Optimal Market Making under Endogenous Information Risk in Prediction Markets

Partial Observation, Flow-Modulated Jumps, and Hybrid Liquidity Control

Conrad O. Voigt
University of California, Santa Barbara

August 2026

Abstract

Prediction-market makers earn the spread while carrying two linked risks: fills can reveal private information, and the marked probability can jump before inventory is unwound. This paper formulates liquidity provision as a partially observed stochastic control problem in which public order flow, probability jumps, and the maker's own executions jointly identify a latent information regime. The public price is a martingale on the probability simplex; a bounded flow state changes the regime-dependent jump compensator and fill toxicity. The maker chooses bid and ask offsets, may withdraw either side, and may use costly market orders to reduce inventory. These controls lead to a nonlocal Hamilton–Jacobi–Bellman quasi-variational inequality. Under boundedness, regularity, viability, and no-free-loop conditions, the separated value function is characterized as the unique viscosity solution. For mutually exclusive outcomes, settlement risk is governed by the categorical covariance $\mathrm{diag}(p) - pp^{\mathsf{T}}$, which makes complete-set inventory a null direction and permits a conditional reduction from K inventory coordinates to $K-1$.

The model is evaluated only with controlled synthetic experiments; no market data are used. A finite-state implementation solves the filter and the hybrid control problem by backward induction and evaluates eight policies with common random numbers. An independent implementation of the diffusion-only passive benchmark of Feil and Nendel [2026] reproduces its optimal-policy P&L standard deviation (10.38 versus 10.34) and mean absolute terminal inventory (15.26 versus 15.23). In the correctly specified leakage scenario, conditioning on flow raises the mean control objective by 0.00163 (paired standard error 0.00033), the impulse option adds 0.00676 (0.00038), and partial observation costs 0.00810 (0.00055) relative to a full-information policy. A duration-aware filter reduces the Brier score from 0.1408 to 0.1318 under non-geometric regime durations. In a three-outcome diagnostic, a diagonal variance approximation misranks 27.2% of random inventory pairs. The results identify where endogenous information signals, intervention rights, and exact settlement geometry alter optimal market-making decisions.

Keywords: prediction markets; market making; stochastic control; partial observation; impulse control; jump diffusion; hidden semi-Markov model.
JEL classification: C61, C63, D82, G13, G14.

1 Introduction

A binary event contract pays one unit if an event occurs and zero otherwise. Its arbitrage-free price therefore lies in $[0, 1]$, and under a pricing measure it is naturally represented as a conditional probability. This bounded payoff changes the market-making problem. Ordinary inventory risk disappears after the event is known but becomes categorical settlement risk before resolution. Volatility is state dependent, probability jumps are economically meaningful rather than exceptional, and a large trade may be informative precisely because it precedes a public repricing. A maker who treats flow, price risk, and execution as separate modules can therefore quote aggressively in the state in which fills are most adverse.

The control problem studied here has three features. First, the public probability is a martingale on the simplex, with a jump compensator that depends on a bounded order-flow state and an unobserved information regime. Second, the maker learns about that regime from external flow, public price changes, and its own fills. Because the likelihood of a fill depends on the quote, the filter is control dependent. Third, the maker can quote, withdraw one or both sides, or cross the spread to reduce inventory. The last choice is an impulse control, not a limiting quote. Together these elements produce a partially observed hybrid control problem and a Hamilton–Jacobi–Bellman quasi-variational inequality (HJB–QVI).

The setting is an order-book market, not a scoring-rule or constant-function automated market maker. The distinction matters. A scoring-rule mechanism chooses a deterministic cost function and absorbs trader demand, whereas an order-book maker chooses prices at which executions arrive stochastically. Work on bounded-loss and uniform-loss automated makers addresses the former design problem [Hanson, 2003, Chen and Pennock, 2007, Moallemi et al., 2026]. The present problem concerns an inventory-bearing agent whose limit orders face execution and adverse-selection risk.

The paper makes four contributions.

1. It gives a probability-simplex market model in which external flow changes both the distribution of public jumps and the informativeness of executions, while preserving the martingale restriction. The latent regime may be represented by a phase-type chain, allowing non-geometric durations without losing a finite-dimensional filter.

2. It derives the separated stochastic-control problem when the maker's observations include its own quote-dependent fills. Quote placement has a dual role: it generates spread revenue and changes the information arriving to the controller. Passive quoting, withdrawal, and market-order inventory reduction are treated in one HJB–QVI.

3. It identifies the exact multi-outcome settlement geometry. For a categorical payoff, the covariance $C(p) = \mathrm{diag}(p) - pp^{\mathsf{T}}$ has the complete-set direction $\mathbf{1}$ in its null space. This yields a conditional dimension reduction and shows why diagonal or Euclidean inventory penalties can give the wrong ordering of portfolios.

4. It supplies a reproducible finite-state solver and a controlled numerical design. A separate diffusion-limit experiment reproduces the closest prediction-market stochastic-control benchmark and verifies the nested passive specialization. The remaining experiments isolate the value of flow information, impulse execution, partial observation, duration specification, and behavior under heavy-jump misspecification. These are numerical experiments under specified data-generating processes, not estimates of live-market profitability.

The contribution lies in the joint control problem. Prediction-market stochastic control has recently been studied with continuous probability diffusions and passive bid–ask choice [Feil and

Nendel, 2026]. Partial-information market making with hidden fill intensities is established for conventional securities [Campi and Zabaljauregui, 2020]; persistent flow has been handled with Hawkes processes [Jusselin, 2021]; and mixed limit/market-order control is known to produce QVIs [Guilbaud and Pham, 2013]. None of these components alone addresses a probability-simplex price, flow-dependent information jumps, controlled filtering through fills, categorical settlement covariance, and impulse inventory reduction in the same model.

The numerical results are deliberately diagnostic. In the diffusion-only restriction, the implementation closely reproduces the optimal-policy risk and inventory results reported by Feil and Nendel [2026]; agreement is the relevant target because the admissible passive-control problem is the same. In the leakage data-generating process, flow conditioning improves the objective relative to the same QVI solved without flow information, but its contribution is smaller than the contribution of market-order inventory reduction. Under a null in which flow is uninformative, the flow-aware controller loses a small amount because it reacts to noise. Under scheduled news, the same controller is also slightly inferior to the no-flow specification because the event is time driven rather than flow driven. These negative controls are important: the model does not make flow information valuable by construction. Under non-geometric durations, augmenting the hidden state with regime age improves probability calibration and halves the quiet-state false-positive rate, but also lowers the fraction of leakage periods classified above a fixed 0.5 threshold. Under heavy jumps outside the controller's model, the filtered QVI remains materially better than passive and myopic policies, without implying robustness to arbitrary misspecification.

The rest of the paper proceeds as follows. Section 2 positions the model in the relevant literature. Section 3 specifies the probability, flow, regime, fill, and settlement processes. Section 4 derives the controlled filter. Section 5 gives the HJB–QVI and its characterization. Section 6 treats complete-set geometry. Section 7 describes the finite-state approximation, Section 8 establishes the nested diffusion benchmark, and Section 9 reports the synthetic experiments. Section 10 states the model's limits, and Section 11 concludes. Proofs and implementation details appear in the appendices.

Large language models (primarily OpenAI's ChatGPT) were used to assist with technical calculations, code development, and text refinement, while the author conducted the majority of the research, evaluated all outputs, and takes full responsibility for the final results.

2 Related literature

Prediction markets aggregate dispersed information into prices and have been studied as forecasting and mechanism-design institutions [Wolfers and Zitzewitz, 2004, Arrow et al., 2008]. The price of an event claim need not equal a physical probability when risk premia, collateral constraints, fees, or segmented participation are present. This paper therefore works under an explicitly chosen pricing measure $\mathbb{Q}$ and treats the quoted probability as a martingale price, not as an unrestricted claim about population beliefs.

Two recent strands are especially close. Dalen [2025] specifies a logit jump diffusion whose drift is fixed by the requirement that the probability price be a $\mathbb{Q}$-martingale. Feil and Nendel [2026] formulates optimal passive market making for a binary prediction contract, derives an HJB equation, proves a classical-solution result for their transformed diffusion setting, and quantifies the roles of inventory, belief, time to resolution, and risk aversion. Xi et al. [2026] develops a structural volatility model with deadline-resolution and informed-order-flow channels and evaluates it on Kalshi contracts. Information-based asset-pricing models also provide a general route from signal revelation to conditional-probability prices [Brody et al., 2008]. The model below retains the martingale restriction but focuses on latent leakage, marked jumps, and quote-dependent learning.

Latent-state methods have also been applied directly to bounded prediction-market prices. A groupwise Beta hidden Markov model uses regime-dependent Beta emissions and cross-contract structure to classify price behavior [Voigt, 2025]. The present formulation uses a finite hidden regime to represent variation in jump and fill laws, and phase augmentation to accommodate non-geometric durations while preserving a finite-dimensional filter. These objects are embedded in an explicit observation filtration and control problem.

The conventional market-making literature begins with dealer inventory and transaction uncertainty [Ho and Stoll, 1981] and includes tractable intensity-based formulations [Avellaneda and Stoikov, 2008, Guéant et al., 2013, Guéant, 2016]. The adverse-selection channel goes back to models in which order flow conveys information about value [Glosten and Milgrom, 1985]. Multi-asset extensions show that covariance structure, rather than separate marginal variances, determines inventory pressure [Bergault et al., 2021]. The categorical covariance in Section 6 is a prediction-market specialization with an exact null direction.

Partial information and persistent flow are established control topics. Campi and Zabaljauregui [2020] filter hidden Markov fill intensities and characterize a reduced value function by a viscosity HJB equation. Jusselin [2021] models persistent order flow by Hawkes processes and gives a consistent numerical method. The current formulation differs in two linked respects: the hidden regime governs both public probability jumps and executions, and the maker's quote changes the fill likelihood used in Bayes' rule. The filter therefore cannot be estimated once and passed to a controller without retaining the control in the observation law.

Finally, impulse control is required when a maker can cross the spread for immediate execution. Guilbaud and Pham [2013] formulate limit and market orders as regular and impulse controls and obtain a QVI. Standard references cover jump-diffusion and impulse-control methods [Pham, 2009, Øksendal and Sulem, 2007, Cont and Tankov, 2004]. Here the intervention operator is coupled to a partially observed probability-jump state and to terminal categorical risk.

