - \frac{T(a^2+b^2-2abC)}{2S^2}, \end{aligned} \tag{6.8}$$

$$I_v(\alpha) := \int_0^T \dot{Q}_t^2\,\mathrm{d}t = \frac{\kappa\big((a^2+b^2)C - 2ab\big)}{2S} + \frac{\kappa^2 T(a^2+b^2-2abC)}{2S^2}. \tag{6.9}$$

Consequently,

$$M(\alpha) = W_0 + \frac{\theta}{2}\big(Q_T^2 - Q_0^2\big) - \eta I_v(\alpha) + \mu I_0(\alpha), \qquad V(\alpha) = \sigma^2 I_2(\alpha). \tag{6.10}$$

Proof. The calculations are elementary and collected in Appendix A. □

6.3 Scalar fixed point

At a stationary path, the scalar α in (5.5) must agree with the moments generated by the corresponding closed-form path. Substituting (6.10) into (5.5) gives a one-dimensional fixed-point equation.

Proposition 6.6 (Scalar self-consistency condition)**.** *Any stationary schedule of the certainty-equivalent problem must satisfy*

$$\alpha = \frac{\gamma\sigma^2 M(\alpha)}{2M(\alpha)^2 + \gamma V(\alpha)}, \qquad \alpha > 0, \quad M(\alpha) > 0. \tag{6.11}$$

Equivalently, α is a positive root of the cross-multiplied form

$$F(\alpha) := \alpha\big(2M(\alpha)^2 + \gamma V(\alpha)\big) - \gamma\sigma^2 M(\alpha) = 0. \tag{6.12}$$

Conversely, any positive root of (6.11) *generates a stationary inventory path through* (6.2).

Proof. If Q is stationary, Proposition 5.2 shows that it solves (5.4) with α given by (5.5). By Proposition 6.1, that solution is the shifted sinh-bridge parametrized by α. Evaluating the moments along this path and substituting into (5.5) gives (6.11). Cross-multiplying yields the equivalent form (6.12). The converse is immediate. □

Equation (6.11) is the only nonlinear step in the construction. Once a positive root is found, the full trajectory is explicit. When convenient we write $M(\kappa) := M(\eta\kappa^2)$, $V(\kappa) := V(\eta\kappa^2)$, and

$$F_\kappa(\kappa) := \eta\kappa^2\big(2M(\kappa)^2 + \gamma V(\kappa)\big) - \gamma\sigma^2 M(\kappa),$$

so that the positive roots of F_κ correspond one-to-one with those of F in (6.12) via $\alpha = \eta\kappa^2$.

Remark 6.7 (Existence, feasibility, and root selection). As $\kappa \downarrow 0$, the candidate path tends to the risk-neutral quadratic $Q_{\mathrm{RN}}(t)$ from Remark 6.4, and

$$M(0^+) = W_0 + \frac{\theta}{2}(Q_T^2 - Q_0^2) - \frac{\eta(Q_T - Q_0)^2}{T} + \frac{\mu(Q_0 + Q_T)T}{2} + \frac{\mu^2 T^3}{48\eta}.$$

If $M(0^+) > 0$, then $F_\kappa(0^+) = -\gamma\sigma^2 M(0^+) < 0$. For any nontrivial problem with $Q_0^2 + Q_T^2 > 0$, the impact cost $I_v(\kappa) \sim \frac{1}{2}(Q_0^2 + Q_T^2)\kappa$ as $\kappa \to \infty$, so $M(\kappa) \to -\infty$. Let $\kappa_0 > 0$ be the first zero of $M(\kappa)$; at that point $F_\kappa(\kappa_0) = \eta\kappa_0^2\gamma V(\kappa_0) > 0$. By the intermediate value theorem, at least one feasible positive root exists on $(0, \kappa_0)$.

We do not claim uniqueness under arbitrary parameters. In the calibrations examined, there is typically a unique positive root. If several positive roots exist, evaluate $J(Q_\alpha) = M(\alpha) - \frac{\gamma}{2}V(\alpha)/M(\alpha)$ at each and retain the root with the largest objective value.

Remark 6.8 (Second-order conditions). The stationarity condition (5.4) is necessary but not sufficient for an interior maximum of J. We do not prove local or global sufficiency in the full path space. Numerically, the selected root yields the largest value of the reduced scalar objective $J(\kappa)$ among all positive roots we observe.

Remark 6.9 (Small-relative-risk regime). Using (5.5),

$$\alpha = \frac{\gamma\sigma^2}{2M(\alpha)}\left(1 + \frac{\gamma V(\alpha)}{2M(\alpha)^2}\right)^{-1}.$$

When $V(\alpha) \ll M(\alpha)^2$, this becomes

$$\alpha \approx \frac{\gamma\sigma^2}{2M(\alpha)}.$$

This approximation connects to the classical Almgren–Chriss intuition: higher expected wealth lowers the effective curvature parameter, while higher volatility and higher risk aversion raise it.

Remark 6.10 (Log utility). For $\gamma = 1$, the CRRA utility becomes logarithmic and all formulas remain valid with the certainty-equivalent weight $\gamma/2 = 1/2$.

Remark 6.11 (Monotonicity constraints). The admissible class $\mathcal{A}$ fixes only the endpoints. It does not forbid transient overshooting or round trips. For sufficiently strong drift, the stationary schedule from (6.2) can become non-monotone. Imposing monotonicity or no-round-trip constraints is a natural extension, but it is outside the present open-loop benchmark. (The free-endpoint cosh-bridge of Section 7 avoids this concern: it is automatically monotone; see Remark 7.4.)

