The Private Enforcer: Algorithmic Deterrence and the Shadow Tax on Insider Trading*

Ruoran Lai† Tingyi Lin‡ Weiyu Qi§ Karl Yu¶

Abstract

Can transparent order-flow environments discipline informed trading even in the absence of public enforcement? We study a DeFi market in which sandwich-capable MEV searchers impose immediate, state-dependent execution losses on aggressive informational trading. In a baseline equilibrium execution model, predation risk enters the trader's problem through a predictable attack intensity and state-contingent jump/loss terms, generating an endogenous shadow tax on aggressive trading and a safe-harbor region of near-zero execution. We then discipline these reduced-form objects through a bot-side partial-observation impulse-control problem, which rationalizes how greater trading aggressiveness and detectability increase predation exposure. As a pricing extension, we show that the same predation channel can load equilibrium prices with additional jump-driven tail risk under a projected Cox representation and an LP-based pricing kernel. Our contribution shows that transparent DeFi microstructure can implement a private deterrence technology with real distributional and pricing consequences.

JEL Classification: G14, D47, C73, G12

Keywords: MEV; sandwich attacks; algorithmic predation; private deterrence; informed trading; optimal execution; HJB–QVI; impulse control.

*We are grateful to Professor Lin Wanlin (Sun Yat-sen University) and Professor Wang Peng (The University of Hong Kong) for stimulating discussions and insightful feedback that helped refine our theoretical framework. We also thank Dr. Dong Ke and the faculty at the PBC School of Finance, Tsinghua University (Tsinghua PBCSF) for their detailed preliminary review and constructive comments. Any remaining errors are our own. **Funding and Disclosure:** This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The author declares that has no relevant or material financial interests that relate to the research described in this paper.

†Sun Yat-sen University. Email: lairr5@mail2.sysu.edu.cn

‡Central University of Finance and Economics. Email: hayashi@email.cufe.edu.cn(**Corresponding author.**)

§The University of Chicago. Email:weiyuq@uchicago.edu

¶Boston College. Email:yulin@bc.edu

"Private Vices, Public Benefits."

— *Bernard Mandeville (1714)*

"The primitive ancestor of smart contracts is the humble vending machine."

— *Nick Szabo (1997)*

1 Introduction

Decentralized finance (DeFi) has created trading venues in which execution and settlement are governed by open-source smart contracts rather than centralized intermediaries. While this architecture promises transparency and automation, it also exposes a stark microstructure vulnerability: *transaction-ordering dependence*. In many DeFi markets, pending orders are publicly visible before confirmation, and validators or sophisticated "searchers" can insert, delete, or reorder transactions to extract profit at the expense of ordinary traders. The phenomenon—now widely labeled maximal (or miner/validator) extractable value (MEV)—has been documented as pervasive in decentralized exchanges (DEXs), with strategic front-running, back-running, and sandwiching forming a competitive ecology around order flow (Weintraub et al. 2022). This paper argues that the economics of MEV cannot be understood purely as "latency arbitrage" or a frictional transaction cost; instead, MEV can function as a private deterrence technology that disciplines informed trading—albeit at the price of new tail risks and, under the transfer-neutral benchmark developed in Section 5.7, genuine welfare distortions when predation losses contain dissipative components (Bahrani, Garimidi, and Roughgarden 2024).

1.1 A private-deterrence view of MEV

A central difference between traditional securities markets and DeFi markets is the timing and mechanism of deterrence. In conventional markets, deterrence of insider trading and predatory conduct relies in part on public enforcement—investigations, prosecutions, and sanctions applied with delay and subject to legal and institutional constraints. Becker's classic deterrence framework emphasizes that behavior responds to the expected punishment (probability of detection times sanction severity) and that enforcement is itself an economic allocation problem (Becker 1968). DeFi does not replicate public law, but

code and competition create immediate, state-contingent execution losses for traders who reveal valuable information or urgency. In this sense, MEV searchers implement a private deterrence technology: they do not merely trade against informed orders; they can impose mechanical execution externalities on aggressive informational trading (Cong and He 2019).

This perspective reframes a long-standing concern in market microstructure: information asymmetry can degrade trade and market quality. Akerlof's "lemons" logic shows how asymmetric information can lead to adverse selection and unraveling of markets (Akerlof 1970). In DeFi, transparency reduces informational frictions for adversarial traders (searchers), but may worsen the effective trading environment for agents whose orders are informationally "high quality" (e.g., informed or urgent trades), because those orders become easier to target. The result can be an endogenous thinning of liquidity—traders optimally delay, hide, or refrain from trading when the expected "predation tax" dominates the informational rent (Akerlof 1970; Daian et al. 2020).

1.2 Relation to insider trading and continuous-time microstructure

Our modeling approach builds on the canonical informed-trading tradition. In Kyle (1985), a monopolistic informed trader optimally trades over time against competitive market makers, and prices incorporate private information gradually. Back (1992) provides a continuous-time formalization, delivering tractability and characterizing equilibrium pricing rules under broad distributions. These models, and their descendants, have become foundational in thinking about order flow toxicity, price impact, and strategic trading.

DeFi introduces a qualitatively new strategic actor: a predatory bot (or a supply chain of searchers/builders/validators) that can intervene discretely through transaction reordering and sandwiching. Such interventions are naturally modeled as impulses rather than smooth diffusion frictions, and their arrival intensities depend on observability and profitability. This discrete, strategic nature is consistent with the empirical picture in Flash Boys 2.0, which highlights transaction reordering and fee-bidding games around block position (Daian et al. 2020). Our contribution is to place this DeFi predation mechanism into a disciplined equilibrium framework that couples (i) an informed trader's

continuous-time execution problem and (ii) a strategic bot's optimal intervention problem.

1.3 Connecting MEV to market design and high-frequency trading

The "race for priority" in DeFi also resonates with the economics of high-frequency trading (HFT) and market design. Budish, Cramton, and Shim (2015) argue that continuous-time serial processing in limit order books mechanically creates arbitrage rents to speed and induces an arms race; they propose frequent batch auctions as a design response. DeFi order flow exhibits a parallel tension: continuous-time mempool visibility and sequential execution create exploitable priority opportunities (Daian et al. 2020). However, the mechanism differs in a key respect. Whereas HFT rents often arise from fleeting discrepancies in public information across venues, MEV rents arise from endogenous transaction-ordering dependence embedded in smart contract execution. Recent mechanism-design work formalizes why fee mechanisms become more difficult to design once block producers have private values for reordering—one interpretation of MEV—yielding impossibility results for simultaneously incentive-compatible mechanisms absent additional structure (Bahrani, Garimidi, and Roughgarden 2024). Our model complements this design perspective by characterizing a microstructural channel—private deterrence via predation—through which those incentive issues translate into trading behavior and price dynamics.

1.4 DeFi microstructure, AMMs, and the economics of liquidity provision

A second distinctive feature of DeFi is the dominance of automated market makers (AMMs) as trading venues, with Uniswap as the leading example. Recent finance research has begun to characterize AMM equilibria and compare AMMs to traditional limit order markets using comprehensive on-chain data, highlighting how arbitrage and liquidity provision interact in these venues (Lehar and Parlour 2025; Mohan 2022). Complementary work emphasizes that liquidity providers in AMMs face systematic adverse selection from arbitrageurs who pick off stale prices; the "loss-versus-rebalancing" (LVR) framework provides tractable formulas for this running cost and connects it to classical

continuous-time finance intuitions (Milionis et al. 2022). A formal foundations literature also models AMMs as strategic mechanisms whose outcomes depend on arbitrage and ordering, reinforcing the relevance of discrete-time interventions in DeFi settings (Bartoletti, Chiang, and Lluch-Lafuente 2022).

Our paper focuses on a related but distinct channel: rather than modeling arbitrage against stale AMM quotes, we study predation against toxic order flow, where the bot's discrete intervention changes the trader's realized execution and acts as an endogenous deterrence mechanism. This channel is particularly relevant in environments where order flow is observable prior to execution and can be re-ordered, as emphasized in MEV measurement and the evolution of private pools (Weintraub et al. 2022). Empirical work on Ethereum also documents that transaction fees, congestion, failure rates, and concentration can be substantial—conditions that shape both the feasibility and incidence of predatory strategies and the burden borne by smaller traders (Cong et al. 2023; Chiu, Kahn, and Koeppl 2022).

1.5 Model overview: HJB–QVI coupled equilibrium and the MEV shadow tax

We develop a continuous-time equilibrium model with two strategic agents. The first is a risk-averse informed trader ("insider") who chooses a trading/execution policy over a finite horizon. The second is a strategic MEV bot that observes a noisy signal of the insider's informational advantage and chooses discrete intervention times and sizes (e.g., a sandwich attack) to maximize its expected profits net of costs. The insider's problem yields a Hamilton–Jacobi–Bellman (HJB) equation; the bot's intervention problem yields a quasi-variational inequality (QVI) typical of impulse control. The equilibrium is a fixed point in which the bot's optimal policy endogenizes the intensity and severity of predation faced by the insider, while the insider's optimal policy determines the bot's opportunity set. This approach is consistent with the view that blockchain environments create strategic distortions by altering the information environment and feasible contracting space (Cong and He 2019; Biais et al. 2019).

Our first main result is the emergence of a MEV shadow tax—an endogenous wedge in the insider's first-order condition. Analogous to Becker-style deterrence, predation risk functions like a state-dependent expected penalty, but implemented privately and

instantaneously through adverse execution (Becker 1968). This shadow tax is increasing in execution aggressiveness and in the expected profitability of predation. Importantly, it generates an endogenous safe-harbor (inaction) region: for sufficiently small informational mispricing, optimal trading is to not trade. We occasionally refer to this no-trade region as a "liquidity trap" as shorthand, but the economic object is a state-contingent safe harbor generated by private predation.

1.6 Limits to arbitrage, but with bots: discipline and predation

The private-deterrence interpretation also casts new light on limits-to-arbitrage. In textbook finance, arbitrage is frictionless and riskless; in reality it requires capital and is risky, and it can fail in extreme states (Shleifer and Vishny 1997). In DeFi, MEV bots are often portrayed as exacerbating inefficiency and extracting rents; yet they can also discipline informed traders by making bot predation low-latency and immediate, thereby curbing informed exploitation of uninformed liquidity. In our equilibrium, MEV can be simultaneously (i) a source of deadweight loss through over-extraction and strategic congestion and (ii) a mechanism that limits the insider's informational rent extraction by taxing aggressive trading. This duality is consistent with the broader theme in blockchain economics that decentralized infrastructures can expand contracting possibilities while creating new strategic distortions and distributional consequences (Biais et al. 2019).

1.7 Asset-pricing implications: endogenous jumps, fat tails, and the cost of private deterrence

Private deterrence is not free. Predatory interventions introduce discontinuities into price dynamics and trading outcomes. In classical continuous-time finance, Merton (1976) provides the foundational framework for pricing when underlying returns include jumps; jump risk carries pricing consequences distinct from diffusion risk (Merton 1976). In our model, the bot's equilibrium intervention policy generates an endogenous jump component whose intensity and magnitude depend on trading incentives and the information environment. When agents price risk, the change from the physical to the pricing measure amplifies jump intensity through a jump-risk premium channel, producing fatter tails and elevated tail risk in equilibrium.

These implications relate to a growing asset-pricing literature on crypto assets. Evidence suggests that crypto returns reflect market-specific forces (network-related factors, attention, and momentum), highlighting that risk premia in crypto markets can be distinct from those in traditional assets (Liu and Tsyvinski 2021). More broadly, token-economy valuation theories emphasize feedback between adoption and token prices, implying that risk and valuation are intertwined with platform dynamics rather than cash flows (Cong, Li, and Wang 2021; Catalini and Gans 2018). Our results identify a complementary channel: a deterrence–tail-risk tradeoff, where MEV-based predation can reduce informed trading but increases jump-driven tail risk through endogenous jump exposure. In this respect, our findings complement broader concerns about the economic limits and systemic implications of permissionless trust mechanisms in blockchain systems (Budish 2025).

1.8 Related literature and contributions

Related literature. Our paper sits at the intersection of three strands of work. First, we connect the economics of deterrence and adverse selection to DeFi's transaction-ordering vulnerability (Becker 1968; Akerlof 1970). Second, we build on the informed-trading and continuous-time microstructure tradition (Kyle 1985; Back 1992), but introduce a DeFi-native strategic actor—an MEV predator—whose actions are discrete and timing-based (Weintraub et al. 2022). Third, we complement market-design and blockchain-economics perspectives: continuous-time priority creates rents and arms races, fee mechanisms become harder to design under MEV (Bahrani, Garimidi, and Roughgarden 2024), and decentralized infrastructures reshape the information environment and feasible contracting space (Cong and He 2019; Biais et al. 2019).

Contributions. We make three contributions. (i) *Mechanism:* we show that MEV can operate as a private deterrence technology, imposing an endogenous shadow tax on aggressive informational trading, generating a state-contingent safe-harbor region, and inducing a conditional welfare distortion relative to a transfer-neutral benchmark when predation losses contain dissipative components. (ii) *Method:* we provide a tractable coupled equilibrium in which the theorem-level deterrence results are established in reduced form and disciplined by a bot-side partial-observation impulse-control microfoundation. The key bridge is a projected predation intensity and projected jump/loss objects under

the insider's information set, not an identification of the compensator with the inverse of an expected boundary-hitting time. (iii) *Pricing:* we study a pricing extension in which the same projected predation process loads equilibrium prices with additional jump-driven tail risk under an LP-selected equivalent martingale measure.