3 Market and information model

3.1 Event payoffs and probability prices

Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{Q})$ satisfy the usual conditions. There are $K \geq 2$ mutually exclusive and exhaustive outcomes. The settlement payoff is

$$Y = e_J \in \{e_1, \dots, e_K\},$$

where e_k is the kth coordinate vector. Trading takes place on $[0, T]$ and resolution occurs immediately after the trading horizon. The public price vector is

$$P_t = \mathbb{E}^{\mathbb{Q}}[Y \mid \mathcal{F}_t^{\text{pub}}] \in \Delta^{K-1}, \qquad \Delta^{K-1} = \{p \in [0,1]^K : \mathbf{1}^\mathsf{T} p = 1\}. \tag{1}$$

Thus P is a bounded $\mathbb{Q}$-martingale. The measure $\mathbb{Q}$ absorbs any pricing wedge relevant to the traded contracts. No claim that P_t equals a representative agent's physical belief is needed.

Let $S_t \in \mathcal{S} = \{1, \dots, m\}$ be an unobserved information state. A two-state interpretation is quiet versus leakage. More states can represent phases within a duration distribution, a distinction used in Section 4. The maker does not observe S. The regime carries information about the timing and magnitude of repricing, not its direction; we impose $\mathbb{E}^{\mathbb{Q}}[Y \mid \mathcal{G}_t] = P_t$ for the maker's observation filtration $\mathcal{G}_t$. If private fills also reveal directional information about settlement, the maker's posterior outcome probability must replace P_t as an additional controlled state.

We model the public price as

$$dP_t = \Sigma(P_{t-}, H_{t-}, t)\, dW_t + \int_{\mathcal{Z}(P_{t-})} z\, \widetilde{\mu}^P(dt, dz), \tag{2}$$

where W is a Brownian motion, H is an observed flow state defined below, and

$$\widetilde{\mu}^P(dt, dz) = \mu^P(dt, dz) - \nu_{S_{t-}}(P_{t-}, H_{t-}, t, dz)\, dt.$$

The regime-dependent compensator separates public-news intensity from leakage intensity:

$$\nu_s(p, h, t, dz) = \left\{\lambda_s^{\text{exo}}(p, t) + \lambda_s^{\text{endo}}(p, A(h), |I(h)|)\right\} F_s(p, h, t, dz). \tag{3}$$

The first term accommodates scheduled or background news. The second is increasing in selected specifications of activity A and imbalance magnitude $|I|$. The jump mark law has support

$$\mathcal{Z}(p) = \{z \in \mathbb{R}^K : p + z \in \Delta^{K-1},\ \mathbf{1}^\mathsf{T} z = 0\}.$$

For the core filtering result, Σ is common across hidden states. Otherwise continuous observation of quadratic variation can reveal S immediately when regime volatilities differ. Regime information enters the observable jump and point-process laws instead.

Assumption 1 (Simplex viability and integrability)**.** *For every (p, h, t), $\mathbf{1}^\mathsf{T}\Sigma(p, h, t) = 0$. If $p_k = 0$, then $e_k^\mathsf{T}\Sigma(p, h, t) = 0$. Each $F_s(p, h, t, \cdot)$ is supported on $\mathcal{Z}(p)$, has finite second moment, and satisfies $\int z\, F_s(p, h, t, dz) = 0$. The coefficients are predictable and the stochastic integrals in (2) are well defined.*

Proposition 1 (Simplex martingale)**.** *Under Assumption 1, a solution of (2) starting in Δ^{K-1} remains in Δ^{K-1}. If it is nonexplosive, it is a true $\mathbb{Q}$-martingale.*

The result separates economic and geometric restrictions. The compensator makes each coordinate a local martingale; boundedness upgrades it to a martingale. Tangency, boundary degeneracy, and admissible jump support keep the process in the simplex. A direct proof is in Section A.

For $K = 2$, write $p = S(x) = (1 + e^{-x})^{-1}$. A logit jump diffusion

$$dX_t = \mu_X(X_{t-}, H_{t-}, t)dt + \sigma_X(X_{t-}, H_{t-}, t)dW_t + \int_{\mathbb{R}} z\, \widetilde{\mu}^X(dt, dz) \tag{4}$$

induces a martingale probability precisely when its predictable drift satisfies

$$\mu_X(x, h, t) = -\frac{\frac{1}{2}S''(x)\sigma_X^2(x, h, t) + \int_{\mathbb{R}}\{S(x+z) - S(x) - S'(x)z\}\nu_X(x, h, t, dz)}{S'(x)}. \tag{5}$$

This is the risk-neutral logit restriction used by Dalen [2025]. The direct simplex representation

(2) is more convenient for mutually exclusive outcomes; (5) shows that the binary formulation is consistent with it.

3.2 Observed order flow

Let $\mu^O(dt, d\zeta)$ be the marked point measure of public market orders, with marks $\zeta = (\varepsilon, v)$ containing sign $\varepsilon \in \{-1, +1\}$ and a nonnegative size or impact statistic v. Its regime-dependent compensator is

$$\rho_s(p, h, t, d\zeta)dt.$$

The maker compresses recent order flow into exponentially decaying states

$$dH_t^{\pm} = -\kappa H_t^{\pm} dt + \int_{\{\varepsilon=\pm 1\}} w(v)\, \mu^O(dt, d\varepsilon, dv), \tag{6}$$

and bounded transforms

$$I(H_t) = \frac{H_t^+ - H_t^-}{\epsilon + H_t^+ + H_t^-}, \qquad A(H_t) = \frac{H_t^+ + H_t^-}{c + H_t^+ + H_t^-}. \tag{7}$$

The bounded transforms prevent a single scale choice from making jump intensities unbounded. A logit-impact weighting $w(v)$ is one possible specification near probability boundaries; the theory requires only predictable, integrable marks and bounded intensities after transformation.

3.3 Limit-order fills and adverse selection

For outcome k, the maker posts a bid $b_t^k = P_{t-}^k - \delta_t^{b,k}$ and ask $a_t^k = P_{t-}^k + \delta_t^{a,k}$. A side can be withdrawn, denoted by $\delta = \infty$. Conditional on $S_{t-} = s$, maker fills are counting processes $M^{b,k}$ and $M^{a,k}$ with intensities

$$\Lambda_s^{b,k}(P_{t-}, H_{t-}, \delta_t^{b,k}), \qquad \Lambda_s^{a,k}(P_{t-}, H_{t-}, \delta_t^{a,k}). \tag{8}$$

The intensities decrease with the quote offset and may increase with activity. Adverse selection can be represented either by regime-dependent intensities or by a marked joint law for a fill and the next public price jump. The numerical model uses the latter: a bid fill is more likely when the next probability move is down, and an ask fill is more likely when it is up.

For a complete set, quote vectors must exclude static arbitrage:

$$\sum_{k=1}^{K} b_t^k \leq 1, \qquad \sum_{k=1}^{K} a_t^k \geq 1. \tag{9}$$

Selling one share of every outcome pays one unit at settlement, so aggregate bids above one admit an immediate arbitrage against the maker. Aggregate asks below one create the corresponding buy-side arbitrage.

3.4 Inventory, cash, and settlement risk

Let $Q_t \in \mathbb{Z}^K$ be inventory and X_t cash. Unit maker fills change the state by

$$\begin{aligned}\text{bid fill in } k: &\quad (X, Q) \mapsto (X - b^k, Q + e_k),\\ \text{ask fill in } k: &\quad (X, Q) \mapsto (X + a^k, Q - e_k).\end{aligned}$$

At intervention times τ_n, a market order $\xi_n \in \mathbb{Z}^K$ gives

$$Q_{\tau_n} = Q_{\tau_n -} + \xi_n, \qquad X_{\tau_n} = X_{\tau_n -} - P_{\tau_n -}^{\mathsf{T}} \xi_n - c(\xi_n; P_{\tau_n -}, H_{\tau_n -}), \tag{10}$$

where c includes half-spread, fees, and impact. In the numerical implementation, interventions can only reduce the magnitude of each binary inventory position; they cannot initiate directional risk.

The predictable covariance rate of the public probability is

$$\mathcal{A}_s(p, h, t) = \Sigma\Sigma^{\mathsf{T}}(p, h, t) + \int zz^{\mathsf{T}} \nu_s(p, h, t, dz), \qquad \overline{\mathcal{A}}(p, h, \pi, t) = \sum_s \pi_s \mathcal{A}_s(p, h, t). \tag{11}$$

The maker uses a mean–variance control criterion

$$\begin{aligned} J^u(t, x, p, h, \pi, q) = \mathbb{E}^u_{t,x,p,h,\pi,q}\Bigg[& X_T + Q_T^{\mathsf{T}} P_T - \frac{1}{2}\int_t^T \gamma(r) Q_r^{\mathsf{T}} \overline{\mathcal{A}}(P_r, H_r, \pi_r, r) Q_r \, dr \\ & - \frac{\gamma_T}{2} Q_T^{\mathsf{T}} C(P_T) Q_T \Bigg], \end{aligned} \tag{12}$$

where $C(p) = \mathrm{diag}(p) - pp^{\mathsf{T}}$. The running term penalizes local marked-price variation. The terminal term is the exact conditional covariance of categorical settlement, scaled by γ_T. Transaction costs from impulses already enter X. The value function is $V = \sup_u J^u$ over admissible quote, withdrawal, and impulse policies.

4 Controlled filtering of the information regime

The maker's observation filtration $\mathbb{G} = (\mathcal{G}_t)$ contains P, public order-flow marks, its own fill processes, and its past controls. Define

$$\pi_t^s = \mathbb{Q}(S_t = s \mid \mathcal{G}_t), \qquad s \in \mathcal{S}. \tag{13}$$

Because own executions are observations, their likelihood must be evaluated at the posted quote. Suppose an observed event of type r with mark y has state-dependent predictable density or intensity $\ell_s^r(y; P_{t-}, H_{t-}, \delta_{t-})$. At the event time,

$$\pi_t^s = \frac{\pi_{t-}^s \ell_s^r(y; P_{t-}, H_{t-}, \delta_{t-})}{\sum_j \pi_{t-}^j \ell_j^r(y; P_{t-}, H_{t-}, \delta_{t-})}. \tag{14}$$

For a maker bid fill, ℓ_s^r is $\Lambda_s^{b,k}(p, h, \delta^{b,k})$; for survival without a fill, the absence of an event contributes the usual intensity-compensator term. If $G = (g_{ij})$ is the generator of S, the between-event posterior dynamics contain

$$d\pi_t^s = \sum_j \pi_t^j g_{js}\, dt - \pi_t^s \left(\ell_s^0(\delta_t) - \sum_j \pi_t^j \ell_j^0(\delta_t) \right) dt + \text{innovation terms}, \tag{15}$$

where ℓ_s^0 is total observed-event intensity and the innovation terms include compensated observed point measures. Standard finite-state filtering gives the complete expression [Elliott et al., 1995, Cappé et al., 2005]. Equation (14) is the part that matters most for control: changing the quote changes both execution revenue and posterior evolution.