7 Free terminal inventory and the cosh-bridge

The fixed-endpoint problem of Section 6 requires Q_T to be specified in advance. In many practical situations, the terminal inventory is itself a decision variable: the trader wants to reduce a position but need not fully liquidate. This section derives the stationary schedule when Q_T is free.

7.1 Transversality condition

Replace the fixed-endpoint admissible class $\mathcal{A}$ with

$$\mathcal{A}_{\mathrm{free}} := \{Q \in C^2([0,T]) : Q(0) = Q_0\}.$$

Variations h now satisfy only $h(0) = 0$; the value $h(T)$ is unrestricted.

Proposition 7.1 (Transversality condition)**.** *If Q is a stationary point of* (4.2) *in the class $\mathcal{A}_{\text{free}}$ with $M[Q] > 0$, then Q satisfies the ODE* (5.4) *on* $(0, T)$ *together with the boundary condition*

$$\dot{Q}_T = \frac{\theta}{2\eta} Q_T. \tag{7.1}$$

Proof. When $h(T) \neq 0$, two additional boundary terms appear in the first variation of J. Integration by parts in the temporary-impact term retains $-2\eta\dot{Q}_T\, h(T)$, and the permanent-impact boundary term $\frac{\theta}{2}Q_T^2$ contributes $\theta Q_T\, h(T)$. The full first variation is therefore

$$\delta J = \Lambda\Big[\int_0^T \big(2\eta\ddot{Q}_t + \mu - 2\alpha Q_t\big)h_t\,\mathrm{d}t + \big(-2\eta\dot{Q}_T + \theta Q_T\big)h(T)\Big],$$

where $\Lambda = 1 + \gamma V/(2M^2) > 0$. Setting $\delta J = 0$ for all admissible h yields (5.4) from the integral term and (7.1) from the boundary term. □

The transversality condition (7.1) balances two marginal effects at the terminal instant: the temporary-impact cost of a nonzero terminal trading rate ($2\eta\dot{Q}_T$ per share) against the marginal permanent-impact cost of further liquidation—equivalently, the marginal mean benefit θQ_T of retaining terminal inventory.

7.2 The cosh-bridge ($\theta = 0$)

When permanent impact is absent, the transversality condition reduces to $\dot{Q}(T) = 0$: arrive at the terminal inventory with zero trading speed. Any nonzero terminal rate wastes temporary-impact cost for no benefit.

Proposition 7.2 (Free-endpoint cosh-bridge)**.** *For fixed $\alpha > 0$ and $\theta = 0$, the unique C^2 solution of*

$$\eta\ddot{Q}_t - \alpha Q_t = -\frac{\mu}{2}, \qquad Q(0) = Q_0, \quad \dot{Q}(T) = 0,$$

is

$$Q_t = q_\mu + (Q_0 - q_\mu)\,\frac{\cosh(\kappa(T-t))}{\cosh(\kappa T)}. \tag{7.2}$$

The stationary terminal inventory is

$$Q_T^* = q_\mu + \frac{Q_0 - q_\mu}{\cosh(\kappa T)}, \tag{7.3}$$

and the trading rate is

$$v_t = -\kappa(Q_0 - q_\mu)\,\frac{\sinh(\kappa(T-t))}{\cosh(\kappa T)}. \tag{7.4}$$

Proof. Write $Q_t = q_\mu + R_t$, reducing the ODE to $\ddot{R}_t - \kappa^2 R_t = 0$. The general solution is $R_t = C_1\cosh(\kappa(T-t)) + C_2\sinh(\kappa(T-t))$. The Neumann condition $\dot{Q}(T) = \dot{R}(T) = -\kappa C_2 = 0$ forces $C_2 = 0$. Then $Q(0) = Q_0$ gives $C_1 = (Q_0 - q_\mu)/\cosh(\kappa T)$. □

When $\mu = 0$ (so $q_\mu = 0$), the formulas simplify to

$$Q_t = \frac{Q_0\cosh(\kappa(T-t))}{\cosh(\kappa T)}, \qquad Q_T^* = \frac{Q_0}{\cosh(\kappa T)}. \tag{7.5}$$

The fraction of the initial position retained is $1/\cosh(\kappa T)$—a single number that captures the stationary partial liquidation as a function of the composite parameter κT.

Remark 7.3 (Limiting behavior). The zero-drift formula $Q_T^* = Q_0/\cosh(\kappa T)$ has clean limits. As $\kappa T \to 0$ (short horizon or low risk aversion), $Q_T^* \to Q_0$: don't trade, because the cost of trading exceeds the risk-reduction benefit. As $\kappa T \to \infty$ (long horizon or high risk aversion), $Q_T^* \to 0$: fully liquidate, because risk dominates. For intermediate κT, the residual position decreases smoothly. With positive drift, Q_T^* is pulled toward $q_\mu = \mu/(2\alpha) > 0$, so the trader retains more inventory to capture anticipated price appreciation.

Remark 7.4 (Monotonicity of the cosh-bridge). Unlike the sinh-bridge, the cosh-bridge with $\theta = 0$ is automatically monotone: v_t from (7.4) does not change sign on $[0, T]$. When $Q_0 > q_\mu$, the path is a genuine liquidation schedule ($v_t < 0$ throughout); when $Q_0 < q_\mu$, it is an accumulation schedule. No round trips occur regardless of parameter values.