1.9 Roadmap

The remainder of the paper is organized as follows. Section 2 describes the institutional environment and the insider's execution problem. Section 3 studies the baseline reduced-form deterrence model, in which predation enters through a predictable attack intensity and state-contingent jump/loss terms, and derives the shadow-tax and safe-harbor results. Section 4 studies the MEV bot's partial-observation impulse-control problem and introduces the structural hazard and projection objects that discipline the reduced-form model. Section 5 defines equilibrium as a coupled HJB–QVI fixed point linking insider optimality, bot optimality, the projected reduced-form predation objects, and a conditional welfare benchmark. Section 6 studies pricing implications under the same physical projected Cox representation and an LP-based selection of the pricing measure. Section 7 provides numerical illustrations, and the appendices contain proofs and implementation details.

2 Institutional Background & Stylized Facts

2.1 The Mechanics of MEV: Mempool and Sandwich Attacks

A defining institutional feature of many DeFi trading environments is *pre-trade observability*: user transactions are broadcast to the network and become visible before they are finalized on-chain. In public settings, pending transactions typically reside in a transaction pool (the *mempool*), where they can be monitored in near real time by validators, builders, and specialized searchers. This visibility matters because the realized execution outcome of a trade is not determined solely by the trader's intent and the automated market maker (AMM) pricing function; it also depends on the eventual *ordering* of transactions within the block that confirms them. Transaction-ordering dependence therefore creates an endogenous opportunity set: when a pending swap implies a predictable price

movement along the AMM curve, a bot can attempt to monetize that predictability by inserting strategically chosen transactions around the victim's trade.

The canonical predatory pattern is the *sandwich attack*, which can be described as a minimal three-step sequence executed within a narrow confirmation window. First, the victim transaction enters the mempool. The attacker monitors the mempool and identifies pending swaps whose implied price impact is sufficiently large—given pool depth, AMM curvature, and prevailing market conditions—to create a profitable *bracketing* opportunity. Second, the bot observes and estimates profitability. It computes an expected profit net of trading fees and the cost of securing priority, taking into account the probability of successful insertion and any competition from other searchers. Third, conditional on being profitable, the bot constructs a "front-run → victim → back-run" bundle. The bot *front-runs* by executing a trade that moves the AMM price against the victim (raising the price the victim will pay when buying, or lowering the price the victim will receive when selling). The victim then executes at the worsened price. The bot subsequently *back-runs* by unwinding its initial position after the victim's trade has further moved the price, capturing the spread created by the victim's price impact (net of fees and costs).

The economic content of the sandwich is therefore a mechanical transfer induced by ordering. In an AMM, the mapping from trade size to price movement is deterministic given the pool's state. By changing the sequence, the bot effectively changes the point on the pricing curve at which the victim executes. Put differently: the victim's trade generates predictable price pressure, and the bot converts that predictability into an exploitable round-trip wedge. Crucially, this is not merely "high-frequency trading" in the sense of reacting to public information; it is transaction ordering that directly shapes execution outcomes and can systematically penalize urgent or large trades.

2.2 Why Mempool Ordering and Priority Fees Imply Impulse Intervention

Two institutional features of DeFi execution make predation naturally *discrete* rather than continuous.

First, ordering is scarce and contested. Only one ordering is realized in the block, and to secure a desired position the bot typically must outbid others via priority fees, bundle construction, or builder/validator pipelines. This turns predation into a *threshold*

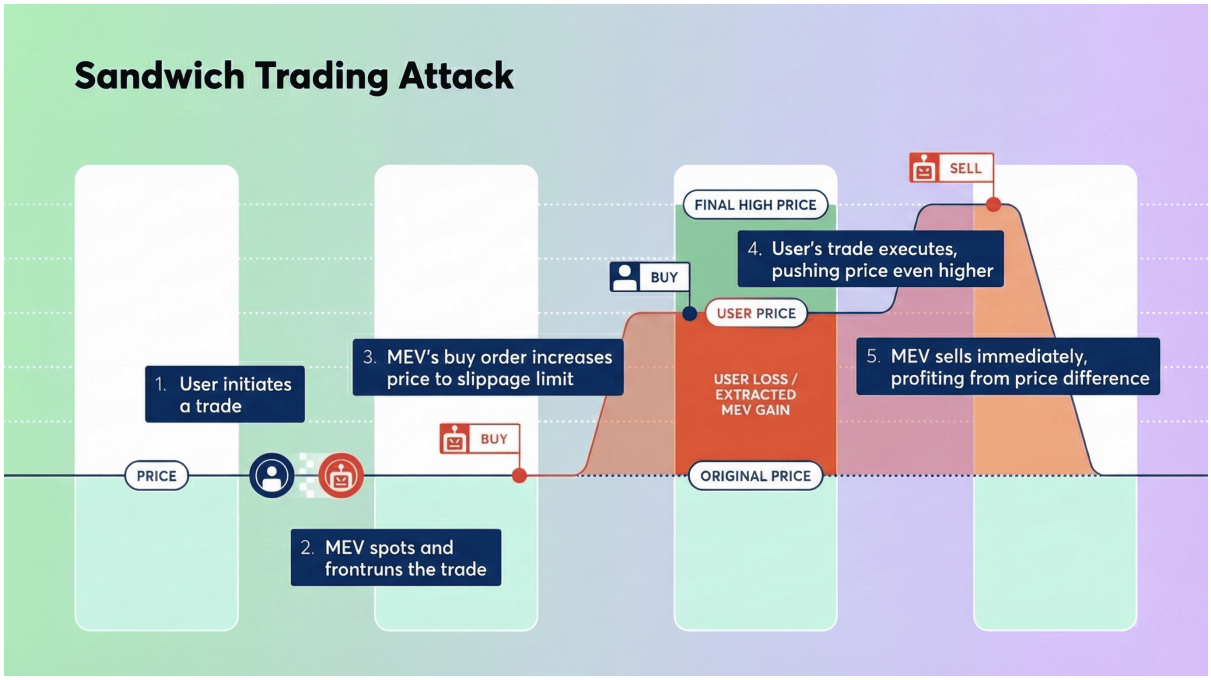


Figure 1: **Timeline of a Sandwich Attack Mechanism.**

Note: This figure illustrates the discrete lifecycle of a predatory MEV extraction, tracing how a searcher exploits mempool visibility to strategically order front-running and back-running transactions around a victim, thereby endogenously determining the execution price.

problem: the bot intervenes only when the expected surplus from bracketing the victim is large enough to justify the incremental cost of obtaining priority (and the risk of losing the race). When the opportunity is small, the optimal response is often inaction. This is precisely the logic of an *inaction region*—a hallmark of impulse-style decision problems.

Second, the timing clock differs across the trader and the bot. The informed trader (insider) can be modeled as choosing an execution intensity that can, in principle, adjust continuously over time as state variables evolve. In contrast, the bot's profitable action is tied to discrete confirmation events—specifically, block inclusion and within-block position. The bot's intervention is best thought of as arriving at discrete times (whenever a block builder can place the bracket) rather than as a smooth drift or proportional fee.

Together, these two features generate a coupled, *dual-clock environment*. The trader's continuous execution choice shapes what the bot can see and whether an attack is profitable; the bot's discrete intervention, when it happens, changes the trader's realized execution outcome and thereby alters the trader's effective marginal cost of trading. This coupling is central to the interpretation of MEV as a private deterrence technology: the market's microstructure does not merely add a constant transaction cost; it creates a state- and action-dependent predation threat that can discipline trading behavior through

the anticipation of discrete penalties.

2.3 Mapping Institutions to Reduced-Form Predation Primitives

The institutional mechanics described above map directly into the reduced-form primitives used in our model. Let θ denote the trader's execution aggressiveness (trading speed/intensity) and let p denote the relevant price state. The reduced-form objects $\lambda(\theta)$, $\gamma(p, \theta)$, and $T(\theta)$ are not arbitrary: each is a disciplined proxy for a concrete institutional component of sandwich predation.

$\lambda(\theta)$**: Visibility / Detection Technology.** This function captures the idea that more aggressive execution is easier to detect and more attractive to target. In a mempool environment, larger and more urgent swaps typically imply larger predictable price movements and provide stronger incentives for searchers to compete for ordering. Reduced-form, this is represented as an arrival intensity (or hazard) of predation opportunities that increases with $|\theta|$. Economically, $\lambda(\theta)$ is the model's representation of how execution choice translates into exposure to predation risk.

$\gamma(p, \theta)$**: Attack-Induced Price Displacement (Jump).** This term captures the discrete execution displacement induced by bracketing. When a sandwich occurs, the realized execution price is shifted relative to the no-predation counterfactual because the bot's front-run moves the AMM price before the victim, and the back-run interacts with the post-victim state. This is naturally represented as a jump-like or non-local displacement that can depend on the prevailing price level p and on the aggressiveness θ, reflecting the scale of the trade's implied movement along the AMM curve.

$T(\theta)$**: Discrete Wealth Loss upon Attack.** This function aggregates the victim's monetary harm conditional on an attack. Conceptually, this includes (i) worsened execution (slippage) due to executing at a less favorable point on the pricing curve, (ii) fee components associated with priority and execution, and (iii) adverse price impact attributable to reordering. Reduced-form, $T(\theta)$ is the model's single object for the state-contingent penalty the market imposes when the trader's order is successfully bracketed.

Key takeaway. This mapping serves as the institutional bridge for the rest of the paper: predation is endogenously triggered by execution intensity through $\lambda(\theta)$, it affects realized prices discontinuously through $\gamma(p, \theta)$, and it imposes a state-contingent penalty through $T(\theta)$. Together, these objects deliver the core *shadow tax* logic: MEV acts like a market-imposed expected penalty that is immediate and mechanically implemented through execution outcomes.

2.4 Stylized Fact 1: Leptokurtosis in DeFi Returns

A first, widely observed feature of DeFi asset returns is leptokurtosis—a distribution that is more sharply peaked around zero and exhibits fatter tails than a Gaussian benchmark. This stylized fact is central for our paper because the model's core mechanism introduces endogenous jump risk: predatory interventions generate discontinuous execution outcomes and state-dependent price displacements, and the associated jump component is amplified under the pricing measure. In other words, the private-deterrence channel we study is not merely a level shift in transaction costs; it leaves a distributional footprint in returns that naturally takes the form of excess tail risk.

Data construction. Our stylized-fact evidence is constructed at a daily frequency, ensuring consistency with the granularity of the MEV and volume aggregates used below while mitigating high-frequency microstructure noise. Let P_t denote the daily closing price (or end-of-day reference price) of a representative DeFi asset (or index). We compute daily log returns as

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right). \tag{1}$$

We then standardize returns as $z_t = (r_t - \bar{r})/\hat{\sigma}$, where $\bar{r}$ and $\hat{\sigma}$ are the sample mean and standard deviation. The benchmark Gaussian is chosen to match $(\bar{r}, \hat{\sigma})$, so that any deviation in the center and tails is attributable to higher moments rather than first- or second-moment differences. We obtain daily Open-High-Low-Close price data for Ethereum (ETH) and representative DeFi governance tokens (e.g., UNI, AAVE, MKR) from *Yahoo Finance*. The sample covers the period from 2021 to 2025. We explicitly use the adjusted closing prices to compute daily log returns, as defined in Equation (1). Yahoo Finance serves as a standard, accessible source for aggregated crypto-asset pricing, providing a robust reference for calculating daily return distributions and volatility metrics.

Visual evidence. Figure 2 provides two complementary diagnostics. Panel A plots a histogram (or kernel density) of standardized daily returns overlaid with the matched normal density. A leptokurtic distribution appears as excess mass near $z = 0$ and higher probability in the tails relative to the normal curve. Panel B displays the time-varying nature of tail risk (or rolling kurtosis). Panel C reports a QQ-plot against the normal distribution. Under normality, the QQ-plot lies close to the 45-degree line; systematic departures in the extreme quantiles indicate tail thickness.