4.1 Duration dependence through phase augmentation

A two-state continuous-time Markov chain implies exponential holding times; its discrete-time analogue implies geometric durations. Leakage episodes may instead have increasing or decreasing exit hazards. A finite-dimensional solution is to replace each economic regime by phases. Let (s, d) denote regime s and duration phase $d \in \{1, \ldots, D_s\}$. Transitions within a phase block and exits to another block define a finite generator. The resulting holding time is phase type, a dense class on the positive half-line [Asmussen et al., 1996]. The posterior remains finite dimensional, now over $\sum_s D_s$ states. A hidden semi-Markov model with a finite duration cap is the discrete counterpart.

This augmentation is substantive rather than cosmetic. Two observations with the same current leakage probability can imply different next-period hazards if one posterior concentrates on early phases and the other on late phases. The numerical diagnostic in Section 9.4 compares a constant-hazard filter with an age-augmented filter under a non-geometric data-generating process.

Assumption 2 (Filtering regularity)**.** *The hidden state space is finite. State transition rates and observed-event intensities are bounded. Observation densities are measurable and, on active quote sets, their mixture likelihood is strictly positive. The common continuous covariance of P is known to the controller; hidden-state dependence enters through observed jump and point-process characteristics. Controls are predictable with respect to $\mathbb{G}$.*

Theorem 1 (Finite-dimensional separation)**.** *Under Assumption 2, the posterior π obeys a finite-dimensional controlled filtering equation. For every admissible history-dependent policy there is a policy Markov in $(t, P_t, H_t, \pi_t, Q_t)$ with the same conditional law of future rewards. Hence the partially observed problem is equivalent to a fully observed control problem on the separated state.*

The theorem does not assert certainty equivalence. The separated state is sufficient, but the control still affects the generator of π. A policy obtained by filtering under one fixed quote and then optimizing as if the filter were exogenous generally solves a different problem.

5 Hybrid stochastic control and the HJB–QVI

Cash translation gives

$$V(t, x, p, h, \pi, q) = x + v(t, p, h, \pi, q). \tag{16}$$

Let $z = (p, h, \pi)$ denote the continuous separated state. For a regular quote control $\delta \in \mathcal{D}(z, q)$, let $\mathcal{A}^\delta$ be the controlled generator of (P, H, π), including exogenous observed jumps and posterior updates. Maker-fill operators must update cash, inventory, and belief simultaneously. For example,

$$\mathcal{F}^{\delta}_{b,k}v(t,z,q) = \overline{\Lambda}^{b,k}(z,\delta)\big[v(t,z^{+}_{b,k}(\delta),q+e_k) - v(t,z,q) \\ - (p_k - \delta^{b,k})\big], \tag{17}$$

where $\overline{\Lambda}^{b,k} = \sum_s \pi_s \Lambda^{b,k}_s$ and $z^{+}_{b,k}$ replaces π by the Bayes update conditional on that fill. The ask operator has inventory $q - e_k$ and cash increment $p_k + \delta^{a,k}$.

Define the regular-control Hamiltonian

$$\mathcal{H}v(t,z,q) = \sup_{\delta\in\mathcal{D}(z,q)} \left\{ \partial_t v + \mathcal{A}^{\delta}v + \sum_{k=1}^{K} \left(\mathcal{F}^{\delta}_{b,k}v + \mathcal{F}^{\delta}_{a,k}v\right) - \frac{\gamma(t)}{2} q^{\mathsf{T}}\overline{\mathcal{A}}(z,t)q \right\}. \tag{18}$$

The admissible quote set enforces inventory bounds and (9). Withdrawal is an ordinary action with zero fill intensity.

For an impulse ξ in a finite admissible set $\mathcal{I}(z,q)$, define

$$\mathcal{M}v(t,z,q) = \sup_{\xi\in\mathcal{I}(z,q)} \left\{ v(t,z,q+\xi) - p^{\mathsf{T}}\xi - c(\xi;p,h) \right\}. \tag{19}$$

Zero-size impulses are excluded. The value function solves

$$\max\left\{\mathcal{H}v(t,z,q),\ \mathcal{M}v(t,z,q) - v(t,z,q)\right\} = 0, \tag{20}$$

with terminal condition

$$v(T,p,h,\pi,q) = q^{\mathsf{T}}p - \frac{\gamma_T}{2} q^{\mathsf{T}}C(p)q. \tag{21}$$

The continuation region has $v > \mathcal{M}v$ and satisfies the controlled integro-HJB equation. In the intervention region, an immediate inventory-reducing market order is optimal. Pulling a quote is distinct from intervening: withdrawal avoids future fills but does not change current inventory.

Assumption 3 (Control problem)**.** *The probability and belief simplexes and the transformed flow domain are compact; inventory is bounded. Regular action sets are nonempty and compact. Coefficients, rewards, jump kernels, and observation likelihoods are bounded and uniformly continuous in the continuous state, with the Lipschitz conditions required for existence of the controlled process. The simplex viability conditions hold. The impulse set is finite at each inventory, $c(\xi) \geq c_0 > 0$ for $\xi \neq 0$, and intervention cycles have strictly positive total cost.*

Theorem 2 (Dynamic programming and QVI characterization)**.** *Under Assumptions 2 and 3, the separated value function satisfies the dynamic programming principle and is a bounded continuous viscosity solution of (20)–(21). On compact interior state domains with the state-constraint boundary condition induced by simplex viability, the comparison principle holds for the associated proper nonlocal QVI; consequently the viscosity solution is unique. If a sufficiently smooth solution exists and measurable maximizing selectors are admissible, the selectors and the intervention rule $v = \mathcal{M}v$ are optimal.*

The proof follows the controlled jump-diffusion and impulse-control arguments in Pham [2009] and Øksendal and Sulem [2007], with the filtered state included in the controlled generator. The positive intervention cost prevents accumulation of costless impulses. Nonlocal comparison uses the same doubling-of-variables structure as viscosity results for integro-differential equations [Barles and Imbert, 2008]. Section A records the steps and the points at which the assumptions are used. Unlike the diffusion-only result of Feil and Nendel [2026], the natural regularity class here is viscosity rather than classical because public jumps, posterior jumps, withdrawal boundaries, and interventions create nonsmooth value functions.

6 Complete-set settlement geometry

For $Y = e_J$ with $\mathbb{Q}(J = k \mid P = p) = p_k$,

$$\mathrm{Cov}(Y \mid P = p) = C(p) = \mathrm{diag}(p) - pp^\mathsf{T}. \tag{22}$$

Thus the conditional variance of inventory settlement is

$$\mathrm{Var}(q^\mathsf{T} Y \mid P = p) = q^\mathsf{T} C(p) q = \sum_{k=1}^{K} p_k \left(q_k - p^\mathsf{T} q\right)^2. \tag{23}$$

Proposition 2 (Complete-set invariance)**.** *For every* $p \in \Delta^{K-1}$*,* $C(p)\mathbf{1} = 0$*. Consequently,*

$$(q + c\mathbf{1})^\mathsf{T} C(p)(q + c\mathbf{1}) = q^\mathsf{T} C(p) q \quad \text{for every } c \in \mathbb{R}.$$

If p *is in the relative interior of the simplex, the null space of* $C(p)$ *is exactly* $\mathrm{span}\{\mathbf{1}\}$*.*

Equal inventory in every outcome is a riskless complete set: it pays the same amount in every state. A diagonal penalty $\sum_k p_k(1 - p_k)q_k^2$ ignores the negative cross-covariances $-p_k p_j$ and incorrectly penalizes this direction. The same issue affects Euclidean penalties.

Let $B \in \mathbb{R}^{K\times(K-1)}$ have columns $e_k - e_K$, $k < K$. Every inventory can be written $q = B\widetilde{q} + c\mathbf{1}$. If trading costs, constraints, and rewards are invariant to adding a complete set at its no-arbitrage price, the risky control problem can be expressed in $\widetilde{q} \in \mathbb{R}^{K-1}$. Frictions can break exact control invariance even though settlement-risk invariance remains exact. Accordingly, the dimension reduction is conditional on the complete-set execution convention; the covariance identity is unconditional.

7 Finite-state numerical method

The experiments use a deliberately transparent discrete model rather than an estimated continuous-time calibration. It preserves the information structure and the hybrid action set while allowing exact backward induction.

7.1 State, observations, and actions

There are $H = 24$ decision dates. The binary public probability lies on

$$\mathcal{P} = \{0, 0.05, \ldots, 0.95, 1\},$$

inventory lies in $\{-4, \ldots, 4\}$, observed activity lies in $\{0, 1, 2\}$, and the leakage posterior is discretized on 21 equally spaced points. The hidden state is quiet ($s = 0$) or leakage ($s = 1$), with transition matrix

$$G = \begin{pmatrix} 0.94 & 0.06 \\ 0.30 & 0.70 \end{pmatrix}. \tag{24}$$

The initial public probability is 0.5, initial leakage probability is 0.1, and initial inventory is zero.

At each date the maker chooses a bid level and ask level from $\{\text{off}, \text{tight}, \text{wide}\}$, giving nine regular actions. Tight and wide offsets are 0.0125 and 0.025. The one-step observation is

$$O_n = (F_n, D_n, B_n, A_n) \in \{0,1\} \times \{-1,0,1\} \times \{0,1\}^2, \tag{25}$$

where F_n is high external flow, D_n is the sign of the public probability move, and B_n, A_n are bid and ask fills. Its state-conditional likelihood factors as

$$L_s(o \mid p, a, \delta) = g_s(f)\, m_s(d \mid p, a)\ \mathrm{Bern}(b; \phi_s^b)\ \mathrm{Bern}(r; \phi_s^a). \tag{26}$$

The factorization is conditional, not an assertion that fills are uninformative about price changes: ϕ_s^b and ϕ_s^a depend on d. A bid is more likely to fill before a down move and an ask before an up move, with a stronger asymmetry in leakage.