7.3 Moments and scalar fixed point

Define $a := Q_0 - q_\mu$ and let $S := \sinh(\kappa T)$, $C := \cosh(\kappa T)$ as before. For the cosh-bridge (7.2) with $\theta = 0$,

$$I_0^{\text{free}}(\alpha) := \int_0^T Q_t \,\mathrm{d}t = q_\mu T + \frac{aS}{\kappa C}, \tag{7.6}$$

$$I_2^{\text{free}}(\alpha) := \int_0^T Q_t^2 \,\mathrm{d}t = q_\mu^2 T + \frac{2q_\mu aS}{\kappa C} + \frac{a^2(SC + \kappa T)}{2\kappa C^2}, \tag{7.7}$$

$$I_v^{\text{free}}(\alpha) := \int_0^T \dot{Q}_t^2 \,\mathrm{d}t = \frac{\kappa a^2(SC - \kappa T)}{2C^2}. \tag{7.8}$$

The calculations are analogous to Appendix A, using $\int_0^T \cosh^2(\kappa u)\,\mathrm{d}u = (SC + \kappa T)/(2\kappa)$ and $\int_0^T \sinh^2(\kappa u)\,\mathrm{d}u = (SC - \kappa T)/(2\kappa)$.

The mean and variance are then

$$M^{\text{free}}(\alpha) = W_0 - \eta I_v^{\text{free}}(\alpha) + \mu I_0^{\text{free}}(\alpha), \qquad V^{\text{free}}(\alpha) = \sigma^2 I_2^{\text{free}}(\alpha),$$

and the scalar self-consistency condition takes the same form as (6.11):

$$\alpha = \frac{\gamma\sigma^2 M^{\text{free}}(\alpha)}{2M^{\text{free}}(\alpha)^2 + \gamma V^{\text{free}}(\alpha)}, \qquad \alpha > 0, \quad M^{\text{free}}(\alpha) > 0. \tag{7.9}$$

The computational recipe of Appendix B applies unchanged: scan for sign changes of the cross-multiplied form, refine by bisection, and select the root with the largest objective value.

Remark 7.5 (Existence of a free-endpoint root). The existence argument of Remark 6.7 carries over. Writing $F_\kappa^{\text{free}}(\kappa) = \eta\kappa^2(2M_{\text{free}}^2 + \gamma V_{\text{free}}) - \gamma\sigma^2 M_{\text{free}}$, we have $F_\kappa^{\text{free}}(0^+) < 0$ when $M_{\text{free}}(0^+) > 0$, and as $\kappa \to \infty$ the impact cost drives M_{free} through zero. At the first zero κ_0 of M_{free}, $F_\kappa^{\text{free}}(\kappa_0) > 0$. By the intermediate value theorem, at least one feasible positive root exists on $(0, \kappa_0)$.

Remark 7.6 (Permanent impact and round trips). When $\theta > 0$, the transversality condition (7.1) requires $\dot{Q}_T > 0$ whenever $Q_T > 0$—meaning the trader is *buying* at the terminal instant even while executing a net liquidation. The stationary path involves a small round trip near T: sell during most of the horizon, then buy back slightly at the end, because the permanent-impact cost of the last units of selling exceeds their risk-reduction benefit. This is mathematically valid but practically problematic, as most execution algorithms prohibit round trips. When the dimensionless ratio $\theta T/\eta$ is small, the terminal round trip is negligible ($\dot{Q}_T \approx 0$). The $\theta = 0$ case is therefore the cleanest and most practically relevant free-endpoint result. (For $\theta > 0$, the Robin boundary condition (7.1) replaces the Neumann condition, and the solution is a mixed cosh/sinh combination rather than the pure cosh-bridge; the problem remains closed-form for fixed α.)

8 Structural invariance

The shifted sinh-bridge and cosh-bridge are not artifacts of the certainty-equivalent approximation. The following result shows they are the exact stationary forms for *any* open-loop objective that depends on terminal wealth only through its first two moments.

Proposition 8.1 (Structural invariance of the sinh-bridge)**.** *Under ABM with deterministic controls, let* $\Phi(M, V)$ *be any smooth objective with* $\Phi_M > 0$ *and* $\Phi_V < 0$*. Then every stationary path of* $\Phi(M[Q], V[Q])$ *satisfies the same constant-coefficient ODE*

$$\eta \ddot{Q}_t - \alpha Q_t = -\frac{\mu}{2}, \tag{8.1}$$

with

$$\alpha = -\frac{\sigma^2 \Phi_V}{\Phi_M}. \tag{8.2}$$

In particular, the fixed-endpoint trajectory is always the shifted sinh*-bridge* (6.2)*; only the scalar* α *depends on the choice of* Φ*.*

Proof. Under ABM with deterministic controls, W_T is Gaussian with mean M and variance V. For any smooth $\Phi(M, V)$, the first variation is

$$\delta\Phi = \Phi_M\, \delta M + \Phi_V\, \delta V = \int_0^T \Big[\Phi_M \big(2\eta \ddot{Q}_t + \mu\big) + 2\sigma^2 \Phi_V\, Q_t\Big] h_t \, \mathrm{d}t,$$

using the variations from Lemma 5.1. Setting this to zero for all admissible h and dividing by $2\Phi_M$ yields the ODE with $\alpha = -\sigma^2 \Phi_V / \Phi_M$. □