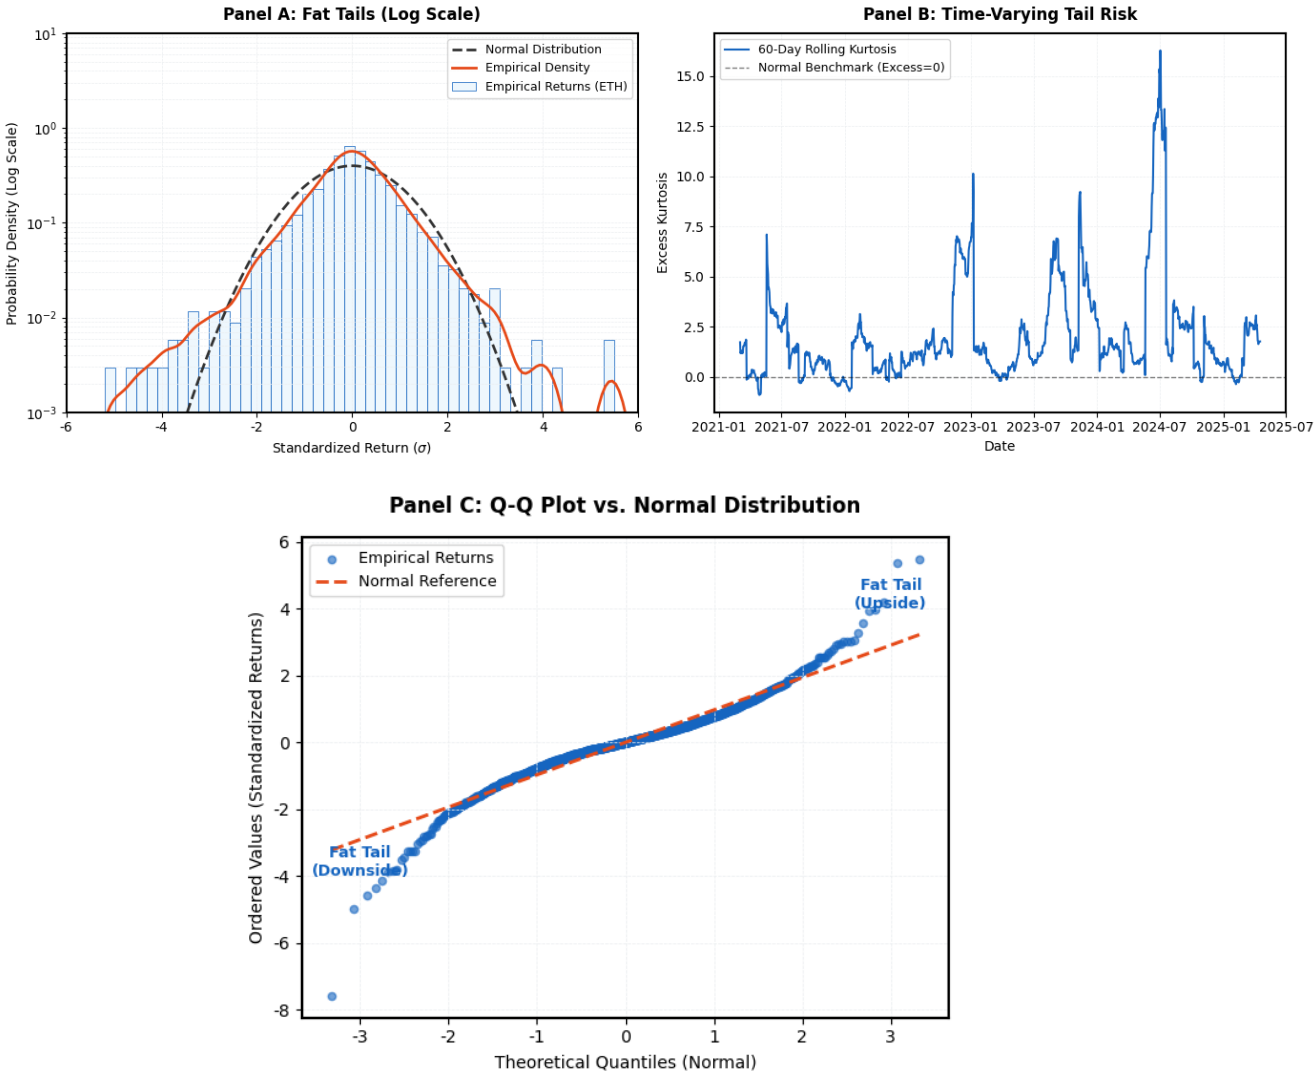


Figure 2: Leptokurtosis in DeFi returns. Panels A and B report the empirical density versus a matched Gaussian benchmark and the time-varying behavior of tail risk. Panel C reports the QQ-plot against the normal benchmark.

One-line statistic. The leptokurtosis diagnosis can be summarized by a single statistic: the sample kurtosis of daily returns is above the Gaussian benchmark of 3. Specifically,

the Pearson kurtosis is

$$\hat{\kappa} = \mathrm{Kurt}(r_t) = 8.37, \quad \hat{\kappa} - 3 = 5.37. \tag{2}$$

Equivalently, the empirical frequency of extreme observations, e.g. $\Pr(|z_t| > q)$ for a high threshold q, exceeds the matched normal probability $\Pr_N(|Z| > q)$ for the same q. (In the paper we report $q = 3$ and the corresponding empirical vs. normal exceedance rates are 1.66% vs 0.27%.)

Bridge to the model. This stylized fact motivates the model's jump-based pricing implications. The theory later shows how strategic predation generates a state-dependent jump component and how, under the pricing measure, the effective jump arrival rate can be amplified by a jump-risk-premium channel. The resulting equilibrium return distribution exhibits exactly the qualitative signature documented here: a sharper center and fatter tails than a Gaussian benchmark.

2.5 Stylized Fact 2: Trading Volume and MEV Sandwich Intensity

A second stylized fact connects market activity to predation risk. In on-chain data, periods of higher trading activity coincide with a higher intensity of *successful* sandwich attacks. We measure predation intensity using the daily count of attacker-side sandwich transactions, which corresponds to realized predation events that successfully enter the chain. We then relate this count to contemporaneous on-chain trading volume aggregated at the same daily frequency. Our empirical analysis relies on granular on-chain data sourced from **Dune Analytics**, a platform that indexes and parses raw Ethereum blockchain data. The sample period covers 2021 to 2025/4.

Figure 3 Panel A plots (log) daily trading volume against the daily count of successful sandwich attacks. To reduce the influence of outliers and emphasize the systematic relationship, the figure overlays decile-binned means and confidence intervals. The positive association is visually clear: higher trading activity is accompanied by more frequent successful sandwich attacks. Panel B provides a complementary event-time diagnostic. Treating extreme-volume days (e.g., the top 10% by volume) as "events," we plot the average sandwich count in a symmetric window around these days. Sandwich activity is elevated around high-activity windows, consistent with the view that predation opportunities expand when order-flow intensity is high.

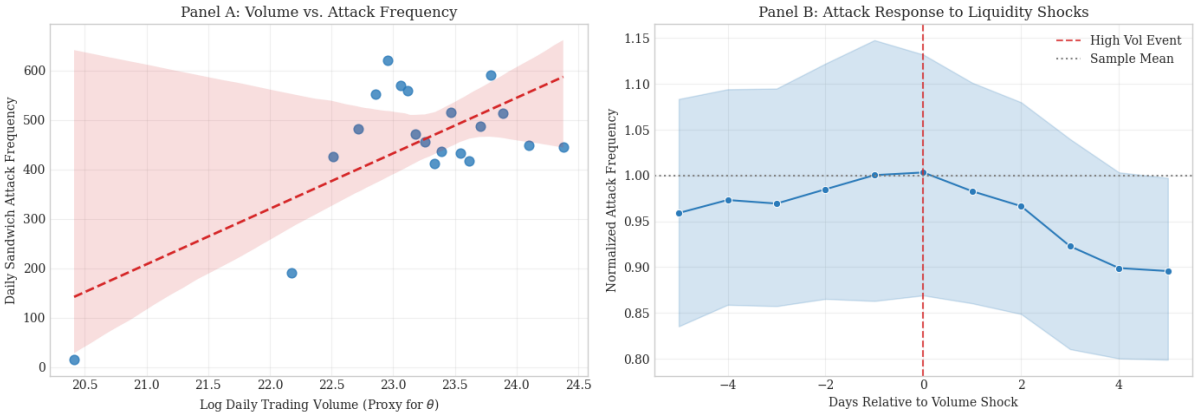


Figure 3: **Empirical Stylized Facts of Sandwich Attacks.** Panel A presents a binned scatter plot showing the positive correlation between log daily trading volume (proxy for execution-related order-flow intensity) and the frequency of sandwich attacks. Panel B illustrates an event study of the normalized attack frequency around high-volume shocks (defined as days exceeding the 90th percentile of volume). Shaded areas represent 95% confidence intervals.

We emphasize that these patterns are *stylized associations* rather than causal estimates: both trading activity and predation can respond to common shocks such as volatility, news, or congestion. Our goal is more modest—to motivate the monotonicity and curvature restrictions imposed on the reduced-form predation intensity in the execution and pricing framework.

Bridge to the model. Taken together, the second stylized fact is consistent with modeling the predation arrival intensity $\lambda(\theta)$ as increasing in execution aggressiveness and order-flow magnitude, and it provides empirical intuition for the comparative statics of the Shadow Tax mechanism developed below.

3 The Insider's Problem: Optimal Execution under Risk of Predation

We consider a finite time horizon $T < \infty$. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space equipped with a filtration $\mathcal{F} = \{\mathcal{F}_t\}_{t\geq 0}$ satisfying the usual conditions.We denote by $\mathcal{F}_t^I \subseteq \mathcal{F}_t$ the insider's information filtration.

Fundamental value. We denote the *static* fundamental by

$$V \equiv V_T.$$

The insider knows V at $t = 0$, while other market participants do not.

3.1 State Variables and Market Dynamics

The insider observes the market state and the AMM price. The effective state can be taken as (t, X_t, P_t). If pool liquidity L_t is stochastic and affects execution costs, the Markov state should be augmented to (t, X_t, P_t, L_t). For tractability, we treat L_t as publicly observable and exogenous; thus $k(L_t)$ is a known coefficient process governing the local scale of price impact and execution costs.

The Wealth Process X_t

The insider's wealth evolves according to the trading rate θ_t, net of execution costs and stochastic predation penalties:

$$dX_t = (V - P_t)\,\theta_t\,dt - \frac{1}{2}\,k(L_t)\,P_t\,\theta_t^2\,dt - T(\theta_t)\,dN_t. \tag{3}$$

where:

- $k(L_t) > 0$ is the convex slippage parameter.
- $T(\theta)$ represents the reduced-form discrete wealth loss incurred during a sandwich attack (e.g., adverse price movement + priority fees). We assume $T(\cdot)$ is strictly convex.

Reduced-form convention: $T(\theta)$ is the insider-side projected loss induced by the bot's optimal intervention policy in Section 4. Formally, one may start from a structural loss $T_{\text{struct}}(\theta, a)$ and define

$$T(\theta) \equiv T^I(\theta) = \mathbb{E}\big[T_{\text{struct}}\big(\theta, a^\star(\Xi_t)\big) \mid \theta_t = \theta\big]\,.$$

The Price Process P_t

The AMM price dynamics are driven by the insider's continuous flow and the bot's discrete interventions:

$$dP_t = k(L_t)\,P_t\,\theta_t\,dt + \sigma_P\,P_t\,dW_t + \gamma(P_{t-}, \theta_t)\,dN_t. \tag{4}$$

where:

- $\sigma_P > 0$ is the diffusive volatility of the AMM price.
- $\gamma(p, \theta)$ denotes the reduced-form instantaneous price displacement caused by a sandwich attack. It is increasing in $|\theta|$ because larger victim orders are more profitable to exploit.
- W_t is a standard Brownian motion.

Reduced-form convention: In Section 4, the bot chooses an impulse intensity $a > 0$ generating a structural jump $\Gamma(p, \theta, a)$. The reduced-form jump in this section is the projected insider-side object

$$\gamma(p, \theta) \equiv \gamma^I(p, \theta) := \mathbb{E}\big[\Gamma\big(p, \theta, a^\star(\Xi_t)\big) \mid \theta_t = \theta,\ P_t = p\big]\,,$$

where $a^\star$ is the bot's optimal impulse control and $\Xi_t = (V_t, \Sigma_t, P_t, L_t)$ is the bot's structural state.

Remark 1 (**Impact normalization**). *The coefficient $k(L_t)p$ should be read as a reduced-form local impact scale. In the baseline notation it appears both in the continuous price drift dP_t and in the quadratic running execution cost in dX_t. This is a parsimonious normalization rather than a claim that temporary and permanent impact are literally identical. One can equivalently write*

$$dX_t = (V - P_t)\theta_t\, dt - \frac{1}{2}k_X(L_t)P_t\theta_t^2\, dt - T(\theta_t)\, dN_t,$$

$$dP_t = k_P(L_t)P_t\theta_t\, dt + \sigma_P P_t\, dW_t + \gamma(P_{t-}, \theta_t)\, dN_t,$$

with two positive regular functions k_X and k_P. All HJB, QVI, and equilibrium arguments below go through with the obvious substitutions. We set $k_X \equiv k_P \equiv k$ to keep the baseline notation compact.

Remark 2 (**From Microstructure to Projected Predation Objects**). *Section 4 introduces two filtrations. Let $\mathcal{F}_t^B$ denote the bot filtration generated by the structural state $\Xi_t = (V_t, \Sigma_t, P_t, L_t)$ and the block-level attack clock, and let $\mathcal{F}_t^I$ denote the insider/public filtration generated by publicly observed states and the insider's own execution policy. Under the bot filtration, the optimally controlled attack process admits a predictable structural hazard kernel $h^B(\Xi_t)$ such that*

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}^B) = h^B(\Xi_t)\, dt + o(dt).$$

The insider-relevant object is the predictable projection

$$\lambda_t^I \ \equiv\ \mathbb{E}\big[h^B(\Xi_t)\,\big|\,\mathcal{F}_{t-}^I\big]\,. \tag{5}$$

Conditioning further on local execution aggressiveness and the public state yields the reduced-form representation used in Section 3:

$$\lambda(t,p,L,\theta) := \mathbb{E}\big[h^B(V_t,\Sigma_t,p,L)\,\big|\,\theta_t=\theta,\ P_t=p,\ L_t=L\big]\,, \tag{6}$$

where $\lambda(\theta)$ is shorthand when the public-state arguments are suppressed. Formally, the reduced-form attack intensity used here is the insider-filtration predictable projection of the structural bot-side hazard defined in Section 4; the boundary-hitting statistic is used only as a diagnostic summary of predation exposure, not as the compensator itself. The same projection logic applies to the reduced-form jump and loss objects, written in Section 3 as $\gamma(p,\theta)$ and $T(\theta)$.