The nonzero move probability at p is

$$\chi_s(p, a) = 4p(1-p)\left(\chi_s^0 + \chi_s^A a/2\right), \tag{27}$$

clipped below one. Conditional on a nonzero move, up and down have equal probability and magnitude 0.05, so the public price is an exact martingale. The factor $4p(1-p)$ shuts down moves at the absorbing boundaries. High-flow probabilities are 0.12 in quiet and 0.72 in leakage. Baseline nonzero-move probabilities at $p = 0.5$ are 0.08 and 0.42, before the activity increment.

After observation o, the filter is

$$\widehat{\pi}_n(s) = \frac{\pi_n(s) L_s(o \mid p, a, \delta)}{\sum_j \pi_n(j) L_j(o \mid p, a, \delta)}, \qquad \pi_{n+1} = \widehat{\pi}_n G. \tag{28}$$

The next posterior value is linearly interpolated between belief-grid nodes in backward induction.

7.2 Bellman recursion

Let $V_n(p, a, \pi, q)$ be value before the intervention decision. A market order may move q only toward zero. If q' is the post-intervention inventory, its cash increment is

$$-p(q' - q) - h|q' - q| - \eta(q' - q)^2, \tag{29}$$

with $h = 0.018$ and $\eta = 0.0015$. Given q', a regular action maximizes expected quote cash flow plus continuation value minus

$$\frac{\gamma}{2}(q')^2\, \mathrm{Var}(\Delta P_n \mid p, a, \pi), \qquad \gamma = 8.$$

The terminal value is

$$V_H(p, a, \pi, q) = qp - \frac{\gamma_T}{2} p(1-p) q^2, \qquad \gamma_T = 0.35. \tag{30}$$

Backward induction enumerates the admissible impulse targets, nine quote actions, and all 24 observations. This is an exact solution of the stated finite Markov decision problem up to floating-point arithmetic and belief interpolation.

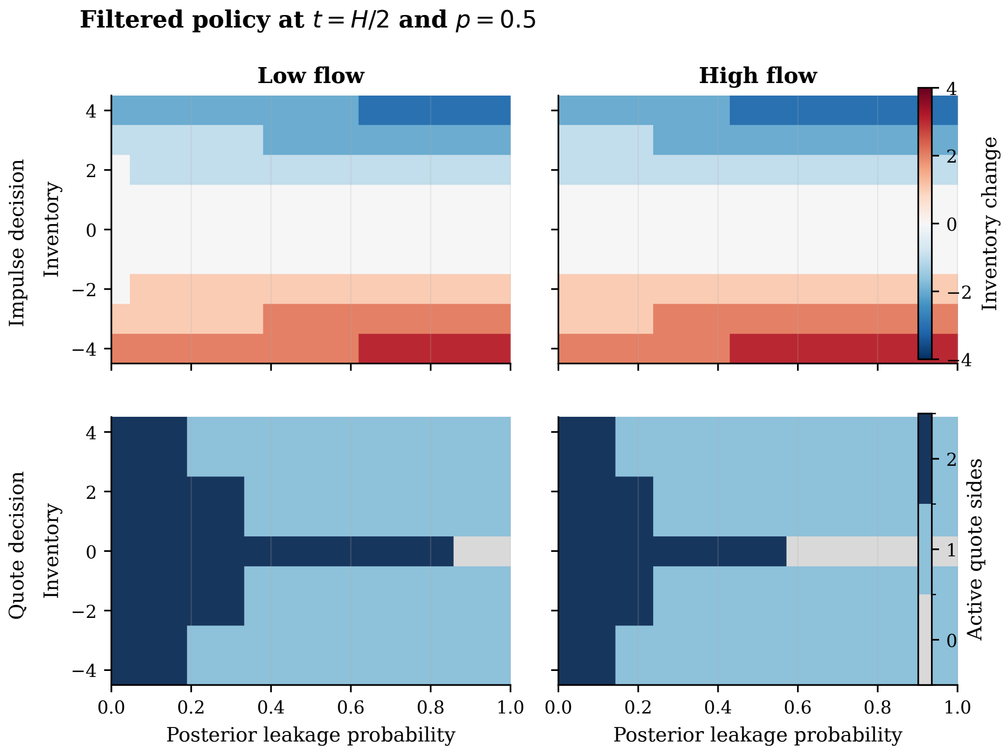

Figure 1: A slice of the filtered policy at the midpoint of the horizon and $p = 0.5$. The upper panels show the inventory change selected by the impulse decision; the lower panels show the number of active quote sides. Columns compare low and high observed activity. High leakage belief and nonzero inventory expand the intervention region and reduce two-sided quoting.

Figure 1 displays a policy slice. The impulse rule is monotone in inventory magnitude over most of the slice but changes discretely with posterior leakage probability. Quoting is not equivalent to a fixed posterior threshold: inventory, flow, and remaining horizon jointly determine whether zero, one, or two sides remain active.

7.3 Benchmarks

Eight policies are evaluated.

1. *Filtered QVI*: the partial-information policy with flow observations and impulses.
2. *Filtered passive*: the same filter and regular control, with market-order impulses disabled.
3. *No-flow QVI*: a QVI solved with flow made uninformative and the activity effect removed.
4. *Pooled passive*: a passive policy that pools the two regimes.
5. *Oracle QVI*: a policy that observes the true current regime.

6. *Fixed threshold*: two-sided tight quotes until leakage belief exceeds 0.68, then withdrawal and inventory reduction.
7. *Flow heuristic*: tight quotes at low activity, wide quotes at medium activity, and withdrawal at high activity.
8. *Myopic*: tight quotes with an inventory-bound skew but no dynamic inventory reduction.

The oracle is an information upper benchmark only when its transition and observation model match the data-generating process. In deliberately misspecified scenarios, it is simply a policy that sees the current regime while using the baseline model; it need not bound other policies.

8 External benchmark and nested diffusion limit

The closest direct benchmark is the binary passive-control problem of Feil and Nendel [2026]. This section serves two purposes. First, it verifies that the present framework contains their HJB as a restricted case. Second, it implements their published synthetic specification independently and tests whether the reported optimal-policy behavior can be recovered without using market data or tuning parameters to their output table.

To obtain the restriction, set $K = 2$ and represent inventory by the number q of claims on outcome one. Suppress the flow and posterior states, set the public jump measure to zero, exclude the intervention set, and permit only passive bid and ask quotes. Let

$$dp_t = \varsigma(t, p_t)\, dW_t$$

and let fills of size Δ arrive with intensities $\Lambda^b(t, p, \pi^b)$ and $\Lambda^a(t, p, \pi^a)$. Because the risk terms in (12)–(21) use a factor of one half, set $\gamma(t) = 2\gamma_{\mathrm{FN}}$ and $\gamma_T = 2\gamma_{T,\mathrm{FN}}$. After subtracting marked inventory qp from the cash-separated value, the QVI continuation equation reduces to

$$\begin{aligned} 0 = & -\partial_t V - \frac{1}{2}\varsigma(t,p)^2 \partial_{pp} V + \gamma_{\mathrm{FN}} q^2 \varsigma(t,p)^2 \\ & - \mathbf{1}_{\{q<Q\}} H^b\Big(t, p; \frac{V(t,p,q) - V(t,p,q+\Delta)}{\Delta}\Big) - \mathbf{1}_{\{q>-Q\}} H^a\Big(t, p; \frac{V(t,p,q) - V(t,p,q-\Delta)}{\Delta}\Big), \end{aligned} \tag{31}$$

with $V(T, p, q) = -\gamma_{T,\mathrm{FN}} q^2 p(1-p)$ and

$$H^\circ(t, p; z) = \Delta \sup_{\pi \in [0,1]} \Lambda^\circ(t, p, \pi)\big(G^\circ(p, \pi) - z\big), \quad G^b = p - \pi, \quad G^a = \pi - p.$$

This is the reduced HJB of Feil and Nendel [2026]. The nested policy should therefore agree with their optimum up to numerical approximation; exceeding it under the same state, objective, and passive action set would not constitute an economic improvement.

The reproduction uses their reported diffusion, intensities, and parameters: $\Delta = 10$, $Q = 100$, $\gamma_{\mathrm{FN}} = 4 \times 10^{-3}$, and $\gamma_{T,\mathrm{FN}} = 10^{-3}$. We solve the HJB by damped implicit Euler on a 400×201 time–price grid with 401 quotes. The 10,000 paths begin at $p_0 = 0.5$ and settle as Bernoulli(p_T). The source bounds only the optimal policy; we report its unbounded myopic convention and a common-bound variant.

Table 1: Replication and common-constraint comparison with Feil–Nendel

Source	Policy	Mean P&L	P&L s.d.	$\mathbb{E}\lvert q_T\rvert$	$\mathrm{VaR}_{0.05}$	$\mathrm{ES}_{0.05}$
Published	Myopic	12.47	28.11	49.37	32.41	40.50
Reproduction	Myopic, unbounded	12.55	34.66	57.05	44.52	61.13
Reproduction	Myopic, common bound	12.15	27.36	47.41	32.79	40.03
Published	Optimal passive	12.39	10.34	15.23	4.20	9.68
Reproduction	Restricted HJB	12.55	10.38	15.26	4.06	9.28

Published entries reproduce Table 2 of Feil and Nendel [2026]. Reproduction entries use their reported model and parameter values, 10,000 synthetic paths, 400 Euler steps, a 201-point price grid, and a 401-point HJB quote grid. VaR and ES are positive loss magnitudes. The published myopic strategy is described as unbounded; the common-bound row applies $|q| \leq 100$ to both policies. The Restricted HJB is the passive, diffusion-only specialization of the model in this paper.