Corollary 8.2 (Invariance of the free-endpoint condition)**.** *For any* $\theta \geq 0$*, the transversality condition* $\dot{Q}(T) = \theta Q_T/(2\eta)$ *from Proposition 7.1 holds for any smooth objective* $\Phi(M, V)$ *with* $\Phi_M \neq 0$*. In particular, when* $\theta = 0$ *the free-endpoint stationary family is always the shifted* cosh*-bridge* (7.2)*, and* Q_T^* *always has the form* (7.3)*; when* $\theta > 0$ *the stationary family is always the mixed* cosh/sinh *solution of the Robin problem. In both cases, only the scalar* α *(and hence* κ*) depends on the choice of objective.*

Proof. In the first variation of $\Phi(M[Q], V[Q])$ with free terminal inventory, the boundary term at $t = T$ is $\Phi_M \cdot (-2\eta \dot{Q}_T + \theta Q_T)\, h(T)$. The variance functional $V[Q] = \sigma^2 \int_0^T Q_t^2 \, \mathrm{d}t$ produces no endpoint contribution, so Φ_V does not appear in the boundary term. Since $\Phi_M \neq 0$, setting this to zero gives $\dot{Q}(T) = \theta Q_T/(2\eta)$. □

This means a trader using CE-CRRA, mean–variance, CARA, or any other (M, V)-based criterion would follow the same cosh-bridge trajectory family and the same formula for Q_T^*; only the quantitative value of α (and hence how aggressively the path trades and what numerical Q_T^* results) differs across objectives.

The certainty-equivalent approximation (4.2) is one particular choice: $\Phi_{\mathrm{CE}}(M, V) = M - \frac{\gamma}{2} V/M$, giving $\alpha_{\mathrm{CE}} = \gamma\sigma^2 M/(2M^2 + \gamma V)$, which recovers (5.5). But any other objective in the class $\Phi(M, V)$—including mean–variance or exponential utility—yields the same trajectory form. The approximation is therefore not load-bearing for the trajectory shape; it matters only for root selection.

8.1 CARA utility as an exact example

The invariance result is particularly clean for CARA (exponential) utility. Under Gaussian terminal wealth, exact expected CARA utility is

$$\mathbb{E}\big[-\tfrac{1}{\lambda}e^{-\lambda W_T}\big] = -\tfrac{1}{\lambda}\exp\left(-\lambda M + \frac{\lambda^2}{2}V\right),$$

which is a well-defined smooth objective $\Phi(M, V)$ for all (M, V) with $\Phi_M > 0$ and $\Phi_V < 0$. Applying (8.2) gives

$$\alpha_{\mathrm{CARA}} = \frac{\lambda\sigma^2}{2}, \tag{8.3}$$

a constant independent of M and V. No fixed-point equation is needed: the CARA stationary schedule is the shifted sinh-bridge with $\kappa_{\mathrm{CARA}} = \sqrt{\lambda\sigma^2/(2\eta)}$.

This contrasts with the CE objective, where α depends on M and V through (5.5) and therefore requires solving the scalar fixed-point equation (6.11). In the small-risk regime $V \ll M^2$, matching the local Arrow–Pratt coefficient $\lambda = \gamma/M$ gives $\alpha_{\mathrm{CARA}} = \gamma\sigma^2/(2M)$, which is precisely the small-V/M^2 approximation to α_{CE} (Remark 6.9). The two objectives therefore produce nearly identical schedules whenever V/M^2 is small—which is exactly the regime in which the CE approximation is valid.

Remark 8.3 (Why not exact CRRA under ABM?). Under ABM with nonzero volatility, W_T is Gaussian and $P(W_T \le 0) > 0$. Since CRRA utility is defined only on $w > 0$, exact expected CRRA utility $\mathbb{E}[U(W_T)]$ is not well-defined without an additional domain convention (e.g., a wealth floor or utility extension). Moreover, even restricting the Gaussian density to $w > 0$, the derivative $\Phi_M = \mathbb{E}[W_T^{-\gamma}]$ diverges for $\gamma \ge 1$ because $w^{-\gamma}$ is not integrable near zero against a density that is strictly positive there. CARA utility avoids both problems entirely, and the CE approximation avoids the domain issue by working locally around positive mean wealth.

9 Comparative statics

This section separates *local* effects visible directly in (6.2) from *global* effects operating through the fixed point (6.11).

9.1 Local effects in the ODE

At fixed α, the inventory path is determined by a simple shifted bridge. Corollary 6.2 shows that positive drift raises the inventory path pointwise. This is the correct economic sign for a liquidation problem: if the asset is expected to drift upward, the trader optimally holds more inventory for longer. Negative drift has the opposite effect.

The temporary-impact coefficient η enters locally through $\kappa = \sqrt{\alpha/\eta}$. Holding α fixed, a larger η lowers κ and flattens the bridge. Intuitively, when rapid trading is more expensive, the schedule becomes smoother.

By contrast, neither W_0 nor θ appears directly in the local ODE. Their effect is entirely global.

9.2 Global effects through the fixed point

The scalar equation (6.11) makes the global channel transparent. At fixed V, differentiating (5.5) gives

$$\left.\frac{\partial\alpha}{\partial M}\right|_V = \frac{\gamma\sigma^2(\gamma V - 2M^2)}{(2M^2 + \gamma V)^2}. \tag{9.1}$$

This is negative precisely when $V/M^2 < 2/\gamma$—the small-relative-risk regime in which the CE approximation is most meaningful (Remark 6.9). The directional statements below hold in that regime.

Volatility σ. A larger σ raises the numerator in (6.11) and therefore tends to increase α. The bridge steepens and inventory is reduced more quickly. This is the classical risk channel.

Risk aversion γ. A larger γ scales the certainty-equivalent penalty and also tends to increase α. Thus more risk-averse traders liquidate faster in the deterministic benchmark.