The MEV Hazard Rate (Reduced-Form Controlled Intensity)

Predation arrivals are modeled by a point process N_t with insider-filtration predictable intensity $\lambda(\theta_t)$. Over an infinitesimal interval:

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}) = \lambda(\theta_t)\,dt + o(dt).$$

Motivated by Stylized Fact 2 in Section 2.5, we assume that $\lambda(\theta)$ is increasing and convex in $|\theta|$.

Assumption 1 (Visibility / Mempool Detection)**.** *$\lambda(\theta)$ is strictly increasing and convex in $|\theta|$. Specifically,*

$$\lambda(\theta) = \lambda_0 + \lambda_1|\theta|^\beta, \qquad \beta \geq 1.$$

3.2 The Value Function (CARA Preferences)

The insider maximizes expected utility of terminal wealth with CARA utility:

$$U(x) = -e^{-\alpha x}, \qquad \alpha > 0.$$

The value function is:

$$H(t,x,p) = \sup_{\theta\in\mathcal{A}} \mathbb{E}\big[-e^{-\alpha X_T}\,\big|\,X_t=x,\ P_t=p\big]\,.$$

3.3 Separation Ansatz and Certainty Equivalent $J(t,p)$

We use the exponential-CARA separation:

$$H(t,x,p) = -\exp\{-\alpha(x + J(t,p))\},$$

where $J(t,p)$ is the certainty-equivalent continuation value ("remaining information rents").

Terminal condition.

$$J(T,p) = 0,$$

since $H(T,x,p) = -e^{-\alpha x}$.

3.4 Derivation of the HJB Equation

By the Dynamic Programming Principle,

$$\frac{\partial H}{\partial t} + \sup_{\theta}\{\mathcal{L}^{\theta} H\} = 0,$$

where $\mathcal{L}^{\theta}$ is the controlled generator (diffusion + jumps).

Step 1: Diffusive generator

Compute derivatives:

$$H_x = -\alpha H, \qquad H_p = -\alpha J_p H, \qquad H_{pp} = (-\alpha J_{pp} + \alpha^2 J_p^2)\, H.$$

The continuous generator is:

$$\mathcal{L}^{\theta}_{\text{cont}} H = H_x \left((V - p)\theta - \frac{1}{2} k(L_t) p \theta^2 \right) + H_p \left(k(L_t) p \theta \right) + \frac{1}{2} \sigma_P^2 p^2 H_{pp}.$$

Substituting gives:

$$\mathcal{L}^{\theta}_{\text{cont}} H = (-\alpha H) \Big[(V - p)\theta - \frac{1}{2} k(L_t) p \theta^2 + k(L_t) p \theta\, J_p + \frac{1}{2} \sigma_P^2 p^2 \big(J_{pp} - \alpha J_p^2 \big) \Big].$$

Step 2: Jump generator (MEV attack)

At an attack:

$$x \mapsto x - T(\theta), \qquad p \mapsto p + \gamma(p, \theta).$$

Thus,

$$H\big(t, x - T, p + \gamma\big) = H(t, x, p) \cdot \exp\left\{\alpha\big[T(\theta) + J(t, p) - J(t, p + \gamma)\big]\right\}.$$

3.5 Optimal Control, FOC, and the "Fear Premium"

Define the certainty-equivalent jump loss

$$\Psi(\theta, p) = T(\theta) + J(t, p) - J\big(t, p + \gamma(p, \theta)\big). \tag{7}$$

Let the Hamiltonian be:

$$\mathcal{H}(\theta) = (V - p)\theta - \frac{1}{2}k(L_t)p\theta^2 + k(L_t)p\theta J_p - \frac{\lambda(\theta)}{\alpha}\big(e^{\alpha\Psi} - 1\big). \tag{8}$$

For an interior solution, $\theta^\star$ satisfies $\partial_\theta \mathcal{H}(\theta^\star) = 0$. Define the jump risk factor:

$$\Delta_{\text{jump}} \equiv e^{\alpha\Psi(\theta^\star, p)} - 1.$$

Differentiate (7):

$$\Psi_\theta(\theta, p) = T'(\theta) - J_p\big(t, p + \gamma\big)\,\gamma_\theta(p, \theta).$$

Then the FOC becomes:

$$V - p = k(L_t)p\,\theta^\star - k(L_t)p\,J_p + \text{MPC}(\theta^\star),$$

where the marginal predation cost is

$$\text{MPC}(\theta) = \frac{\lambda'(\theta)}{\alpha}\big(e^{\alpha\Psi} - 1\big) + \lambda(\theta)\,e^{\alpha\Psi}\Big[T'(\theta) - J_p\big(t, p + \gamma\big)\gamma_\theta(p, \theta)\Big] \tag{9}$$

$$= \underbrace{\frac{\lambda'(\theta)}{\alpha}\big(e^{\alpha\Psi} - 1\big)}_{\text{Visibility / Detection}} + \underbrace{\lambda(\theta)\,e^{\alpha\Psi}T'(\theta)}_{\text{Sandwich geometry}} + \underbrace{\lambda(\theta)\,e^{\alpha\Psi}\big(-J_p(t, p + \gamma)\big)\gamma_\theta(p, \theta)}_{\text{Endogenous state collapse}}. \tag{10}$$

3.6 Proposition: The Generalized MEV Shadow Tax

Proposition 1 (MEV Shadow Tax under Risk Aversion)**.** *The optimal execution speed is suppressed by a composite shadow tax that scales exponentially with the potential MEV loss Ψ. The tax decomposes into:*

- ***Visibility cost:*** $\frac{\lambda'(\theta)}{\alpha}\big(e^{\alpha\Psi}-1\big)$,
- ***Sandwich geometry:*** $\lambda(\theta)\, e^{\alpha\Psi} T'(\theta)$,
- ***Endogenous state compression (safe-harbor amplification):*** $\lambda(\theta)\, e^{\alpha\Psi}\big(-J_p(t, p+\gamma)\big)\gamma_\theta(p,\theta)$.

A larger trade pushes the post-attack price $p+\gamma$ further from the fundamental, eroding continuation rents. Under risk aversion this state-contingent erosion is exponentially penalized.

Risk-neutral check. As $\alpha \to 0$,

$$\frac{e^{\alpha\Psi}-1}{\alpha} \to \Psi,$$

so the jump cost becomes $-\lambda(\theta)\Psi$, consistent with expected-loss intuition.

4 The MEV Bot's Problem: Filtering and Impulse Control

We model the MEV bot as a strategic predator solving a partial-observation impulse control problem. Given the timescale disparity between block-level MEV extraction (milliseconds) and the insider's execution horizon (minutes/hours), we formulate the bot's problem under an infinite-horizon stationary approximation. Since learning is non-stationary when V is static, we restore time-homogeneity by augmenting the state with the posterior variance Σ.

4.1 The Observation Process

The bot does not observe the fundamental value V or the insider's latent motive directly. What it observes is a noisy aggregate of exploitable order-flow signals generated within

the short block-level decision window: submitted transaction size, splitting patterns, queue position, competing flow, uncertain inclusion, and within-block ordering risk. We summarize that observation channel by

$$dY_t = \theta_t \, dt + \sigma_Y \, dB_t,$$

where B_t is a standard Brownian motion in the bot's signal channel and $\sigma_Y > 0$ measures residual observational noise. This specification is not meant to deny mempool visibility. Rather, it captures the gap between seeing a transaction payload and inferring, in real time, the trader's latent informational motive and the profitability of a sandwich.

Assumption 2 (Local strategic linearization)**.** *Fix a short block-level window on which the public state (P_t, L_t) is locally frozen to first order. In the active-trading region, the insider's feedback rule admits a local slope summary*

$$\theta_t = \theta^\star(t, P_t; V, L_t) \approx \beta^g \, (V - P_t), \qquad \beta^g > 0,$$

where β^g is computed from the current candidate insider policy g. This approximation is used only inside the bot's short-window filtering problem and is not asserted as a global identity for the insider's equilibrium policy.

Remark 3 (Local slope extraction and approximation error)**.** *Write mispricing as $m := V - P_t$ and let $F^g(t, p, L, m, \theta) = 0$ denote the insider first-order condition induced by the HJB. On any compact active-trading window away from the safe-harbor boundary, the strict-concavity condition used in Appendix 8.2 implies that $\partial_\theta F^g$ is bounded away from zero. Hence the implicit-function theorem yields local differentiability of the feedback rule in m around any operating point (t_0, p_0, L_0) in the active region:*

$$\theta^\star(t, p; V, L) = \theta_0^\star + \beta^g \, (m - m_0) + O(|m - m_0|^2), \qquad \theta_0^\star := \theta^\star(t_0, p_0; V, L_0).$$

In the analytical approximation, β^g denotes this local derivative. In the numerical implementation, it is the corresponding local least-squares slope over a short active-region window. The filtering problem uses only this block-level local representation.

4.2 Filtering Problem

Let

$$V_t = \mathbb{E}[V \mid \mathcal{F}_t^Y], \qquad \Sigma_t = \mathbb{E}[(V - V_t)^2 \mid \mathcal{F}_t^Y]$$

denote the bot's posterior mean and posterior variance.

Innovation process. Conditional on the local-linearization approximation and on the no-jump continuation within the infinitesimal observation window, define the normalized innovation by

$$dZ_t = \frac{1}{\sigma_Y} \left(dY_t - \beta^g (V_t - P_t)\, dt \right),$$

which is a Brownian motion under $\mathcal{F}_t^Y$ on the continuation region.

Posterior mean dynamics.

$$dV_t = K(\Sigma_t)\, dZ_t, \qquad K(\Sigma) = \frac{\beta^g \Sigma}{\sigma_Y}.$$

Equivalently,

$$dV_t = \frac{\beta^g \Sigma_t}{\sigma_Y^2} \left(dY_t - \beta^g (V_t - P_t)\, dt \right).$$

Posterior variance dynamics. Under the same local continuation approximation,

$$\frac{d\Sigma_t}{dt} = -K(\Sigma_t)^2 = -\frac{(\beta^g)^2 \Sigma_t^2}{\sigma_Y^2}.$$

This Riccati equation is a short-window approximation for the bot's continuation problem; it is not a claim that the global filtering problem remains Gaussian across attack times or across the full nonlinear insider policy.

Continuation filtering with attack updates. Let $\{\tau_i\}_{i \geq 1}$ denote the attack times of N. The diffusion equations above describe the continuation filter on each interval $[\tau_i, \tau_{i+1})$, where the attack process is locally constant. At an observed attack time τ_i, or more generally when the bot observes an intervention-related mark m_i at τ_i, the posterior law

$$\pi_t(dv) := \mathbb{P}\big(V \in dv \mid \mathcal{F}_t^Y \vee \mathcal{F}_t^N\big)$$

updates discontinuously by Bayes' rule:

$$\pi_{\tau_i}(dv) = \frac{\ell(m_i \mid v, \Xi_{\tau_i-})\, \pi_{\tau_i-}(dv)}{\int \ell(m_i \mid \tilde{v}, \Xi_{\tau_i-})\, \pi_{\tau_i-}(d\tilde{v})}.$$

Accordingly, (V_t, Σ_t) should be read as a piecewise-diffusive belief state with discrete Bayesian updates at attack times. The local Gaussian equations above characterize the continuation leg of this filter; intervention events inject additional information through the jump-update map.

4.3 Impulse Control Problem

The bot selects a sequence of attack times $\{\tau_i\}_{i\geq 1}$ and sandwich intensities $\{a_i\}_{i\geq 1}$ (front-running order size) to maximize discounted profit.

Objective (stationary infinite horizon). Let $\Phi(v, p, \Sigma, L)$ denote the stationary value function with state (v, p, Σ, L):

$$\Phi(v, p, \Sigma, L) = \sup_{\{\tau_i, a_i\}} \mathbb{E}\left[\sum_{i=1}^{\infty} e^{-r\tau_i}\Big(R(a_i, v_{\tau_i} - p_{\tau_i}) - C_{\text{gas}}\Big) \;\Big|\; (v_0, p_0, \Sigma_0, L_0) = (v, p, \Sigma, L)\right],$$

where $r > 0$, $R(\cdot)$ is strictly concave in a, and $C_{\text{gas}} > 0$ is the fixed per-attack cost.

Impulse effect on price. At an attack, the bot chooses $a > 0$ generating a structural price displacement:

$$p \mapsto p + \Gamma(p, \theta, a).$$

Under Assumption 2, the bot substitutes $\theta \approx \beta^g(v - p)$.

4.4 Quasi-Variational Inequalities (QVI)

The value function Φ satisfies the stationary QVI:

$$\min\{\, r\Phi - \mathcal{L}_{\text{bot}}\Phi,\;\; \Phi - \mathcal{M}\Phi \,\} = 0.$$

1) Continuation region (waiting)

Under Assumption 2,

$$\mathbb{E}[\theta_t \mid \mathcal{F}_t^Y] \approx \beta^g (v - p).$$

Using the AMM drift, the bot's perceived price drift is:

$$\mu_P(v, p, L) = k(L)\, p\, \beta^g (v - p).$$

Since Σ evolves deterministically, the generator includes $\dot{\Sigma}\, \Phi_\Sigma$. Hence:

$$\mathcal{L}_{\text{bot}} \Phi = \mu_P(v, p, L)\, \Phi_p + \frac{1}{2} \sigma_P^2 p^2\, \Phi_{pp} + \frac{1}{2} K(\Sigma)^2\, \Phi_{vv} + \dot{\Sigma}\, \Phi_\Sigma,$$

where

$$K(\Sigma) = \frac{\beta^g \Sigma}{\sigma_Y}, \qquad \dot{\Sigma} = -\frac{(\beta^g)^2 \Sigma^2}{\sigma_Y^2}.$$

Cross-term: since B_t and W_t are independent ($\rho = 0$), the mixed term Φ_{vp} is zero. Under correlation, an additional term $\rho\, K(\Sigma) \sigma_P p\, \Phi_{vp}$ would appear.