The restricted HJB reproduces the published optimal row closely. Mean P&L is 12.55 versus 12.39, P&L standard deviation is 10.38 versus 10.34, mean absolute terminal inventory is 15.26 versus 15.23, and the two downside-risk measures differ by less than 0.42. The common-bound myopic reproduction also closely matches the published risk profile: its P&L standard deviation is 27.36 versus 28.11, mean absolute terminal inventory is 47.41 versus 49.37, VaR is 32.79 versus 32.41, and expected shortfall is 40.03 versus 40.50. Relative to that common-bound comparator, the restricted HJB reduces P&L standard deviation by 62.0%, terminal inventory by 67.8%, VaR by 87.6%, and expected shortfall by 76.8%, while its mean P&L is 0.41 higher.

The unbounded myopic implementation has similar mean P&L but more dispersion than the published row: standard deviation 34.66, mean absolute terminal inventory 57.05, and expected shortfall 61.13. It crosses the nominal inventory bound on 25.4% of paths at some point and finishes outside it on 13.9%. This discrepancy is reported rather than removed. It shows that the risk statistics of the myopic comparator are sensitive to the discrete execution and inventory convention, whereas the optimal-policy reproduction is stable. For the extension experiments below, every policy is therefore evaluated under a common inventory constraint and common random numbers.

The benchmark establishes nesting and numerical continuity with the closest existing prediction-market control model. It is not the source of the paper's incremental performance claim: when flow, public jumps, hidden states, and impulses are absent, the expanded framework deliberately collapses to the passive optimum. Incremental value is evaluated only after the corresponding information channels or controls are introduced and is accompanied by negative controls.

9 Controlled synthetic experiments

9.1 Design and estimands

No observed market data enter the experiment. All parameters in Table 2 are normalized inputs. Five data-generating processes are used:

1. *Quiet null*: flow, jump activity, fills, and adverse selection are identical across the nominal hidden states. Flow should have no information value.

2. *Leakage cascade*: the controller's two-state model is correctly specified. Leakage produces more high-flow observations, more public moves, and more adverse fills.

3. *Scheduled news*: flow is uninformative and move probability rises late in the horizon for exogenous time-based reasons.

4. *Duration asymmetry*: state-switching hazards depend on the age of the current regime, while the baseline controller uses constant transition probabilities.

5. *Heavy-jump stress*: the data-generating process permits two-grid moves and stronger adverse selection, while the controller is solved for one-grid moves.

Each scenario-policy cell contains 10,000 paths. Policies within a scenario use the same public states, settlement uniforms, and fill uniforms. Pairwise differences therefore use common random numbers. Reported standard errors for policy levels are ordinary Monte Carlo standard errors; ablation standard errors are computed from pathwise paired differences. Settlement is sampled as $Y \sim \text{Bernoulli}(P_H)$, preserving calibration of the synthetic martingale.

The primary performance measure is the realized finite-horizon analogue of (12):

$$\widehat{J} = X_H + q_H P_H - \frac{1}{2}\sum_{n=0}^{H-1} \gamma q_n^2 \operatorname{Var}(\Delta P_n \mid \text{DGP}) - \frac{\gamma_T}{2} P_H(1 - P_H) q_H^2. \tag{32}$$

Settlement P&L is $X_H + q_H Y$. Expected shortfall $\text{ES}_{0.05}$ is the mean settlement P&L among the worst five percent of paths. Objective values and P&L are in normalized contract-value units; they are not annualized returns, Sharpe ratios, or estimates of implementable dollar profit.

Table 2: Baseline numerical parameters

Parameter	Value
Decision steps	24
Probability-grid spacing	0.05
Inventory bound	± 4
Belief-grid points	21
Quiet-to-leakage probability	0.06
Leakage-to-quiet probability	0.3
High-flow probability, quiet/leakage	0.12 / 0.72
Move probability at $p = 0.5$, quiet/leakage	0.08 / 0.42
Tight/wide quote offsets	0.0125 / 0.025
Running risk coefficient	8.0
Terminal risk coefficient	0.35
Impulse half-spread	0.018
Impulse quadratic impact	0.0015

Parameters are normalized simulation inputs, not estimates from market data. Boundary attenuation multiplies move probabilities by $4p(1 - p)$.

9.2 Policy performance

Table 3 and figure 2 give the main policy levels. In the correctly specified leakage scenario, Filtered QVI has mean objective 0.0656 with standard error 0.0008. No-flow QVI has 0.0640, Filtered passive has 0.0589, and Myopic has -0.0464. The full-information QVI, not shown in the main table, has mean 0.0737. The paired Filtered-QVI advantage over Myopic is 0.1120 with standard error 0.0014.

The ranking is not driven only by the mean penalty. In leakage, Filtered QVI has worst-five-percent settlement P&L of −0.0995 and mean absolute terminal inventory 0.0011. The passive policy has −0.4811 and 0.2121, respectively; Myopic has −1.0420 and 1.1083. The filtered QVI quotes at least one side for 89.7% of all side-time opportunities, so its tail behavior is not obtained by permanent withdrawal.

The negative controls behave as intended. In the quiet null, No-flow QVI exceeds Filtered QVI by 0.00179 on a paired basis. In scheduled news, the no-flow policy exceeds it by 0.00063. A controller that interprets flow through the leakage model pays for reacting to an uninformative signal. Conversely, flow conditioning helps in the duration-asymmetry and heavy-jump scenarios, where flow still carries regime information despite misspecification.

The heavy-jump stress produces the lowest filtered-policy objective, 0.0306, and the most negative filtered expected shortfall, −0.2368. Even there, Filtered passive has expected shortfall −0.5191 and Myopic has −1.1304. This is evidence about one specified stress process. It does not establish a distribution-free robustness guarantee.

Table 3: Policy performance across synthetic scenarios

Scenario	Policy	Objective	Settlement P&L	$\mathrm{ES}_{0.05}$	$\mathbb{E}\lvert q_H\rvert$	Quote uptime
Quiet null	Filtered QVI	0.0937 (0.0008)	0.1062 (0.0008)	-0.0531	0.0003	0.925
	No-flow QVI	0.0955 (0.0008)	0.1079 (0.0008)	-0.0522	0.0002	0.944
	Filtered passive	0.0838 (0.0008)	0.1079 (0.0024)	-0.4689	0.2151	0.776
	Myopic	0.0082 (0.0014)	0.1326 (0.0066)	-1.0032	1.0989	0.867
Leakage cascade	Filtered QVI	0.0656 (0.0008)	0.0800 (0.0008)	-0.0995	0.0011	0.897
	No-flow QVI	0.0640 (0.0008)	0.0786 (0.0008)	-0.1076	0.0013	0.933
	Filtered passive	0.0589 (0.0008)	0.0831 (0.0024)	-0.4811	0.2121	0.757
	Myopic	-0.0464 (0.0016)	0.0773 (0.0067)	-1.0420	1.1083	0.866
Scheduled news	Filtered QVI	0.0780 (0.0009)	0.1008 (0.0010)	-0.1052	0.0023	0.909
	No-flow QVI	0.0787 (0.0009)	0.1015 (0.0009)	-0.0979	0.0017	0.910
	Filtered passive	0.0702 (0.0010)	0.1074 (0.0023)	-0.4532	0.1908	0.767
	Myopic	-0.0487 (0.0020)	0.1217 (0.0066)	-0.9990	1.1031	0.866
Duration asymmetry	Filtered QVI	0.0595 (0.0008)	0.0751 (0.0008)	-0.1054	0.0010	0.891
	No-flow QVI	0.0587 (0.0008)	0.0744 (0.0008)	-0.1116	0.0013	0.930
	Filtered passive	0.0525 (0.0009)	0.0770 (0.0024)	-0.4859	0.2035	0.754
	Myopic	-0.0570 (0.0016)	0.0767 (0.0066)	-1.0338	1.0999	0.867
Heavy-jump stress	Filtered QVI	0.0306 (0.0012)	0.0635 (0.0012)	-0.2368	0.0192	0.897
	No-flow QVI	0.0262 (0.0012)	0.0595 (0.0013)	-0.2561	0.0212	0.931
	Filtered passive	0.0211 (0.0012)	0.0737 (0.0025)	-0.5191	0.2298	0.762
	Myopic	-0.1463 (0.0024)	0.0569 (0.0067)	-1.1304	1.1235	0.864

Means are followed by Monte Carlo standard errors in parentheses. $\mathrm{ES}_{0.05}$ is the mean settlement P&L in the worst five percent of paths. Each scenario-policy cell contains 10,000 paths. Policies within a scenario use common random numbers.

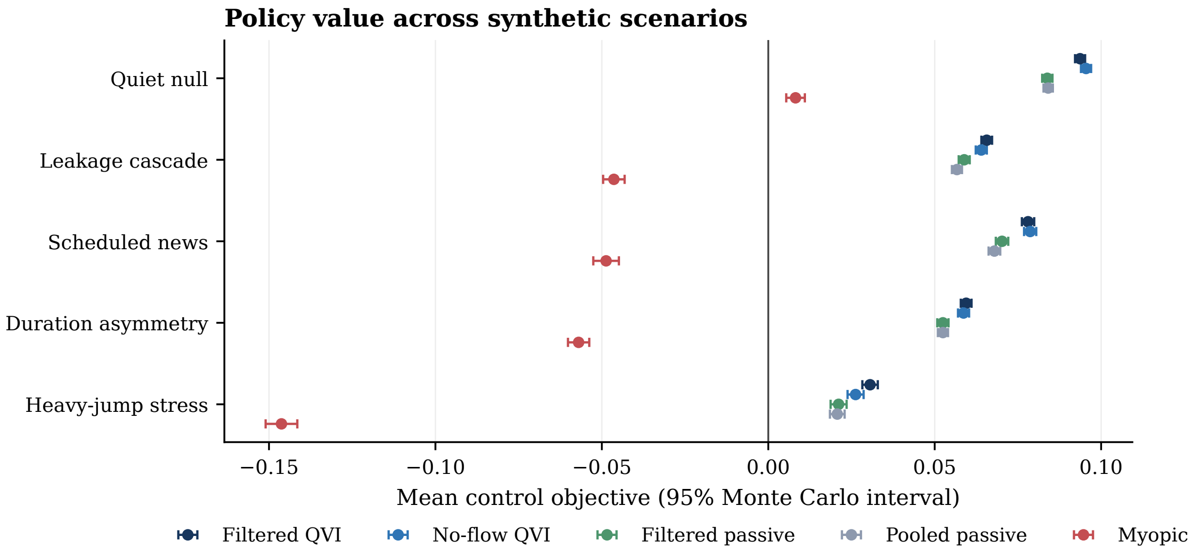

Figure 2: Mean control objective and 95% Monte Carlo intervals. The vertical line marks zero. The figure reports five representative policies; the fixed-threshold, flow-heuristic, and full-information policies are retained in the replication output and selected comparisons below.