Initial total wealth W_0. A larger W_0 raises $M(\alpha)$ without changing the local ODE. By (9.1), in the small-risk regime ($V/M^2 < 2/\gamma$), higher M lowers α. The resulting schedule is flatter. This is a purely global wealth effect.

Permanent impact θ. The permanent-impact contribution in (3.2) is the boundary term $\frac{\theta}{2}(Q_T^2 - Q_0^2)$. For the liquidation case $Q_T = 0 < Q_0$, this term is negative. Hence larger θ lowers $M(\alpha)$, which by (9.1) increases α in the small-risk regime. The key point is conceptual: θ is not a local running-cost term in the stationarity equation. Its influence is indirect and global.

Drift μ. Drift has both a local and a global effect. Locally, it enters through the shift $q_\mu = \mu/(2\alpha)$ and therefore delays liquidation when positive. Globally, positive drift raises $M(\alpha)$, which by (9.1) lowers α in the small-risk regime and further flattens the schedule. In the arithmetic model these two mechanisms point in the same direction.

These comparative-static statements are approximate: the exact dependence runs through the nonlinear equation (6.12), in which M and V both vary simultaneously with the parameter of interest.

9.3 Comparative statics for the free-endpoint problem

In the cosh-bridge with $\mu = 0$, the stationary terminal inventory $Q_T^* = Q_0/\cosh(\kappa T)$ depends on parameters only through κ. Every comparative-static channel above operates in the same direction: parameters that raise α (and hence κ) reduce Q_T^*, and conversely. For instance, larger σ or γ increases κ and reduces the retained fraction $1/\cosh(\kappa T)$; larger W_0 lowers κ and increases Q_T^*. Figure 3(b) illustrates the γ channel directly. With nonzero drift, Q_T^* is additionally pulled toward $q_\mu = \mu/(2\alpha)$, reinforcing the local drift effect of Corollary 6.2.

10 Numerical illustrations

Unless otherwise noted, all numerical examples in this section concern the fixed-endpoint liquidation problem ($Q_T = 0$) and solve the scalar equation (6.11) by a one-dimensional root search on the equivalent κ formulation, with $M(\kappa)$ and $V(\kappa)$ obtained from Section 6. Parameter values are chosen so that $M(\alpha) > 0$ and the illustrated liquidation paths remain monotone.

Table 1 reports the fixed-point solutions. The figures line up with the theory. Panel (a) shows the pointwise drift effect from Corollary 6.2: positive drift shifts the path upward. Panel (b) shows that higher γ increases the fixed-point coefficient α and therefore accelerates liquidation. Panel (c) shows the purely global role of initial wealth: W_0 never appears in the local ODE, yet it materially changes the schedule through the denominator of the certainty-equivalent penalty.

Table 1: Certainty-equivalent fixed-point solutions used in Figure 1. The column α_{CE} is the positive root of (6.12); V/M^2 measures the relative variance at that root. The CE approximation is most reliable when this ratio is small (Remark 6.9).

Panel	Parameter	α_{CE}	V/M^2
Drift	$\mu = -0.2$	0.1018	0.0024
Drift	$\mu = 0$	0.1008	0.0030
Drift	$\mu = 0.2$	0.0997	0.0037
Risk aversion	$\gamma = 1$	0.0505	0.0032
Risk aversion	$\gamma = 2$	0.1008	0.0030
Risk aversion	$\gamma = 4$	0.2013	0.0027
Initial wealth	$W_0 = 5$	0.2002	0.0121
Initial wealth	$W_0 = 10$	0.1003	0.0033
Initial wealth	$W_0 = 20$	0.0501	0.0009

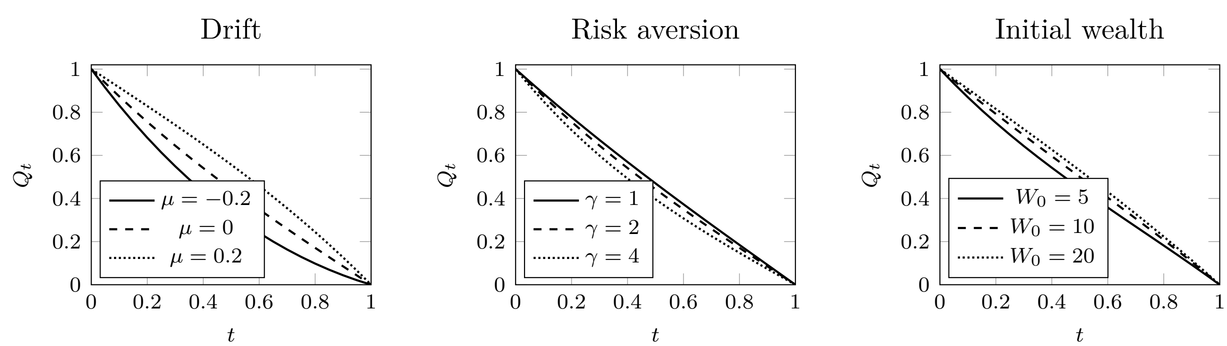

(a) Baseline $Q_0 = 1$, $Q_T = 0$, $T = 1$, $\eta = 0.1$, $\sigma = 1$, $\theta = 0.02$, $W_0 = 10$, $\gamma = 2$. Positive drift slows liquidation.

(b) Same baseline as panel (a) but with $\mu = 0$. Higher risk aversion increases α and steepens the bridge.