2) Intervention region (attacking)

The intervention operator is:

$$\mathcal{M}\Phi(v, p, \Sigma, L) = \sup_{a > 0} \left\{ \Phi\big(v,\ p + \Gamma(p, \beta^g (v - p), a),\ \Sigma,\ L\big) + R(a, v - p) - C_{\text{gas}} \right\}.$$

4.5 Optimal Strategy, Regions, and the Projected Predation Intensity

The QVI partitions the augmented state space into the continuation region

$$\mathcal{C} : \ \Phi > \mathcal{M}\Phi \qquad \text{(wait and filter)}$$

and the intervention region

$$\mathcal{A} : \ \Phi = \mathcal{M}\Phi \qquad \text{(attack immediately)}.$$

Let

$$\tau_{\partial\mathcal{A}} := \inf\{t \geq 0 : (V_t, P_t, \Sigma_t, L_t) \in \mathcal{A}\}$$

denote the first entry time into the intervention region.

Boundary-hitting as an exposure statistic. The quantity

$$e(\xi) := \mathbb{E}_\xi[\tau_{\partial\mathcal{A}}]^{-1}, \qquad \xi = (v, p, \Sigma, L),$$

is used as a summary statistic for how quickly the optimally controlled state reaches intervention. It is informative about comparative statics of predation exposure, but it is *not* identified directly with the compensator of the attack counting process unless one imposes an additional memoryless/Poisson approximation.

Structural attack clock. Let N_t denote the bot's attack counting process on the block-level decision clock. We assume that, under the bot filtration $\mathcal{F}_t^B$, the optimally controlled attack process admits a predictable structural hazard kernel $h^B(\Xi_t)$ such that

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}^B) = h^B(\Xi_t)\, dt + o(dt), \qquad \Xi_t = (V_t, \Sigma_t, P_t, L_t).$$

The QVI determines the sign, monotonicity, and state dependence of h^B: states with greater toxicity, detectability, and profitability lie closer to or inside the intervention region and therefore carry larger structural attack hazard.

Projection to the insider's information set. Let $\mathcal{F}_t^I$ denote the insider/public filtration generated by publicly observed states and the insider's own execution policy. The insider-relevant predation intensity is the predictable projection

$$\lambda_t^I := \mathbb{E}\big[h^B(\Xi_t) \mid \mathcal{F}_{t-}^I\big]\,.$$

In reduced form, we write

$$\lambda(t, p, L, \theta) := \mathbb{E}\big[h^B(V_t, \Sigma_t, p, L) \mid \theta_t = \theta,\ P_t = p,\ L_t = L\big]\,,$$

and suppress (t, p, L) when no confusion arises. The maintained reduced-form restriction used in Section 3 is that λ is increasing and convex in $|\theta|$.

Projected jump and loss objects. The bot's optimal intervention also induces reduced-form execution objects for the insider. Define

$$\gamma^I(p,\theta) := \mathbb{E}\big[\Gamma\big(p,\theta,a^\star(\Xi_t)\big) \mid \theta_t = \theta,\ P_t = p\big]\,,$$

and

$$T^I(\theta) := \mathbb{E}\big[T_{\text{struct}}\big(\theta,a^\star(\Xi_t)\big) \mid \theta_t = \theta\big]\,.$$

These projected objects are the reduced-form jump and loss terms that enter the insider's HJB. For notational simplicity, Section 3 writes them as $\gamma(p,\theta)$ and $T(\theta)$.

Assumption 3 (Projected hazard reduction)**.** *Under the block-level attack clock, the optimally controlled attack process admits the structural predictable hazard kernel $h^B(\Xi_t)$ above. Its predictable projection onto the insider/public filtration exists, is bounded on the truncated state domain, and induces reduced-form objects $\lambda(t,p,L,\theta)$, $\gamma^I(p,\theta)$, and $T^I(\theta)$ that satisfy the monotonicity, regularity, and integrability conditions imposed in Sections 3 and 6.*

Primitive construction under a block-level opportunity clock. Appendix 8.2 derives h^B from a marked block-level opportunity clock rather than leaving it as a free reduced-form shortcut. Let $\mathcal{N}(ds,du)$ denote candidate attack opportunities with predictable compensator $\nu(\Xi_s)F(du \mid \Xi_s)\,ds$, where the mark u summarizes builder access, within-block position, and competing-flow conditions. If the bot attacks at a candidate opportunity exactly when the mark-specific intervention gain is nonnegative, then

$$N_t = \int_0^t \int \mathbf{1}\big\{\Delta^\Phi(\Xi_{s-},u) \geq 0\big\}\,\mathcal{N}(ds,du),$$

where $\Delta^\Phi(\xi,u) := \mathcal{M}_u\Phi(\xi) - \Phi(\xi)$ is the gain from attacking rather than continuing at mark u, and the structural hazard becomes

$$h^B(\Xi_t) = \nu(\Xi_t)\int \mathbf{1}\big\{\Delta^\Phi(\Xi_t,u) \geq 0\big\}\,F(du \mid \Xi_t).$$

In the deterministic-mark special case, this reduces to $h^B(\Xi_t) = \nu(\Xi_t)\mathbf{1}_{\{\Xi_t \in \mathcal{A}\}}$, tying the hazard directly to the QVI intervention region.

5 Equilibrium via a Coupled HJB–QVI Fixed Point

This section characterizes a Markovian equilibrium of the strategic interaction between the insider (Section 3) and the MEV bot (Section 4). The equilibrium is expressed as a fixed point linking insider optimality, bot optimality, and the projected predation objects induced by the bot's optimal policy.

5.1 State, Information, and Reduced-Form Attack Representation

Public state variables

Let the publicly observable market state be

$$S_t \equiv (P_t, L_t).$$

Bot belief state

The bot's belief state is

$$B_t \equiv (V_t, \Sigma_t),$$

and the augmented structural state is

$$\Xi_t \equiv (V_t, \Sigma_t, P_t, L_t).$$

Projected predation objects

Given a candidate insider policy and the bot's optimal intervention policy, the structural attack process admits a bot-filtration hazard kernel $h^B(\Xi_t)$. The insider-relevant intensity is its predictable projection

$$\lambda_t^I = \mathbb{E}\big[h^B(\Xi_t) \mid \mathcal{F}_{t-}^I\big].$$

Similarly, the reduced-form execution objects entering the insider problem are the projected jump and loss terms

$$\gamma^I(p,\theta) := \mathbb{E}\big[\Gamma\big(p,\theta,a^\star(\Xi_t)\big) \mid \theta_t = \theta,\ P_t = p\big]\,,$$

$$T^I(\theta) := \mathbb{E}\big[T_{\text{struct}}\big(\theta,a^\star(\Xi_t)\big) \mid \theta_t = \theta\big]\,.$$

For notational simplicity, we write $\gamma(p,\theta)$ and $T(\theta)$ in the baseline HJB.

5.2 Forward Dynamics (Physical Measure $\mathbb{P}$)

Under $\mathbb{P}$, the AMM price is driven by insider execution and predation jumps:

$$dP_t = k(L_t)P_t\theta_t^\star\,dt + \sigma_P P_t\,dW_t + \gamma^I(P_{t-},\theta_t^\star)\,dN_t,$$

where N_t is a counting process with $\mathbb{P}$-compensator

$$\int_0^t \lambda_s^I\,ds.$$

Equivalently,

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}^I) = \lambda_t^I\,dt + o(dt).$$

5.3 Insider Optimality (HJB, Section 3 Object)

The certainty-equivalent value $J(t,p)$ satisfies:

$$\begin{aligned} 0 = {} & J_t + \frac{1}{2}\sigma_P^2 p^2\left(J_{pp} - \alpha J_p^2\right) \\ & + \sup_{\theta\in\mathcal{A}}\left\{(V-p)\theta - \frac{1}{2}k(L_t)p\theta^2 + k(L_t)p\theta\,J_p - \frac{\lambda(\theta)}{\alpha}\left(e^{\alpha\Psi(\theta,p)}-1\right)\right\}, \qquad J(T,p) = 0, \end{aligned}$$

Here $\lambda(\theta)$, $\gamma(p,\theta)$, and $T(\theta)$ are shorthand for the projected reduced-form objects induced by the bot's optimal policy under Assumption 3. The quantity $\Psi(\theta,p)$ is defined in (7). The insider's equilibrium control is

$$\theta_t^\star = \Theta^\star(t,P_t;V,L_t).$$

5.4 Bot Optimality (QVI, Section 4 Object)

The bot's value function $\Phi(v, p, \Sigma, L)$ satisfies:

$$\min\{\, r\Phi - \mathcal{L}_{\text{bot}}\Phi,\ \ \Phi - \mathcal{M}\Phi \,\} = 0,$$

with

$$\mathcal{L}_{\text{bot}}\Phi = \mu_P(v, p, L)\, \Phi_p + \frac{1}{2}\sigma_P^2 p^2\, \Phi_{pp} + \frac{1}{2}K(\Sigma)^2\, \Phi_{vv} + \dot{\Sigma}\, \Phi_\Sigma,$$

where $\mu_P(v, p, L) = k(L)p\beta^g(v - p)$, $K(\Sigma) = \beta^g\Sigma/\sigma_Y$, and $\dot{\Sigma} = -(\beta^g)^2\Sigma^2/\sigma_Y^2$; and

$$\mathcal{M}\Phi(v, p, \Sigma, L) = \sup_{a>0} \left\{ \Phi\big(v, p + \Gamma(p, \beta^g(v - p), a), \Sigma, L\big) + R(a, v - p) - C_{\text{gas}} \right\}.$$

5.5 Markovian Equilibrium as a Fixed Point

The baseline equilibrium notion in Sections 3–5 is conditional on the public AMM execution environment. In particular, the local pricing rule $k(L_t)P_t$ and the public liquidity state L_t are taken as exogenous primitives of the trading venue, while the strategic fixed point is solved between the insider and the bot. Endogenizing LP supply would add further market-clearing or zero-profit conditions, but those are not part of the baseline execution–predation equilibrium studied here.

Definition 1 (Markovian equilibrium)**.** *A Markovian equilibrium of the execution–predation game is a collection of objects*

$$\Theta^\star,\ \ \Phi^\star,\ \ a^\star,\ \ h^B,\ \ \lambda^I,\ \ \gamma^I,\ \ T^I$$

such that:

- ***Public environment:*** *the AMM execution schedule $k(L_t)P_t$ and the public liquidity state L_t are taken as given common-knowledge primitives of the venue.*
- ***Insider optimality:*** *$\theta_t^\star = \Theta^\star(t, P_t; V, L_t)$ solves the insider's HJB given the projected predation objects $(\lambda^I, \gamma^I, T^I)$.*
- ***Bot optimality:*** *$\Phi^\star$ solves the QVI, yielding the optimal intervention rule $a^\star$ and intervention region $\mathcal{A}$.*

- ***Consistency:*** *the bot's optimal policy induces the structural hazard kernel h^B, whose predictable projection onto the insider/public filtration is λ^I; the same policy induces the projected reduced-form jump and loss objects γ^I and T^I used by the insider.*

5.6 Existence

Assumption 4 (Standing conditions for existence)**.** *On the truncated compact domains used in Appendix 8.2, assume:*

(i) the control and impulse sets are compact, and primitive coefficients are continuous and bounded;

(ii) the insider HJB with the nonlocal CARA jump term satisfies comparison in the bounded state-constraint viscosity class and admits a bounded continuous solution;

(iii) the bot QVI admits a unique bounded continuous viscosity solution, and $C_{gas} > 0$ rules out Zeno accumulation;

(iv) the local slope summary operator $\mathcal{B}$, the structural hazard kernel h_g^B, and the shape-preserving projection Π_λ are continuous under uniform convergence;

(v) the induced reduced-form mapping $\mathscr{T}$ preserves the admissible class.

Detailed sufficient conditions are collected in Appendix 8.2 and Appendix 8.2.

Proposition 2. *Under Assumption 4, there exists a Markovian equilibrium.*

Proof strategy. Define the best-response mapping $\mathscr{T}$ on the admissible reduced-form class and apply Schauder's fixed-point theorem. The main technical inputs are boundedness and comparison for the nonlocal CARA HJB, viscosity well-posedness of the bot QVI, continuity of the local slope summary operator, and continuity of the hazard-projection map.

5.7 Conditional Welfare Benchmark

The strategic fixed point above is positive rather than normative. To identify whether private predation is purely redistributive or genuinely distortionary, decompose the pro-

jected insider loss into a transfer component and a dissipative component:

$$T^I(\theta) = \Pi^I(\theta) + D^I(\theta), \qquad D^I(\theta) \geq 0,$$

where Π^I is a private transfer from the insider to the predator and D^I collects real resource costs such as burned priority fees, congestion externalities, and execution degradation not rebated within the system.

For a fixed public state $s = (p, L)$, define the transfer-neutral flow surplus

$$w_s(\theta) = (V - p)\theta - \frac{1}{2}k(L)p\theta^2 - \lambda(\theta)D^I(\theta),$$

and let $\theta^{FB}(s)$ denote its unique maximizer.