9.3 Component values and information loss

Table 4: Paired increments in the control objective

Scenario	Flow information	Impulse option	Oracle gap
Quiet null	-0.0018 (0.0002)	0.0098 (0.0004)	0.0116 (0.0005)
Leakage cascade	0.0016 (0.0003)	0.0068 (0.0004)	-0.0081 (0.0005)
Scheduled news	-0.0006 (0.0003)	0.0078 (0.0005)	0.0124 (0.0006)
Duration asymmetry	0.0008 (0.0003)	0.0070 (0.0004)	-0.0096 (0.0006)
Heavy-jump stress	0.0044 (0.0005)	0.0095 (0.0006)	-0.0176 (0.0008)

Entries are paired mean differences with standard errors in parentheses. Flow information is Filtered QVI minus No-flow QVI; the impulse option is Filtered QVI minus Filtered passive; the oracle gap is Filtered QVI minus the full-information policy. The latter is an oracle bound only in the correctly specified leakage scenario.

In the leakage cascade, flow information adds 0.001634 (standard error 0.000326), the impulse option adds 0.006761 (0.000381), and the partial-information policy is 0.008098 (0.000550) below the correctly specified full-information policy. These quantities answer different questions. The first holds the impulse architecture fixed and changes the information set. The second holds the filter fixed and adds immediate inventory reduction. The third measures the value of observing the regime, not the value of filtering relative to ignoring regimes.

The impulse increment is positive in every data-generating process, ranging from 0.0068 to 0.0098. Flow information is state dependent: negative in the quiet null and scheduled-news scenario, positive in leakage, duration asymmetry, and the heavy-jump stress. The experiment therefore supports a selective claim. Flow is valuable when it predicts the transition law relevant to adverse fills and public jumps; it is not generically valuable merely because it is observed.

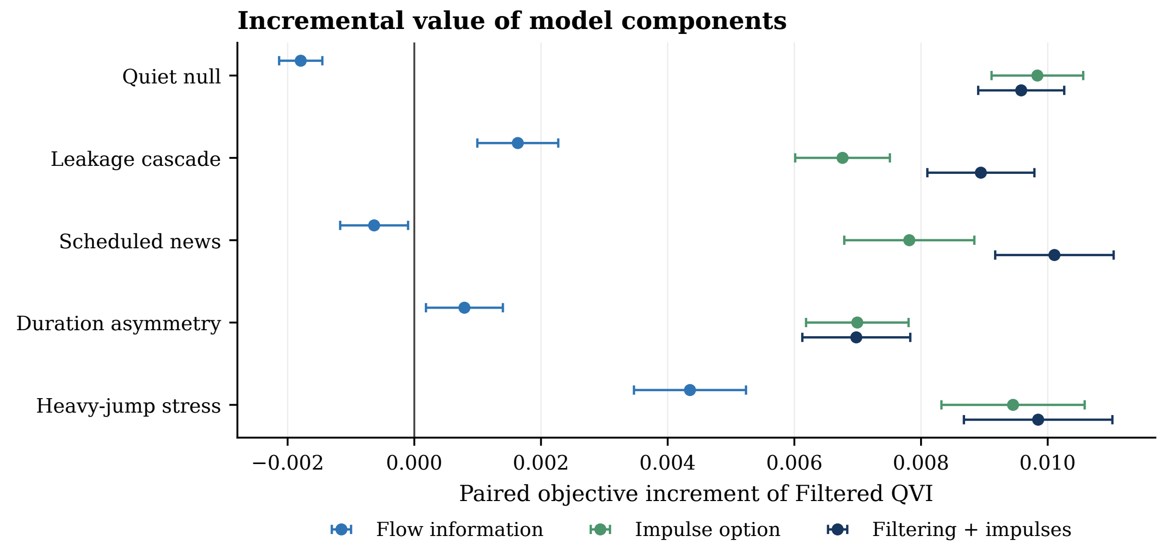

Figure 3: Paired objective increments with 95% Monte Carlo intervals. Flow information is Filtered QVI minus No-flow QVI. The impulse option is Filtered QVI minus Filtered passive. Filtering plus impulses is Filtered QVI minus Pooled passive.

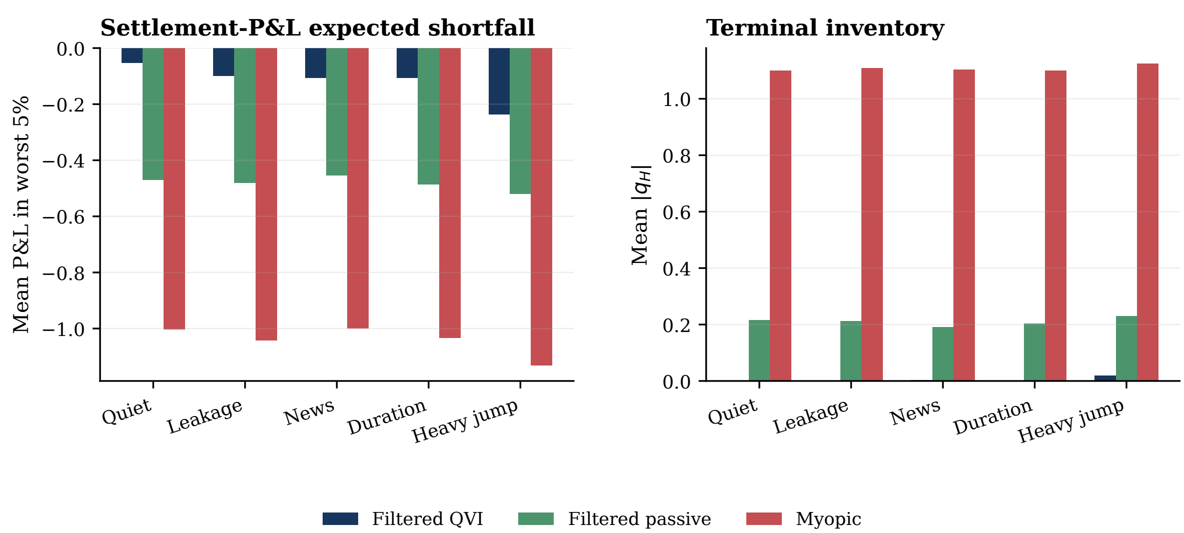

Figure 4: Settlement-P&L expected shortfall and mean absolute terminal inventory for three policies. The impulse-enabled filtered policy closes almost all inventory in the baseline scenarios; residual inventory rises under heavy-jump misspecification.

9.4 Duration-aware filtering

For the duration-asymmetry data-generating process, a separate diagnostic compares the baseline two-state Markov filter with an age-augmented filter. Both see the same high-flow indicator, public move, and fills generated by tight two-sided quotes. The duration-aware state contains 20 age phases per economic regime and uses the data-generating hazards. Scores use the prior leakage probability at each decision date.

Table 5: Filtering under non-geometric regime durations

Filter	Brier	Log score	False positive	True positive	Delay
Markov	0.1408	-0.4524	0.0595	0.2081	1.810
Duration-aware	0.1318	-0.4229	0.0318	0.1823	1.789

Scores use the prior leakage probability at each decision time. Delay is measured in simulation steps from entry into a leakage episode to the first posterior above 0.5, censored at episode end.

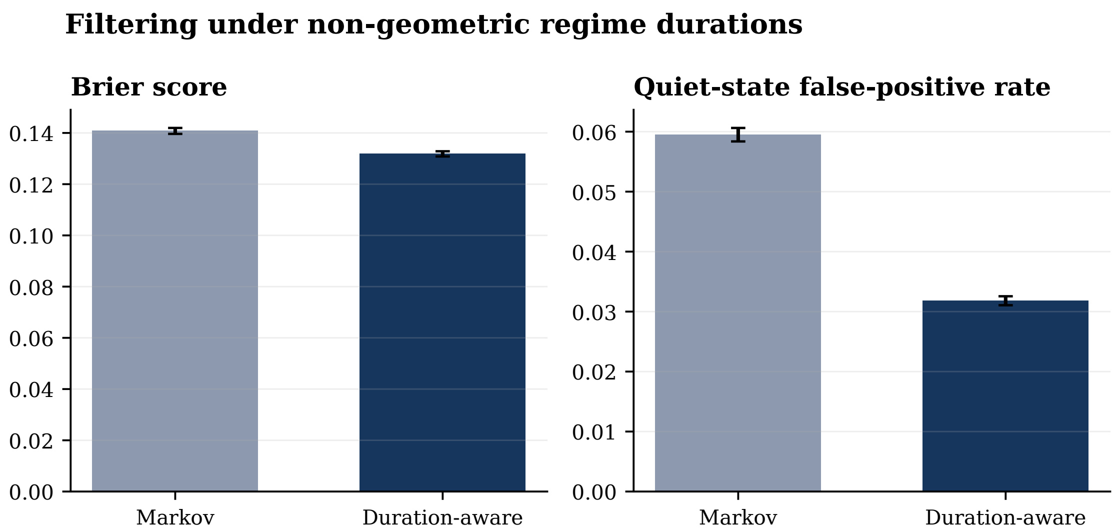

Figure 5: Filtering under non-geometric durations. Bars report path means and 95% Monte Carlo intervals. Lower Brier score and false-positive rate are better.

The duration-aware filter lowers the Brier score from 0.1408 to 0.1318 and improves mean log score from -0.4524 to -0.4229. Its quiet-state false-positive rate falls from 0.0595 to 0.0318. The gain is not uniform across classification metrics: the share of leakage observations above 0.5 falls from 0.2081 to 0.1823. Censored mean detection delay falls only from 1.810 to 1.789 steps, with median two for both filters. Duration augmentation improves calibration and suppresses false alarms in this design; it does not dominate at every threshold-based detection measure.

9.5 Multi-outcome settlement risk

The complete-set diagnostic draws 100,000 three-outcome probability vectors from a Dirichlet distribution with concentration vector $(0.8, 1.2, 1.5)$ and integer inventories from a bounded grid. It compares exact risk (23) with the diagonal approximation $\sum_k p_k(1-p_k)q_k^2$ and with the Euclidean norm. Pairwise misranking is the probability that two random inventories receive opposite risk orderings under the approximation and exact covariance.