(c) Baseline $Q_0 = 1$, $Q_T = 0$, $T = 1$, $\eta = 0.1$, $\sigma = 1$, $\theta = 0.02$, $\mu = 0.1$, $\gamma = 2$. Larger initial wealth lowers the effective risk penalty and slows liquidation.

Figure 1: Numerical illustrations of the shifted sinh-bridge.

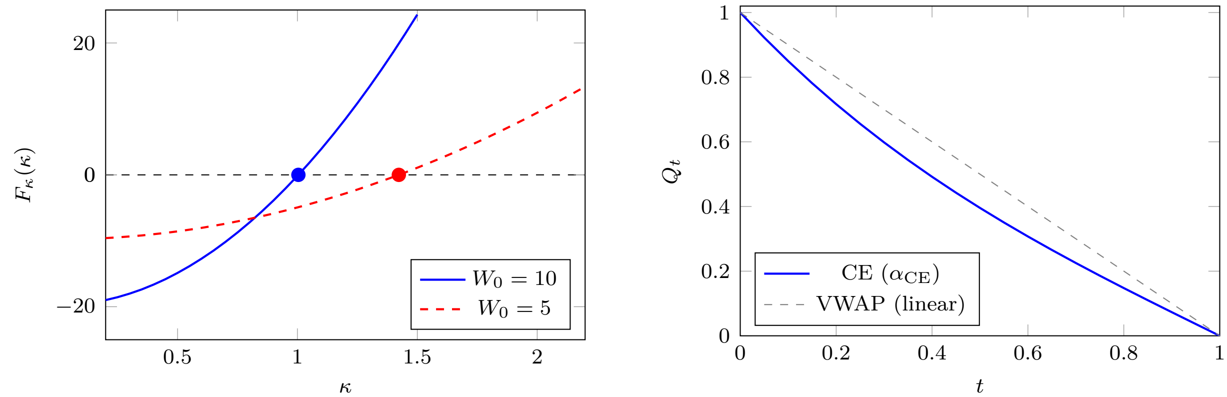

(a) The cross-multiplied fixed-point function $F_\kappa(\kappa)$. Each curve has a single zero crossing in these calibrations, illustrating a unique admissible positive root. Lower initial wealth shifts the root rightward (higher κ, more aggressive liquidation).

(b) CE inventory path for $W_0 = 5$ (stress case, $V/M^2 \approx 0.012$). The CE schedule front-loads liquidation relative to the linear VWAP benchmark.

Figure 2: Root structure and schedule comparison. Panel (a) shows the fixed-point function $F_\kappa(\kappa)$ for two wealth levels; panel (b) shows the CE schedule in the $W_0 = 5$ stress case against the linear VWAP benchmark.

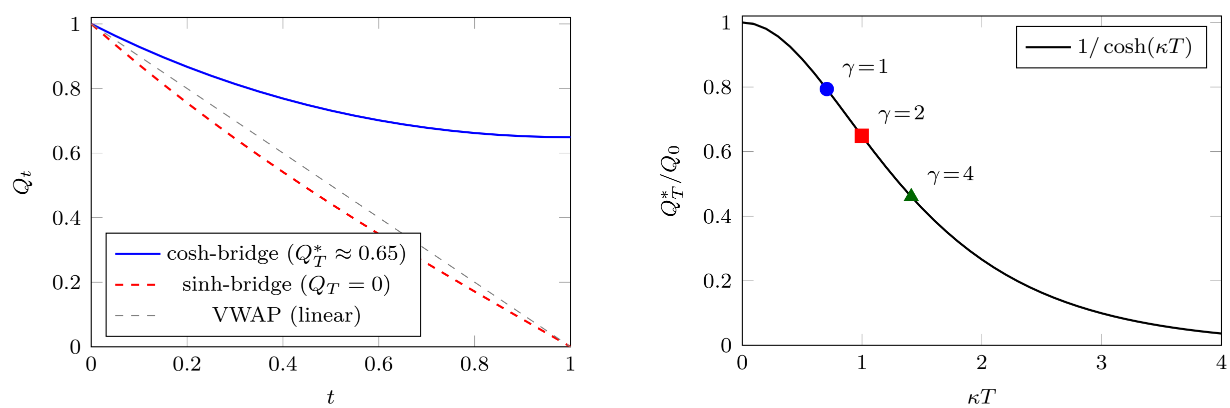

(a) Fixed-endpoint sinh-bridge (full liquidation) versus free-endpoint cosh-bridge (optimal partial liquidation) for baseline parameters with $\mu = 0$, $\gamma = 2$. The cosh-bridge retains $\approx 65\%$ of the initial position and arrives with zero trading speed.

(b) Retained fraction $Q_T^*/Q_0 = 1/\cosh(\kappa T)$ as a universal function of κT (zero drift, $\theta = 0$). Dots mark the CE-CRRA operating points for $\gamma = 1, 2, 4$ at baseline parameters. Higher risk aversion increases κ and reduces the retained fraction.

Figure 3: Free-endpoint cosh-bridge from Section 7. Panel (a) compares the two bridge types; panel (b) shows the universal retained-fraction curve.

The $W_0 = 5$ row is the stress test: with $V/M^2 \approx 0.012$, the CE approximation is at its least reliable. Figure 2(b) plots the CE schedule for this case. Even here, the schedule front-loads liquidation relative to the linear VWAP benchmark in the expected way. As noted in Section 8, the CARA objective with matched Arrow–Pratt coefficient $\lambda = \gamma/M$ gives $\alpha_{\text{CARA}} = \gamma\sigma^2/(2M)$, which differs from α_{CE} only at order V/M^2. In the regime where V/M^2 is small—as in all rows of Table 1—the CE and CARA schedules are nearly indistinguishable.