Let the insider's private static objective at state s be written as

$$\pi_s(\theta) = w_s(\theta) - G_s(\theta),$$

where G_s collects the private-transfer component of predation together with the continuation-value erosion wedge generated by the nonlocal term in Proposition 1.

Proposition 3 (Private predation and conditional welfare distortion)**.** *Suppose w_s is strictly concave. Assume that G_s is convex, differentiable, and normalized by $G_s(0) = 0$, and that*

$$\theta\, G_s'(\theta) \geq 0$$

for every admissible θ. Then the private objective $\pi_s = w_s - G_s$ is strictly concave and has a unique maximizer $\theta^\star(s)$ satisfying

$$|\theta^\star(s)| \leq |\theta^{FB}(s)|$$

for every state s. If, in addition,

$$\theta^{FB}(s)\, G_s'\big(\theta^{FB}(s)\big) > 0,$$

then the inequality is strict. Whenever $\theta^\star(s) \neq \theta^{FB}(s)$*, one has*

$$w_s(\theta^\star(s)) < w_s(\theta^{FB}(s)).$$

Proof idea. Relative to the transfer-neutral benchmark w_s, the insider's private objective subtracts the wedge G_s. Convexity and the sign restriction $\theta G_s'(\theta) \geq 0$ imply that G_s is minimized at zero and that its marginal effect always pushes the optimizer back toward the origin. Because w_s is strictly concave, the maximizer of $\pi_s = w_s - G_s$ must therefore lie weakly between zero and the transfer-neutral maximizer $\theta^{FB}(s)$, with strict inward movement whenever the wedge is active at $\theta^{FB}(s)$. Strict concavity of w_s then yields the strict transfer-neutral welfare loss whenever the two controls differ. A formal proof is given in Appendix 8.2.

Interpretation. This benchmark identifies a genuine welfare distortion, not merely a redistribution. In the present model, G_s contains the transfer component of predation and the nonlocal continuation-value wedge, both of which operate privately but need not coincide with the planner-relevant real resource cost D^I. As a result, private deterrence can induce under-trading relative to a transfer-neutral benchmark even before one endogenizes LP supply or solves a full LP-supply equilibrium.

6 Pricing Extension Under a Projected Cox Representation

This section studies pricing implications of the equilibrium predation mechanism. The baseline deterrence results of Sections 3–5 do not require asset pricing. Here we take the equilibrium physical law under $\mathbb{P}$ as given — in particular, the projected predation compensator λ_t^I and the projected jump term $\gamma^I(P_{t-}, \theta_t^\star)$ — and ask how claims exposed to predation-induced jumps are priced in an incomplete-market environment. The priced object may be interpreted as a claim on the AMM state, or equivalently as LP cash flows whose value is exposed to predation-induced execution jumps.

6.1 Physical Dynamics Under the Projected Predation Law

Under $\mathbb{P}$, the equilibrium price follows

$$dP_t = k(L_t)P_t\theta_t^\star\,dt + \sigma_P P_t\,dW_t + \gamma^I(P_{t-},\theta_t^\star)\,dN_t,$$

where W_t is a $\mathbb{P}$-Brownian motion, $\theta_t^\star = \Theta^\star(t,P_t;V,L_t)$, and N_t is a counting process with compensator

$$\int_0^t \lambda_s^I\,ds.$$

Under Assumption 3, this is a projected Cox representation of equilibrium predation under the insider/public filtration. The QVI determines the state dependence of the underlying structural hazard; the reduced-form pricing object is the projected compensator λ_t^I.

6.2 Pricing Kernel and Measure Change for Predation Jumps

We construct an equivalent martingale measure $\mathbb{Q}\sim\mathbb{P}$ under which discounted traded claims are local martingales. Because jumps make the market incomplete, $\mathbb{Q}$ is not unique; this section selects one pricing measure through an LP-based stochastic discount factor. Let $B_t = e^{rt}$ be the money-market account and let $\Lambda_t = d\mathbb{Q}/d\mathbb{P}|_{\mathcal{F}_t}$ denote the Radon–Nikodym density.

Let ξ_t be the market price of diffusive risk and let $\kappa_t > 0$ be a predictable jump-risk premium multiplier. Define

$$\Lambda_t = \exp\left\{-\int_0^t \xi_s\,dW_s - \frac{1}{2}\int_0^t \xi_s^2\,ds\right\}\cdot\prod_{i=1}^{N_t}\kappa_{\tau_i}\cdot\exp\left\{-\int_0^t(\kappa_s-1)\lambda_s^I\,ds\right\}.$$

Equivalently, using the compensated process

$$\widetilde{N}_t := N_t - \int_0^t \lambda_s^I\,ds,$$

$$\frac{d\Lambda_t}{\Lambda_{t-}} = -\xi_t\,dW_t + (\kappa_t-1)\,d\widetilde{N}_t.$$

Under the usual boundedness and integrability conditions on ξ_t, κ_t, and λ_t^I, Λ is a strictly positive $\mathbb{P}$-martingale and $\mathbb{Q}$ is well-defined.

Intensity transform under $\mathbb{Q}$. By Girsanov's theorem for point processes,

$$\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I.$$

Thus the pricing measure loads additional weight on predation states whenever $\kappa_t > 1$.

This pricing construction should be read as a pricing extension built on the equilibrium physical law under $\mathbb{P}$. It is not obtained by identifying the compensator with an inverse expected hitting time. The relevant physical object is the projected compensator λ_t^I defined in Sections 4–5.

6.3 Selecting the Jump Risk Premium: The LP Channel

In an incomplete market, the equivalent martingale measure is not unique. We select a pricing measure using a representative LP stochastic discount factor, interpreted as a reduced-form device for pricing claims exposed to predation-induced cash-flow losses. This is a selection rule for $\mathbb{Q}$; it is not a claim that the AMM invariant itself is set by a representative dealer in the same way as in a complete-market dealer model.

Assume a representative LP with CRRA utility over wealth W:

$$U(W) = \frac{W^{1-\eta}}{1-\eta}, \qquad \eta > 0.$$

Let a predation event at time τ reduce LP wealth by $\Delta_{\text{LP}} > 0$:

$$W_{\tau+} = W_{\tau-} - \Delta_{\text{LP}}.$$

Pricing-kernel condition (marginal utility ratio). The jump risk premium ratio is

$$\kappa_\tau = \frac{U'(W_{\tau-} - \Delta_{\text{LP}})}{U'(W_{\tau-})} = \left(\frac{W_{\tau-} - \Delta_{\text{LP}}}{W_{\tau-}}\right)^{-\eta} = \left(1 - \frac{\Delta_{\text{LP}}}{W_{\tau-}}\right)^{-\eta}.$$

Proposition 4 (Predation premium)**.** *If $\Delta_{\text{LP}} > 0$ and $\eta > 0$, then $\kappa_t > 1$. For small losses,*

$$\kappa_t \approx 1 + \eta \frac{\Delta_{\text{LP}}}{W_t}.$$

Accordingly, the LP channel provides one economically disciplined selection of the jump-risk premium multiplier κ_t within the family of admissible equivalent martingale measures.

6.4 Priced Claims Under the Selected Measure

Fix a maturity T and consider a claim with terminal payoff

$$H = \Phi(P_T, L_T)$$

and cumulative dividend process

$$D_t = \int_0^t \delta_s \, ds,$$

where Φ is $\mathcal{F}_T^I$-measurable and

$$\mathbb{E}^{\mathbb{Q}}\left[e^{-2rT}|H|^2 + \int_0^T e^{-2rs}|\delta_s|^2 \, ds\right] < \infty.$$

Its ex-dividend price under the selected measure is

$$C_t := \mathbb{E}^{\mathbb{Q}}\left[\int_t^T e^{-r(s-t)}\delta_s \, ds + e^{-r(T-t)}H \,\middle|\, \mathcal{F}_t^I\right].$$

Equivalently, the discounted cum-dividend value

$$M_t := e^{-rt}C_t + \int_0^t e^{-rs}\delta_s \, ds$$

is a $\mathbb{Q}$-martingale.

Proposition 5 (Selected-measure pricing of predation-exposed claims)**.** *Under the integrability condition above, C_t is well-defined and M_t is a $\mathbb{Q}$-martingale. If, in addition, $C_t = c(t, P_t, L_t)$ for a sufficiently regular Markov function c, the dividend flow is absolutely continuous, and the exogenous liquidity state L_t has $\mathbb{Q}$-generator $\mathcal{L}_L^{\mathbb{Q}}$, then c satisfies the valuation equation*

$$0 = c_t + \mathcal{L}_L^{\mathbb{Q}}c + \mu_t^{\mathbb{Q}}p\, c_p + \frac{1}{2}\sigma_P^2 p^2 c_{pp} + \lambda_t^{\mathbb{Q}}\Big(c\big(t, p + \gamma^I(p, \theta_t^\star), L_t\big) - c(t, p, L_t)\Big) - r\, c + \delta_t,$$

with terminal condition

$$c(T, p, L) = \Phi(p, L)$$

whenever the payoff is Markov in the priced state.

This proposition makes explicit what is being priced: a predation-exposed claim on the AMM state or, equivalently, LP cash flows whose continuation value loads on the same projected jump law as the baseline equilibrium.

6.5 Predation and Jump-Driven Tail Thickness Under $\mathbb{Q}$

Let log return be

$$R_t = \log \frac{P_t}{P_0}.$$

Using Itô's formula with jumps, under $\mathbb{Q}$,

$$dR_t = \left(\frac{\mu_t^{\mathbb{Q}}}{P_t} - \frac{1}{2}\sigma_P^2 \right) dt + \sigma_P \, dW_t^{\mathbb{Q}} + J_t \, dN_t, \qquad J_t := \log \left(1 + \frac{\gamma_t^I}{P_{t-}} \right),$$

where $\gamma_t^I = \gamma^I(P_{t-}, \theta_t^\star)$ and N_t has $\mathbb{Q}$-intensity $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$.

Define the jump-driven fourth-cumulant component

$$K_4^{\text{jump}}(T) := \mathbb{E}^{\mathbb{Q}} \left[\int_0^T \lambda_t^{\mathbb{Q}} \, \mathbb{E}^{\mathbb{Q}}[J_t^4 \mid \mathcal{F}_{t-}] \, dt \right].$$

Define also the jump-driven second-cumulant component

$$K_2^{\text{jump}}(T) := \mathbb{E}^{\mathbb{Q}} \left[\int_0^T \lambda_t^{\mathbb{Q}} \, \mathbb{E}^{\mathbb{Q}}[J_t^2 \mid \mathcal{F}_{t-}] \, dt \right],$$

and let

$$V_d(T) := \sigma_P^2 T$$

denote the diffusive variance contribution. In the short-horizon / small-jump regime where $K_2^{\text{jump}}(T) = O(\varepsilon)$ and $K_4^{\text{jump}}(T) = O(\varepsilon)$ while $V_d(T) \geq \underline{v} > 0$ is unchanged at first order, denominator effects enter only at higher order.

Proposition 6 (Leading-order jump-driven tail thickness under $\mathbb{Q}$)**.** *Assume the regime above. Then the excess kurtosis of R_T satisfies*

$$\text{ExKurt}(R_T) = \frac{K_4^{\text{jump}}(T)}{V_d(T)^2} + O(\varepsilon^2).$$

Consequently, along any parameter perturbation that leaves $V_d(T)$ unchanged at first order, any first-order increase in one or more of the following objects raises jump-driven

tail thickness under $\mathbb{Q}$:

- *the pricing multiplier κ_t,*
- *the physical projected predation intensity λ_t^I,*
- *the conditional jump magnitude distribution through $\mathbb{E}^{\mathbb{Q}}[J_t^4 \mid \mathcal{F}_{t-}]$.*

In particular, stronger predation exposure and stronger pricing of predation states jointly amplify the contribution of jumps to higher-order return moments.

Remark 4 (Scope)**.** *This proposition is a leading-order comparative static around a fixed diffusive benchmark, not an unrestricted global monotonicity result for every endogenous perturbation. It is about tail thickness / leptokurtosis, not about directional crash asymmetry. A separate skewness or downside-tail result would require additional restrictions on the sign and state dependence of the jump distribution.*

6.6 Conclusion of Section 6

Under the projected physical law, MEV acts as an endogenous generator of priced jump risk.

- **Physical mechanism:** strategic interaction generates state-dependent predation exposure through the projected intensity λ_t^I and projected jump sizes $\gamma^I(P_{t-}, \theta_t^\star)$.
- **Pricing channel:** an LP-selected equivalent martingale measure loads additional weight on predation states through $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$.
- **Asset-pricing implication:** predation can increase jump-driven tail thickness because both the physical predation channel and the pricing of predation states contribute to higher-order return moments.

7 Numerical Illustration

7.1 Purpose, Parameter Anchoring, and Numerical Implementation

This section provides numerical illustrations of the equilibrium implied by the coupled insider HJB and bot QVI system. The objective is not quantitative fit, forecasting, or

structural estimation. Rather, the numerical analysis serves three standard purposes in continuous-time market microstructure and impulse-control theory: (i) to visualize equilibrium policy functions and the implied safe-harbor / no-trade region, (ii) to document comparative statics with respect to economically interpretable parameters (e.g., risk aversion and detection technology), and (iii) to connect micro-level strategic interaction to distributional implications for returns under the pricing measure (Øksendal and Sulem 2019; Fleming and Soner 2006).

This section does not aim to empirically identify MEV deterrence; rather, it provides equilibrium diagnostics showing that the model's predation primitives endogenously implement a Becker-style deterrence mechanism—an action-dependent probability of predation and a state-contingent penalty—thereby generating a safe-harbor no-trade region, an intervention boundary for the bot, and priced tail risk consistent with a private-deterrence interpretation.