Table 6: Three-outcome inventory-geometry diagnostic

Diagnostic	Value
Maximum complete-set shift error	6.39e-14
Maximum projection error	1.46e-14
Median relative error, diagonal heuristic	38.3%
90th percentile relative error, diagonal heuristic	511.2%
Pairwise misranking, diagonal heuristic	27.2%
Pairwise misranking, Euclidean norm	34.4%

The experiment uses 100,000 independent Dirichlet probability vectors and integer inventory vectors. Numerical errors are absolute. Misranking is the fraction of random inventory pairs whose risk ordering disagrees with the exact categorical covariance.

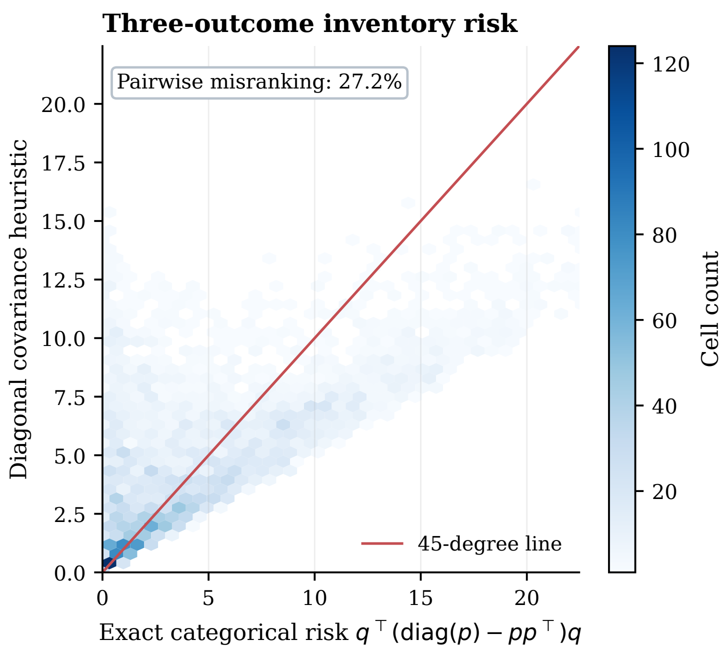

Figure 6: Exact categorical inventory risk versus the diagonal variance heuristic for 100,000 synthetic three-outcome states. Hexagon shading is cell count. The diagonal heuristic omits negative cross-covariances and misranks 27.2% of random inventory pairs.

Adding a complete set changes exact risk by at most 6.39×10^{-14} in floating-point arithmetic, and direct projection to the $K - 1$ risky coordinates changes it by at most 1.46×10^{-14}. The diagonal approximation has median relative error 38.3% and 90th-percentile relative error 511.2%; it misranks 27.2% of pairs. Euclidean inventory norm misranks 34.4%. The result is algebraic, not a fitted regularity: multi-outcome controllers should use categorical covariance or an equivalent reduced-coordinate representation.

9.6 Numerical verification

Four implementation invariants are checked automatically: the public transition is an exact conditional martingale at every probability-grid node; observation probabilities sum to one; every posterior remains in $[0, 1]$; and complete-set shifts leave categorical risk unchanged. In a short-horizon instance whose dynamic-programming initial value is 0.00293037, a 30,000-path simulation of the computed policy gives 0.00294455 with standard error 0.00009853. The difference is 0.144 Monte Carlo standard errors.

Belief-grid refinement holds the economic grid fixed. Initial values are 0.0169559, 0.0169242, and 0.0169159 on 11, 21, and 41 posterior points. The change from 21 to 41 points is 8.28×10^{-6}, or 0.049% of the fine-grid value. These checks do not prove convergence of the particular parameterized experiment to the continuous-time model. They verify probability conservation, policy reproduction, and stability with respect to the principal interpolation grid.

For the external diffusion benchmark, the computed initial value changes from 12.2187 on a $(160, 101, 201)$ time–price–quote grid to 12.2532 on the reported $(400, 201, 401)$ grid, a 0.28% change. The largest terminal fixed-point residual over the fine-grid backward solve is below 2.0×10^{-9}.

More generally, a locally consistent, monotone, stable discretization of (20) converges to its unique viscosity solution as the time and state meshes vanish, provided the intervention operator is discretized monotonically and comparison holds [Kushner and Dupuis, 2001, Barles and Souganidis, 1991]. The finite experiment reported here is solved exactly on its stated grid; it is not presented as a calibrated continuous-time limit.

10 Limitations and scope

The theoretical model is broader than the numerical instance, and the numerical instance is intentionally small. Six limitations delimit the results.

First, no live order-book data are used. Fill probabilities, jump frequencies, regime hazards, spreads, and costs are simulation inputs. The numerical magnitudes cannot support claims about market capacity, net profitability, or external validity. Empirical use would require timestamped quotes, trades, cancellations, queue position, fees, resolution rules, and a causal calibration protocol.

Second, the maker is small. Its quotes affect its own fill likelihood and therefore its observations, but not the public price, external order flow, or other agents' strategic behavior. A large-maker model would require price impact or a game among liquidity suppliers.

Third, the two-state controller is parsimonious. Public news, gradual leakage, manipulation, correlated contracts, and platform outages need not share the same latent law. Phase augmentation handles duration distributions but not arbitrary structural misspecification.

Fourth, the objective is a quadratic risk criterion. It makes local price covariance and terminal settlement covariance explicit, but it is not expected utility and does not impose a pathwise solvency constraint. Exponential utility, drawdown limits, collateral rules, and hard liquidation requirements can change the intervention region.

Fifth, the multi-outcome numerical study isolates risk geometry rather than solving the full K-outcome QVI. The $K-1$ reduction controls dimensional growth only when complete-set execution and costs preserve the relevant translation invariance. Cross-contract fills, logical dependence across separate markets, and joint collateral constraints require additional state variables.

Sixth, the external benchmark is a reproduction from the published model description, not a byte-for-byte execution of the authors' code. The restricted optimal-policy results are close, but the unbounded myopic risk statistics are higher than the published values and are sensitive to the discrete inventory convention. The common-bound comparison is therefore the appropriate controlled contrast within this paper; it should not be read as an audit of the original implementation.

These limits suggest a direct validation sequence: estimate a causal flow and jump model on a training period; evaluate filter calibration before policy returns; estimate quote-dependent fills with queue information; freeze the controller; and test out of sample with fees and latency. Synthetic stress tests should remain part of that process because observed samples contain few extreme information events.

11 Conclusion

Prediction-market making is a control problem under endogenous information risk. Order flow can forecast public probability jumps, and fills can be observations of the same hidden state that makes them adverse. Treating the filter as part of the controlled state leads to a coherent policy over bid and ask offsets, withdrawal, and immediate inventory reduction. The resulting HJB–QVI separates continuation and intervention decisions without separating inference from execution.

Four conclusions survive the controlled experiments. First, in the diffusion-only passive limit, the model recovers the optimal-policy risk and inventory behavior of the closest existing prediction-market control benchmark. Second, flow information has positive value only in data-generating processes where it predicts relevant state changes; it is costly under an uninformative-flow null and slightly costly around purely scheduled news. Third, the option to cross the spread and reduce inventory contributes more than flow conditioning in every scenario examined and materially reduces settlement-tail losses. Fourth, complete-set covariance is not a minor correction in multi-outcome markets: diagonal and Euclidean penalties frequently rank inventories incorrectly.

The paper's theoretical contribution is a standalone stochastic-control formulation that combines probability-simplex martingales, controlled latent-state filtering, flow-modulated jumps, categorical settlement risk, and regular/impulse liquidity decisions. Its numerical contribution is a reproducible diffusion-benchmark replication followed by synthetic diagnostics with negative controls and misspecification stresses. Matching the nested optimum validates the restriction; incremental value is attributed only to additional information or admissible actions under common evaluation conditions. Market-data estimation is a separate step and remains necessary before operational conclusions can be drawn.

A Proofs and theoretical details

A.1 Proof of the simplex-martingale proposition

Because $\mathbf{1}^\top\Sigma = 0$ and every jump mark satisfies $\mathbf{1}^\top z = 0$, applying $\mathbf{1}^\top$ to (2) gives $\mathbf{1}^\top P_t = \mathbf{1}^\top P_0 = 1$. At a boundary face $p_k = 0$, the continuous normal covariance vanishes by $e_k^\top\Sigma = 0$. Every jump maps p to $p + z \in \Delta^{K-1}$. The mean-zero mark condition removes a deterministic compensator drift that could point out of the face between jumps. Standard stochastic-invariance results, or a localization argument applied face by face, then give $P_t^k \geq 0$ for all k. Thus $P_t \in \Delta^{K-1}$.

The stochastic integrals in (2) have zero conditional mean after localization, so each coordinate is a local martingale. Since $0 \le P_t^k \le 1$, it is uniformly integrable and therefore a true martingale. □

A.2 Binary logit drift

Apply Itô's formula with jumps to $S(X_t)$. The predictable finite-variation coefficient is

$$S'(x)\mu_X(x,h,t) + \frac{1}{2}S''(x)\sigma_X^2(x,h,t) + \int \{S(x+z) - S(x) - S'(x)z\}\nu_X(x,h,t,dz).$$

It vanishes if and only if (5) holds. The remaining Brownian and compensated-jump terms are local martingales. Because $S(X) \in (0,1)$, the local martingale is a true martingale under nonexplosion. This also shows why setting the logit drift to zero generally violates the martingale condition in probability space.

A.3 Controlled filtering and separation

Fix an admissible predictable control. Conditional on the hidden state, the continuous public covariance is common and the observed marked point processes have bounded intensities. The Kallianpur–Striebel change-of-measure construction gives unnormalized state weights. Normalization yields (14); compensating the no-event likelihood and applying the hidden generator yields (15). All coefficients depend on the observation history only through the current observable state and the current predictable control, so (P, H, π, Q) is a controlled Markov process.

For any admissible history-dependent policy, regular conditional probabilities determine its continuation distribution given the separated state. Standard Markovization replaces the history-dependent kernel by a relaxed Markov control with the same state-action occupation law and reward. Compactness and continuity permit an ordinary measurable selector. Impulses are included by conditioning at intervention stopping times. This proves Theorem 1. The construction is the controlled-observation analogue of finite-state filtering and does not remove the dependence of π on past controls.