Figure 3 illustrates the free-endpoint cosh-bridge from Section 7. Panel (a) overlays the cosh-bridge against the sinh-bridge for the same baseline parameters with $\mu = 0$ and $\gamma = 2$. The cosh-bridge retains approximately 65% of the initial position ($Q_T^* \approx 0.649$), arriving with zero trading speed, while the sinh-bridge fully liquidates by construction. Panel (b) shows the universal retained-fraction curve $Q_T^*/Q_0 = 1/\cosh(\kappa T)$ from (7.5). At the baseline parameters, the three risk aversion levels $\gamma = 1, 2, 4$ correspond to $\kappa T \approx 0.71, 1.00, 1.41$ and retained fractions of approximately 79%, 65%, and 46% respectively. As expected, more risk-averse traders sell more of their position.

Table 2 reports the free-endpoint fixed-point solutions for varying γ. The retained fractions decrease monotonically with risk aversion, and the relative variance V/M^2 remains small throughout, confirming that the CE approximation is meaningful in these calibrations. Note that α_{CE} scales almost exactly linearly with γ, consistent with the small-risk approximation $\alpha \approx \gamma\sigma^2/(2M)$ from Remark 6.9.

Table 2: Free-endpoint ($\theta = 0$, $\mu = 0$) fixed-point solutions for baseline parameters $Q_0 = 1$, $T = 1$, $\eta = 0.1$, $\sigma = 1$, $W_0 = 10$. Columns: risk aversion γ, CE curvature α_{CE}, stationary terminal inventory Q_T^*, retained fraction Q_T^*/Q_0, and relative variance V/M^2.

γ	α_{CE}	Q_T^*	Q_T^*/Q_0	V/M^2
1	0.0498	0.794	79%	0.0075
2	0.0996	0.649	65%	0.0059
4	0.1991	0.460	46%	0.0042

11 Conclusion

This paper derives a mathematically consistent short-horizon execution benchmark under CRRA preferences. The unaffected price is modeled as an arithmetic Brownian motion, which matches the classical Almgren–Chriss specification and makes the certainty-equivalent first-order condition a constant-coefficient ODE. The paper produces two closed-form stationary schedules. With fixed endpoints, the solution is a shifted sinh-bridge:

$$Q_t = q_\mu + (Q_0 - q_\mu)\phi_0(t) + (Q_T - q_\mu)\phi_T(t).$$

With free terminal inventory and $\theta = 0$, the transversality condition $\dot{Q}(T) = 0$ yields a shifted cosh-bridge:

$$Q_t = q_\mu + (Q_0 - q_\mu)\frac{\cosh(\kappa(T-t))}{\cosh(\kappa T)}, \qquad Q_T^* = q_\mu + \frac{Q_0 - q_\mu}{\cosh(\kappa T)}.$$

The fraction retained, $1/\cosh(\kappa T)$ in the zero-drift case, captures the stationary partial liquidation as a function of a single composite parameter.

Three conceptual points are as important as the closed forms themselves. First, the results are deterministic open-loop benchmarks, not adaptive or time-consistent feedback rules. Second, they solve a second-order certainty-equivalent approximation to CRRA utility, not exact expected-utility maximization. Third, in the fixed-endpoint problem, permanent impact and initial total wealth affect the schedule only through the global scalar fixed point for α; they do not enter as local running-cost terms in the interior stationarity equation. In the free-endpoint extension, θ additionally enters the transversality condition. Moreover, both bridge forms are structurally invariant across all open-loop objectives that depend on terminal wealth only through its mean and variance (Proposition 8.1), so the certainty-equivalent approximation affects only the scalar α, not the trajectory family. The free-endpoint transversality condition $\dot{Q}(T) = \theta Q_T/(2\eta)$ is likewise invariant for all $\theta \geq 0$ (Corollary 8.2). CARA (exponential) utility provides a clean exact example: under Gaussian terminal wealth, $\alpha_{\text{CARA}} = \lambda\sigma^2/2$ is a constant requiring no fixed point, and it coincides with the CE root to first order in V/M^2 when the Arrow–Pratt coefficients are matched.

Several extensions are natural. One could impose monotonicity or no-round-trip constraints, introduce transient impact or resilience as in Obizhaeva and Wang (2013), or move to GBM and accept the resulting time-varying stationary equation. Under transient impact the permanent-impact contribution would no longer reduce to a path-independent boundary term, restoring genuine path-dependence in the mean and making the problem substantially richer. But for the short-horizon arithmetic model with linear impact considered here, the paper provides a coherent closed-form benchmark and a transparent basis for further work.

A Integral formulas for the shifted bridge

This appendix derives the formulas in Proposition 6.5. Write

$$Q_t = q_\mu + R_t, \qquad R_t = a\,\phi_0(t;\kappa) + b\,\phi_T(t;\kappa),$$

with $a = Q_0 - q_\mu$, $b = Q_T - q_\mu$, $S = \sinh(\kappa T)$, and $C = \cosh(\kappa T)$.

Integral of Q_t

Since

$$\int_0^T \phi_0(t;\kappa)\,\mathrm{d}t = \int_0^T \phi_T(t;\kappa)\,\mathrm{d}t = \frac{C-1}{\kappa S},$$

we get

$$\int_0^T Q_t\,\mathrm{d}t = q_\mu T + \frac{(a+b)(C-1)}{\kappa S},$$

which is (6.7).