7.1.1 Why numerical illustration (and why not a full empirical exercise)?

We do not attempt a full empirical implementation for two methodological reasons.

First, the model generates endogenously state-dependent predation risk: both the arrival intensity and the effective jump/loss objects are induced by strategic interaction and, in the bot problem, by partial observation and impulse timing. Translating such objects into a data-based estimation framework requires additional measurement and identification structure—e.g., mapping observed on-chain ordering outcomes to latent beliefs and to counterfactual "no-predation" execution costs—beyond the scope of the present paper. Existing measurement work documents MEV activity at scale, including sandwich attacks and their evolution across Ethereum market design regimes, but those measurements do not by themselves pin down the structural primitives of a coupled HJB–QVI equilibrium.

Second, the theoretical contribution is deliberately general in reduced-form predation primitives: the analysis imposes shape restrictions (monotonicity/convexity/Lipschitz regularity) to establish well-posedness and equilibrium existence. Committing to a specific parametric structure suitable for estimation would shift the paper away from its main contribution—unifying impulse-control predation with execution and pricing—and toward dataset-specific empirical choices. Accordingly, we follow the common practice

(Crandall, Ishii, and Lions 1992; Øksendal and Sulem 2019) in theory papers: we select transparent functional forms consistent with the maintained assumptions, choose order-of-magnitude reasonable baselines, and emphasize robust qualitative patterns across economically relevant parameter ranges.

7.1.2 Anchoring without calibration: how external evidence is used

Although the analysis is not calibrated, we use external evidence to anchor the *order of magnitude* of a small set of key parameters.

First, large-scale empirical and measurement studies document that MEV activity is pervasive on Ethereum and that sandwich and arbitrage strategies represent economically meaningful components of MEV extraction. We therefore choose baseline predation intensities such that the expected number of predation events over the normalized horizon is *order one* at typical execution speeds, rather than vanishingly small or mechanically overwhelming.

Second, empirical characterizations of DEX trading costs show that transaction costs vary sharply with trade characteristics and that adversarial components related to transaction reordering can be economically meaningful, particularly for larger swaps or more speculative assets. We use these results to choose jump magnitudes and discrete loss parameters that are large enough to generate visible strategic effects (and tail implications) but small enough that diffusion remains non-degenerate.

Third, in interpreting policy levers, we map changes in fixed per-attack costs and fee-market conditions to institutional fee-market design (Buterin et al. 2019; Rondelet and Kilbourn 2023) such as EIP-1559 and to evidence on how that mechanism affects fee dynamics and user experience. For privacy/noise levers, we draw on the mempool privacy literature that emphasizes both MEV mitigation potential and market-structure tradeoffs.

7.1.3 Normalizations and truncated state space

To separate mechanisms from units, we normalize the execution horizon to $T = 1$, so intensities may be interpreted as expected event counts over the horizon, and we normalize the initial price to $p_0 = 1$, so jump sizes and slippage components can be interpreted in percentage terms. We compute the equilibrium on truncated compact state and control

domains, consistent with the compact-domain approach used in the existence proof and standard in viscosity/impulse-control computations. Truncation is a numerical device rather than an economic assumption.

7.1.4 Numerical specification of reduced-form predation primitives

The theoretical analysis allows a broad admissible class of reduced-form predation primitives $g = (\lambda, \gamma, T)$. For numerical illustration, we adopt a parsimonious specification that satisfies the maintained shape restrictions and keeps the computation stable under CARA preferences.

Predation intensity (visibility/detection technology). We set

$$\lambda(\theta) = \lambda_0 + \lambda_1 |\theta|^{\beta}, \quad \beta \geq 1, \tag{11}$$

where λ_0 captures background predation risk and λ_1 scales the sensitivity of predation risk to execution aggressiveness (our "technology/visibility" lever). This functional form is increasing and convex in $|\theta|$, reflecting the idea that larger and faster executions are more detectable in transparent transaction pools.

Attack-induced price displacement (jump size). We specify a scale-free jump mapping,

$$\gamma(p, \theta) = \eta \, p \, \theta, \tag{12}$$

where $\eta > 0$ governs the magnitude of predation-induced displacements. This form preserves price-level scale invariance and captures the empirically supported idea that adversarial reordering costs and price distortions become more pronounced for larger swaps (Adams et al. 2025).

Discrete wealth loss at attack. We adopt a strictly convex quadratic loss,

$$T(\theta) = \frac{\tau}{2}\theta^2, \quad \tau > 0, \tag{13}$$

which stabilizes numerics when combined with exponential CARA jump penalties and captures increasing marginal harm as execution becomes more aggressive. This reduced-

form proxy is consistent with evidence that the adversarial component of DEX trading costs rises with trade size and reordering vulnerability.

Baseline magnitudes and scanned ranges. Baseline values are chosen so that: (i) the no-predation benchmark implies active trading, (ii) predation risk is economically meaningful but does not collapse trading everywhere, and (iii) jump effects are visible in return tails without mechanically overwhelming diffusion. We report the baseline computational settings in Appendix 8.2, Table D.1. The ranges are selected to reflect the order-of-magnitude evidence on MEV prevalence and on the size of DEX trading cost components, while maintaining non-degenerate diffusive dynamics.

7.1.5 Numerical solution: HJB, QVI, and fixed-point iteration

Equilibrium is computed as a fixed point between the insider's optimal continuous execution policy and the bot's optimal impulse policy (and the induced predation threat). The algorithm proceeds in three steps.

Step 1 (Insider HJB). Given a candidate reduced-form triple $g = (\lambda, \gamma, T)$, we solve the insider's finite-horizon HJB for the certainty equivalent $J(t, p)$ and the feedback policy $\theta^*(t, p)$. The HJB contains a non-local term through $J(t, p + \gamma)$ and a CARA-induced exponential jump penalty. We discretize the time–state domain and apply policy iteration with interpolation for the non-local term.

Step 2 (Bot QVI). Given the implied insider behavior summarized by the belief-slope proxy used in the strategic linearization, we solve the bot's stationary impulse-control problem as a QVI. The solution yields an intervention region (predation boundary) and an optimal impulse intensity conditional on the bot's state. We implement the QVI via a penalty formulation combined with policy iteration; the fixed per-attack cost ensures no accumulation of impulses (no "Zeno" behavior).

Step 3 (Fixed point). We update the reduced-form objects induced by the bot's best response and iterate until convergence. To stabilize the mapping, we employ damping:

$$g_{n+1} \leftarrow (1 - \omega)\, g_n + \omega\, \mathscr{T}(g_n), \quad \omega \in (0, 1), \tag{14}$$

where $T(\cdot)$ denotes the best-response pipeline defined by the coupled HJB–QVI system. Schauder's theorem guarantees existence of a fixed point but does not imply convergence of this particular damped iteration, so the outer loop should be read as a numerical solver rather than as part of the theorem. In any implementation one should report the chosen damping parameter, the residual sequence $r_n := \|g^{(n+1)} - g^{(n)}\|_\infty$, and robustness across initializations. We declare convergence when the sup-norm distance between successive iterates falls below a pre-specified tolerance on the truncated domain.

Having specified the numerical primitives and solution method, we first illustrate the baseline equilibrium implied by the insider HJB. Figure 4 plots the baseline trading rule $\theta^*(0, p)$ as a function of mispricing $m = V - p$ and the associated continuation value $J(0, p)$. Relative to the no-predation benchmark, MEV risk compresses the execution policy and generates a near-zero safe-harbor / no-trade region, while eroding continuation value. This baseline visualization motivates the comparative statics in Section 7.2.

Economically, the safe harbor in Figure 4 represents the deterrence effect of projected predation exposure. The "Value Erosion" in Panel B is the certainty-equivalent shadow tax associated with expected predation losses, and the safe harbor marks the region in which marginal informational rents are dominated by those losses.

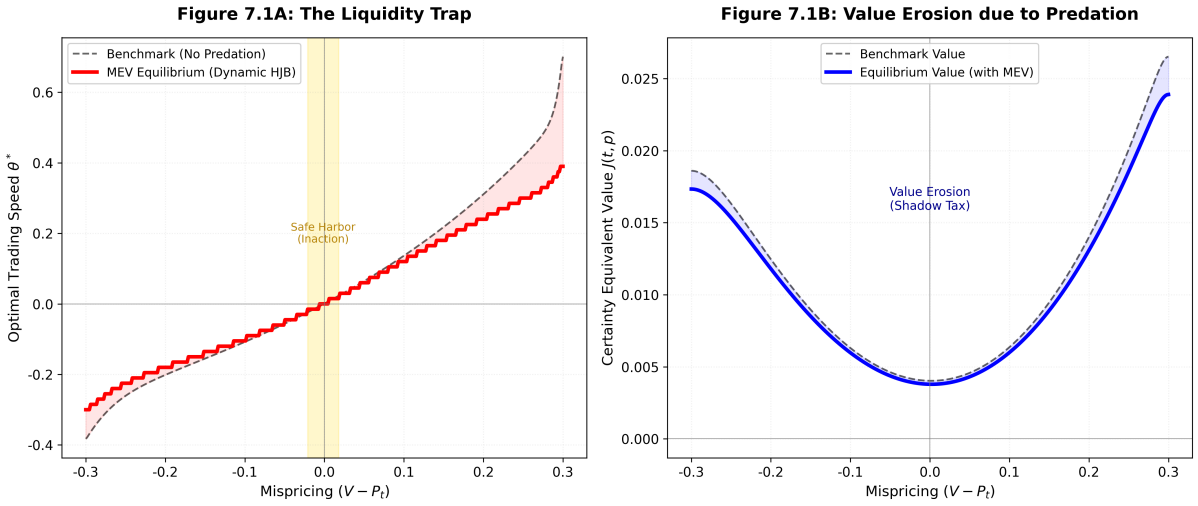


Figure 4: Baseline Equilibrium. Panel A shows the optimal trading speed θ^* versus mispricing, illustrating the safe harbor (inaction) region. Panel B shows the certainty equivalent value $J(t, p)$, highlighting value erosion (shadow tax) relative to the benchmark.

7.1.6 Robustness and interpretation

Because the numerical analysis is intended to illustrate mechanisms rather than fit data, we emphasize qualitative robustness. Appendix 8.2 reports (i) grid refinement and truncation checks, (ii) fixed-point convergence diagnostics, and (iii) sensitivity analysis over η, τ, β in addition to the main comparative statics in α, λ_1. Across these variations, the main patterns highlighted in Figures 4–8 remain: predation risk compresses the insider's execution rule into a safe-harbor region; higher fear (larger α) and stronger detection technology (larger λ_1) widen that region; higher effective fixed costs expand the safe harbor; and increased privacy/noise mitigates predation at the cost of slower information revelation—consistent with fee-market design considerations and the mempool privacy literature (Buterin et al. 2019; Rondelet and Kilbourn 2023).

7.2 Comparative Statics: Fear and Detection Technology

Figure 4 establishes the baseline mechanism: predation risk compresses the insider's equilibrium execution rule and creates a near-zero trading region around $m = V - p = 0$ (the safe-harbor / no-trade region), while reducing continuation value $J(0, p)$. Building on this baseline, this subsection quantifies how the safe harbor widens when we vary two economically interpretable parameters: (i) the insider's fear/risk aversion α, which amplifies the CARA jump penalty, and (ii) the detection-technology scale λ_1 in the intensity function $\lambda(\theta)$, which steepens the state-dependent predation threat.

The objective is to provide visual evidence for the shadow-tax channel implied by the insider HJB, where predation enters through $\lambda(\theta)$ and the exponential term $e^{\alpha\Psi}$, with $\Psi(\theta, p) = T(\theta) + J(t, p) - J(t, p + \gamma(p, \theta))$.

7.2.1 Measuring the Safe-Harbor Region

Let mispricing be $m \equiv V - p$. For each parameter pair α, λ_1, we solve the insider HJB on the same truncated state grid used in the baseline and recover the equilibrium feedback rule $\theta^*(t, p)$. We evaluate the policy at $t = 0$ and define the **safe-harbor (near-zero trading) region** as

$$S(\alpha, \lambda_1) = \{m : |\theta^*(0, p = V - m)| \leq \varepsilon\}, \tag{15}$$

where $\varepsilon > 0$ is a small threshold. We summarize the safe harbor by the **safe-harbor width**

$$W(\alpha, \lambda_1) \equiv \text{Lebesgue-measure}(S(\alpha, \lambda_1)), \tag{16}$$

i.e., the length of the mispricing interval around zero over which optimal execution is negligible. This statistic matches the interpretation of the safe-harbor mechanism in the model: predation risk suppresses trading for small informational advantages, expanding the region in which the insider rationally waits rather than executes.

7.2.2 Comparative-Statics Design

We vary α and λ_1 on a rectangular grid while holding all other primitives fixed at the baseline values reported in Table D.1. For each grid point, we (i) solve the HJB for $J(t, p)$ and $\theta^*(t, p)$, and (ii) compute $W(\alpha, \lambda_1)$ from the $t = 0$ policy. Figure 5 reports the heatmap of $W(\alpha, \lambda_1)$ across the parameter grid and representative cross-sections of W as a function of α for selected values of λ_1.