A.4 Dynamic programming and viscosity characterization

Under Assumption 3, the controlled state has a unique nonexplosive weak solution between impulses. Bounded inventory, compact continuous state, and $c_0 > 0$ exclude infinitely many profitable interventions in finite time. Concatenation of controls at stopping times gives the dynamic programming principle.

To obtain the subsolution property, touch the upper semicontinuous envelope of v from above with a smooth test function at a point in the continuation region. Apply the dynamic programming principle over a short interval before the first impulse and use Itô's formula for the controlled jump process. Optimizing over regular controls gives $\mathcal{H}v \le 0$ in the sign convention of (20); immediate intervention gives $\mathcal{M}v - v \le 0$. The supersolution argument uses a test function touching from below and ϵ-optimal controls. At points in the intervention region, the impulse inequality supplies the required branch. The terminal condition follows from the absence of further trading after T.

For comparison, double the continuous variables and use the finite inventory index componentwise. Properness of the Hamiltonian, bounded jump measures, uniform continuity of the posterior reset maps on positive-mixture-likelihood sets, and the positive impulse-cycle cost prevent a maximizing sequence from alternating indefinitely between intervention branches. The nonlocal terms are controlled by the integro-differential Jensen–Ishii argument. State-constraint test functions handle viable simplex faces. Comparison implies uniqueness. If v is smooth enough for Itô's formula and maximizing selectors exist, applying the formula along the selected controlled process and using the QVI inequalities proves verification.

A.5 Proof of complete-set invariance

Since $\mathbf{1}^\mathsf{T} p = 1$,

$$C(p)\mathbf{1} = \operatorname{diag}(p)\mathbf{1} - pp^\mathsf{T}\mathbf{1} = p - p = 0.$$

Symmetry gives $\mathbf{1}^\mathsf{T} C(p) = 0$, so expanding $(q + c\mathbf{1})^\mathsf{T} C(p)(q + c\mathbf{1})$ leaves $q^\mathsf{T} C(p) q$. For interior p, equation (23) is zero only if $q_k = p^\mathsf{T} q$ for all k, hence q is constant across outcomes. □

B Discrete transition and Bellman details

For regime s, public node i, activity a, and quote action $d = (d_b, d_a)$, the program enumerates $o = (f, m, b, r)$. Let $\alpha(p) = 4p(1-p)$ and

$$u_s(i, a) = \min\{0.92, \alpha(p_i)(u_s^0 + u_s^A a/2)\}.$$

Then

$$\mathbb{Q}(m = 0 \mid s, i, a) = 1 - u_s(i, a), \qquad \mathbb{Q}(m = \pm 1 \mid s, i, a) = \frac{1}{2} u_s(i, a).$$

The activity recursion is $a' = \min(2, a + 1)$ after high flow and $a' = \max(0, a - 1)$ otherwise.

Let $\bar{\phi}_{s,\ell}$ be the ordinary fill probability at quote level $\ell \in \{0, 1, 2\}$, with $\bar{\phi}_{s,0} = 0$. It is multiplied by an activity term and a selection term. For a bid, the selection term is $1 + \zeta_s$ before a down move, $\max(0.2, 1 - 0.32\zeta_s)$ before an up move, and one without a move. The ask rule reverses the signs. Probabilities are clipped at 0.88. Conditional independence of bid and ask Bernoulli draws is imposed after conditioning on (s, f, m).

Let R_n be the post-quote continuation operator, including cash flow and belief interpolation. In compact notation,

$$\begin{aligned} R_n(p, a, \pi, q; d) = \sum_o \bar{L}(o \mid p, a, \pi, d)\{&c_{\text{quote}}(p, d, o) \\ &+ V_{n+1}(p'(o), a'(o), \pi'(o), q'(o))\} - \frac{\gamma}{2} q^2 \sigma^2_{\Delta p}(p, a, \pi). \end{aligned}$$

The impulse-first Bellman equation is

$$\begin{aligned} V_n(p, a, \pi, q) = \max_{\widetilde{q} \in \mathcal{R}(q)} \Bigg\{ &- p(\widetilde{q} - q) - h|\widetilde{q} - q| - \eta(\widetilde{q} - q)^2 \\ &+ \max_{d \in \mathcal{D}(\widetilde{q})} R_n(p, a, \pi, \widetilde{q}; d) \Bigg\}, \end{aligned} \tag{33}$$

where $\mathcal{R}(q)$ is the set of integer inventories between q and zero. The passive policy fixes $\widetilde{q} = q$. The full-information policy replaces π by the observed state and sums over the next state using G.

C Scenario and policy details

The leakage cascade uses the parameters in Table 2. In the quiet null, high-flow probability is 0.22 in both states, nonzero-move probability at $p = 0.5$ is 0.13, the activity increment is zero, the fill multiplier is one, and the adverse-selection coefficient is 0.24 in both states. Scheduled news uses high-flow probability 0.20 in both states, baseline move probability 0.07, and move probability at least $0.58\,4p(1-p)$ from decision date 16 onward.

The duration-asymmetry process switches out of quiet with hazard

$$h_0(d) = \min(0.025 + 0.012d, 0.23)$$

and exits leakage with

$$h_1(d) = \min(0.10 + 0.13d, 0.88),$$

where d is current regime age. The heavy-jump stress changes high-flow probabilities to $(0.16, 0.66)$, baseline move probabilities to $(0.10, 0.46)$, activity increments to $(0.03, 0.20)$, and adverse-selection coefficients to $(0.20, 1.75)$. Conditional on a nonzero move and sufficient distance from the boundary, the move has size two grid points with probability 0.34.

The fixed-threshold and flow heuristics share an inventory skew: at inventory at least two, the bid is removed and the ask is tight; at inventory at most minus two, the ask is removed and the bid is tight. The myopic policy uses the same skew and otherwise posts tight quotes. These rules are intentionally simple; their role is to make the value of dynamic optimization visible, not to represent optimized industry baselines.

D Additional numerical tables

Table 7: Belief-grid refinement with fixed economic discretization

Grid	Posterior points	Initial value	Change from preceding grid
Coarse	11	0.0169559	—
Baseline	21	0.0169242	−0.0000317
Fine	41	0.0169159	−0.0000083

All rows use an eight-step horizon, inventory bound two, and probability-grid spacing 0.05. Only the posterior grid changes.

Table 8: Short-horizon policy-value reproduction

Quantity	Value
Backward-induction initial value	0.00293037
Monte Carlo mean	0.00294455
Monte Carlo standard error	0.00009853
Standardized difference	0.144
Simulation paths	30,000

References

Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz. The promise of prediction markets. *Science*, 320(5878):877–878, 2008. doi: 10.1126/science.1157679.

Søren Asmussen, Olle Nerman, and Marita Olsson. Fitting phase-type distributions via the EM algorithm. *Scandinavian Journal of Statistics*, 23(4):419–441, 1996.

Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008. doi: 10.1080/14697680701381228.

Guy Barles and Cyril Imbert. Second-order elliptic integro-differential equations: Viscosity solutions' theory revisited. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 25(3):567–585, 2008. doi: 10.1016/j.anihpc.2007.02.007.

Guy Barles and Panagiotis E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283, 1991.

Philippe Bergault, David Evangelista, Olivier Guéant, and Douglas Vieira. Closed-form approximations in multi-asset market making. *Applied Mathematical Finance*, 28(2):101–142, 2021. doi: 10.1080/1350486X.2021.1949359.

Dorje C. Brody, Lane P. Hughston, and Andrea Macrina. Information-based asset pricing. *International Journal of Theoretical and Applied Finance*, 11(1):107–142, 2008. doi: 10.1142/S0219024908004749.

Luciano Campi and Diego Zabaljauregui. Optimal market making under partial information with general intensities. *Applied Mathematical Finance*, 27(1–2):1–45, 2020. doi: 10.1080/1350486X.2020.1758587.

Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer, 2005.

Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007. doi: 10.5555/3020488.3020495.

Rama Cont and Peter Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2004.

Shaw Dalen. Toward black–scholes for prediction markets: A unified kernel and market-maker's handbook. *arXiv preprint arXiv:2510.15205*, 2025. URL https://arxiv.org/abs/2510.15205.

Robert J. Elliott, Lakhdar Aggoun, and John B. Moore. *Hidden Markov Models: Estimation and Control*. Springer, 1995.

Dominik Feil and Max Nendel. Optimal market making in prediction markets. *arXiv preprint arXiv:2607.17991*, 2026. URL https://arxiv.org/abs/2607.17991.

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

Olivier Guéant. *The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making*. CRC Press, 2016.

Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Dealing with the inventory risk: A solution to the market making problem. *Mathematics and Financial Economics*, 7(4): 477–507, 2013. doi: 10.1007/s11579-012-0087-0.

Fabien Guilbaud and Huyên Pham. Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1):79–94, 2013. doi: 10.1080/14697688.2012.708779.

Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1): 107–119, 2003. doi: 10.1023/A:1022058209073.

Thomas Ho and Hans R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981. doi: 10.1016/0304-405X(81)90020-9.

Paul Jusselin. Optimal market making with persistent order flow. *SIAM Journal on Financial Mathematics*, 12(3):1150–1200, 2021. doi: 10.1137/20M1376054.

Harold J. Kushner and Paul G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, second edition, 2001.

Ciamac C. Moallemi, Dan Robinson, and Brian Zhu. Uniform-loss automated market making for prediction markets. *arXiv preprint arXiv:2607.17428*, 2026. URL https://arxiv.org/abs/2607.17428.

Bernt Øksendal and Agnès Sulem. *Applied Stochastic Control of Jump Diffusions*. Springer, second edition, 2007.

Huyên Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer, 2009.

Conrad Oskar Voigt. Predicting prediction markets: A beta-hidden markov modeling approach. Bachelor's thesis, Stetson University, 2025.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2): 107–126, 2004. doi: 10.1257/0895330041371321.

Weiye Xi, Ciamac C. Moallemi, Mallesh M. Pai, and Shouqiao Wang. Volatility in prediction markets: A structural approach. *arXiv preprint arXiv:2607.08199*, 2026. URL https://arxiv.org/abs/2607.08199.