Integral of Q_t^2

We have

$$\int_0^T Q_t^2\,\mathrm{d}t = q_\mu^2 T + 2q_\mu \int_0^T R_t\,\mathrm{d}t + \int_0^T R_t^2\,\mathrm{d}t.$$

Using

$$2\sinh u\,\sinh v = \cosh(u+v) - \cosh(u-v),$$

a direct calculation gives

$$\int_0^T R_t^2\,\mathrm{d}t = \frac{(a^2+b^2)C - 2ab}{2\kappa S} - \frac{T(a^2+b^2-2abC)}{2S^2},$$

and hence (6.8).

Integral of $\dot{Q}_t^2$

Because q_μ is constant,

$$\dot{Q}_t = \dot{R}_t = -\kappa a \frac{\cosh(\kappa(T-t))}{S} + \kappa b \frac{\cosh(\kappa t)}{S}.$$

Using

$$2\cosh u \, \cosh v = \cosh(u+v) + \cosh(u-v)$$

and the elementary integral

$$\int_0^T \cosh^2(\kappa t)\, \mathrm{d}t = \frac{T}{2} + \frac{SC}{2\kappa},$$

we obtain

$$\int_0^T \dot{Q}_t^2 \, \mathrm{d}t = \frac{\kappa\big((a^2+b^2)C - 2ab\big)}{2S} + \frac{\kappa^2 T(a^2+b^2-2abC)}{2S^2},$$

which is (6.9).

B Numerical root-finding recipe

The scalar fixed point can be solved in either α or $\kappa = \sqrt{\alpha/\eta}$. The cross-multiplied form (6.12), rewritten in terms of κ, avoids nested fractions and is numerically safer near small M:

$$F_\kappa(\kappa) = \eta\kappa^2\big(2M(\kappa)^2 + \gamma V(\kappa)\big) - \gamma\sigma^2 M(\kappa) = 0.$$

A practical recipe is: (i) scan F on a logarithmic grid in κ to bracket all sign changes; (ii) refine each bracket with Brent's method or bisection; (iii) if multiple positive roots are found, evaluate $J(Q_\alpha) = M(\alpha) - \frac{\gamma}{2}V(\alpha)/M(\alpha)$ at each and retain the root with the largest objective value.

Listing 1: Minimal mpmath implementation.

```
import mpmath as mp

def moments(kappa, Q0, QT, T, eta, mu, sigma, theta, W0):
    alpha = eta * kappa**2
    qmu = mu / (2.0 * alpha)
    a, b = Q0 - qmu, QT - qmu
    S, C = mp.sinh(kappa * T), mp.cosh(kappa * T)

    I0 = qmu * T + (a + b) * (C - 1) / (kappa * S)
    I2 = (qmu**2 * T
          + 2 * qmu * (a + b) * (C - 1) / (kappa * S)
          + ((a**2 + b**2) * C - 2 * a * b) / (2 * kappa * S)
          - T * (a**2 + b**2 - 2 * a * b * C) / (2 * S**2))
    Iv = (kappa * ((a**2 + b**2) * C - 2 * a * b) / (2 * S)
          + kappa**2 * T * (a**2 + b**2 - 2*a*b*C) / (2*S**2))

    M = W0 + 0.5 * theta * (QT**2 - Q0**2) - eta * Iv + mu * I0
    V = sigma**2 * I2
    return M, V

def F(kappa, Q0, QT, T, eta, mu, sigma, theta, W0, gamma):
    M, V = moments(kappa, Q0, QT, T, eta, mu, sigma, theta, W0)
```

```
    if M <= 0:
        return float('inf')  # outside admissible domain
    return eta * kappa**2 * (2*M**2 + gamma*V) - gamma*sigma**2*M

# Practical use:
# 1) scan F on a log grid in kappa, restricting to M(kappa) > 0;
# 2) bracket sign changes of F within the admissible region;
# 3) refine each bracket with Brent or bisection;
# 4) evaluate J = M - 0.5*gamma*V/M at each root;
# 5) keep the root with the largest J.
#
# Note: for very small kappa, qmu = mu/(2*eta*kappa**2)
# diverges and the formulas suffer cancellation. In that
# regime, switch to the risk-neutral limits:
#   I0 -> T/2*(Q0+QT) + mu*T**3/(24*eta)
#   Iv -> (QT-Q0)**2/T + mu**2*T**3/(48*eta**2)
#   I2 -> T/3*(Q0**2+Q0*QT+QT**2)
#         + mu*T**3/(24*eta)*(Q0+QT)
#         + mu**2*T**5/(480*eta**2)
```

References

Almgren, R. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10(1), 1–18.

Almgren, R. and Chriss, N. (2000). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.

Bertsimas, D. and Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets*, 1(1), 1–50.

Forsyth, P. A., Kennedy, J. S., Tse, S. T., and Windcliff, H. (2012). Optimal trade execution: a mean-quadratic-variation approach. *Journal of Economic Dynamics and Control*, 36(12), 1971–1991.

Gatheral, J. and Schied, A. (2011). Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. *International Journal of Theoretical and Applied Finance*, 14(3), 353–368.

Guéant, O., Lehalle, C.-A., and Fernandez-Tapia, J. (2012). Optimal portfolio liquidation with limit orders. *SIAM Journal on Financial Mathematics*, 3(1), 740–764.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), 373–413.

Obizhaeva, A. A. and Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1), 1–32.

Schied, A. (2013). Robust strategies for optimal order execution in the Almgren–Chriss framework. *Applied Mathematical Finance*, 20(3), 264–286.

Schied, A. and Schöneborn, T. (2009). Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, 13(2), 181–204.