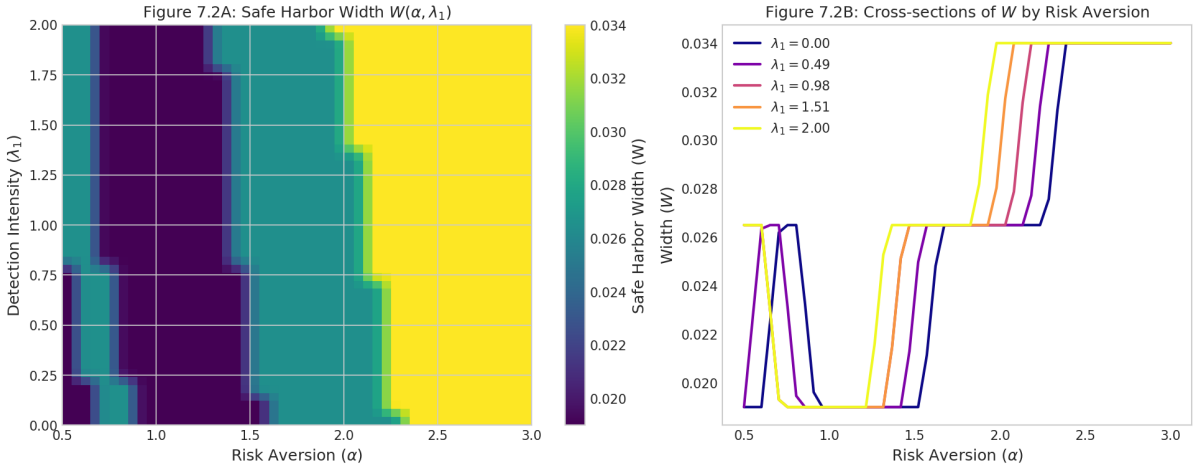


Figure 5: Safe Harbor Width. Panel A plots a heatmap of the safe-harbor width $W(\alpha, \lambda_1)$. Panel B reports cross-sections of W as a function of α. Higher risk aversion α and stronger detection technology λ_1 expand W, indicating a wider region of mispricing in which optimal execution is negligible.

7.2.3 Results and Mechanism

Two patterns are robust. First, **higher fear** (larger α) expands the safe-harbor region: $W(\alpha, \lambda_1)$ increases monotonically in α. Intuitively, higher risk aversion amplifies the effective marginal penalty from predation through the CARA jump term $e^{\alpha\Psi}$, increasing

the shadow-tax wedge and pushing the optimal policy toward inaction for a wider set of small mispricings.

Second, **stronger detection technology** (larger λ_1) also expands the safe-harbor region. A steeper intensity function $\lambda(\theta)$ makes the marginal predation cost rise faster with execution aggressiveness, discouraging trading precisely where informational rents are modest. Together, these forces deepen the safe harbor by widening the state region in which the insider optimally delays execution.

7.2.4 Robustness

We verify that the comparative statics are robust to reasonable variations in the threshold ε, grid refinement, and truncation bounds. Additional diagnostics (including convergence and discretization checks) are reported in Appendix 8.2.

7.3 Predation Boundary: Gas Costs and Observation Noise

This subsection illustrates the bot's intervention threshold ("predation boundary") implied by the impulse-control structure of the MEV searcher's problem. In the model, the bot solves a stationary impulse-control problem characterized by a QVI, which partitions the state space into a continuation (wait) region and an intervention (attack) region. Fixed per-attack costs (e.g., priority fees and related costs) rule out Zeno-type accumulation of impulses and therefore generate a non-trivial inaction region.

We focus on two comparative statics that directly correspond to policy and market-design levers: the fixed per-attack cost C_{gas} and the noise level σ_Y in the bot's observation process. Higher C_{gas} raises the hurdle rate for profitable intervention, shifting the predation boundary outward and enlarging the set of states in which the bot optimally waits. Likewise, higher observation noise reduces the effective signal-to-noise ratio of order flow, slows inference, and delays entry into the attack region, again expanding the bot's inaction region.

Figure 6 visualizes these predictions by plotting the attack threshold as a function of C_{gas} for several values of σ_Y (Panel A) and as a heatmap over C_{gas}, σ_Y (Panel B). The key implication is monotone: increases in either the fixed cost of predation or the noise that obscures detection enlarge the safe-harbor region for traders by pushing the predation boundary outward. The plotted boundary is a local asymptotic proxy obtained

from a quadratic expansion of the intervention gain around the QVI free boundary; the curvature parameter in that expansion is recorded explicitly in Appendix 8.2.

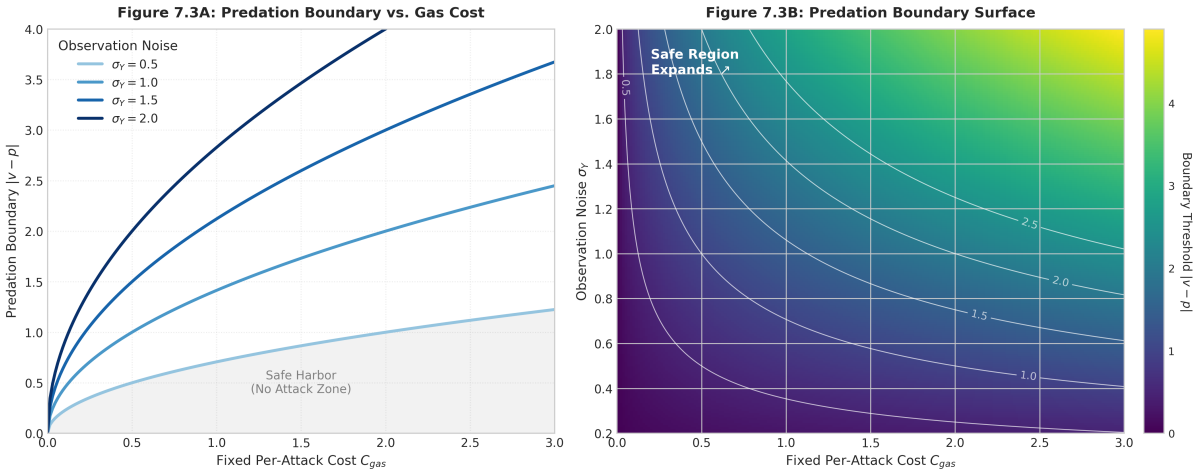


Figure 6: Predation boundary comparative statics. The plotted surface is a leading-order local-boundary proxy obtained from a quadratic expansion of the intervention gain; it is designed to visualize the monotone effects of fixed attack costs and observation noise, not to replace a full numerical solution of the QVI free boundary.

To keep the illustration transparent and computationally light, Figure 6 uses the leading-order local-boundary proxy rather than a direct numerical solution of the QVI boundary. The figure is intended to visualize the model's monotone comparative statics in C_{gas} and σ_Y.

Connection to trader safe harbor (link to Figure 5). An outward shift in the bot's predation boundary translates into a larger safe-harbor region for traders. When the bot requires a larger opportunity to intervene, small-mispricing states are less exposed to predation risk, which weakens the effective predation wedge (the "shadow tax") that enters the insider's HJB via the predation term. As a result, the near-zero trading region around $m \approx 0$ becomes more defensible—precisely the logic behind the safe-harbor statistic $W(\alpha, \lambda_1)$ studied in Figure 5. In this sense, C_{gas} (fee-market conditions that raise the fixed cost of predation) and σ_Y (noise/opacity that lowers detectability) act as market-design levers that mitigate predation pressure by expanding the set of states in which the bot rationally waits—though higher σ_Y may also imply slower information revelation.

7.4 The Market (Fat Tails): Safe-Harbor Peak and MEV-Driven Tail Risk

This subsection connects the microstructure mechanisms highlighted in Sections 7.2 and 7.3 to a market-level implication: **endogenous leptokurtosis** in equilibrium returns. In the model, strategic predation affects prices through two linked channels. First, the insider's equilibrium execution rule $\theta^*(t, p)$ is compressed around small mispricing states (the safe-harbor / no-trade region), reducing typical price impact for many realizations. Second, predation generates **state-dependent jumps** whose intensity and size depend on equilibrium behavior, and these jump arrivals are **amplified under the pricing measure** $\mathbb{Q}$ via the jump risk premium channel. Together, these forces produce a return distribution with **a sharper center and fatter tails** than a Gaussian benchmark.

7.4.1 Pricing-measure amplification and the jump channel

Under the equilibrium reduced-form jump-diffusion, the AMM price inherits a predation jump component whose arrival intensity is increasing in the insider's execution aggressiveness. In the reduced-form representation used for the numerical illustrations, predation arrivals are summarized by $\lambda(\theta)$ and jump sizes by $\gamma(p, \theta)$. Under the pricing measure $\mathbb{Q}$, the projected compensator transforms so that the pricing intensity becomes $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$, where $\kappa_t > 1$ is the jump risk premium ratio. This implies that risk-neutral returns can exhibit materially heavier tails even when diffusive volatility is unchanged: jump arrivals become effectively more frequent under $\mathbb{Q}$, and jump magnitudes are endogenously linked to equilibrium trading behavior through $\gamma(P_{t-}, \theta_t^*)$.

7.4.2 A "money plot": standardized log-density with tail inset

To make the distributional implication visually immediate, we use a **log-density plot** (rather than a pure tail plot). This choice is deliberate: it shows **both** (i) the **center** of the distribution—where safe-harbor dynamics concentrate mass near zero—and (ii) the **tails**—where predation jumps increase the probability of extreme outcomes. We therefore

plot the density of standardized terminal log returns

$$z = \frac{r - \bar{r}}{\sigma}, \quad r = \log \frac{P_T}{P_0}, \tag{17}$$

on a log y-axis, comparing the simulated MEV jump-diffusion under $\mathbb{Q}$ to a matched normal distribution with the same mean and variance. An inset zooms into the tail region $|z| > q_{0.98}$, directly exposing tail amplification.

Figure 7.4: Endogenous Leptokurtosis and Fat Tails

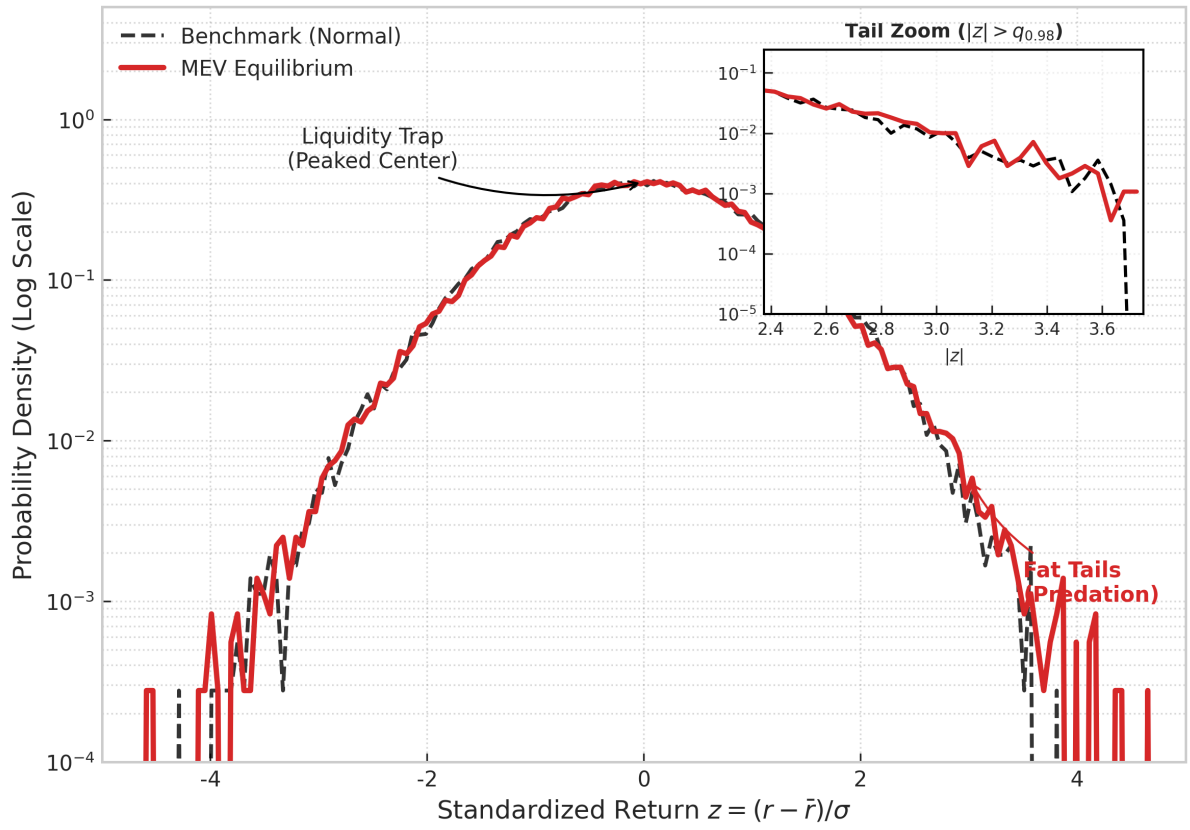


Figure 7: Money Plot: Safe-Harbor Peak and MEV-Driven Fat Tails. Standardized log-density of terminal log returns $z = (r - \bar{r})/\sigma$, comparing a simulated equilibrium MEV jump-diffusion under the pricing measure $\mathbb{Q}$ (solid) to a matched normal distribution (dashed). The inset zooms into the tail region $|z| > q_{0.98}$.

7.4.3 Results: sharp center and fat tails

Figures 7 and 8 deliver a single unified visual message:

Sharper center (mixture signature with equilibrium compression). Relative to the matched normal benchmark, the MEV jump-diffusion exhibits **excess mass near**

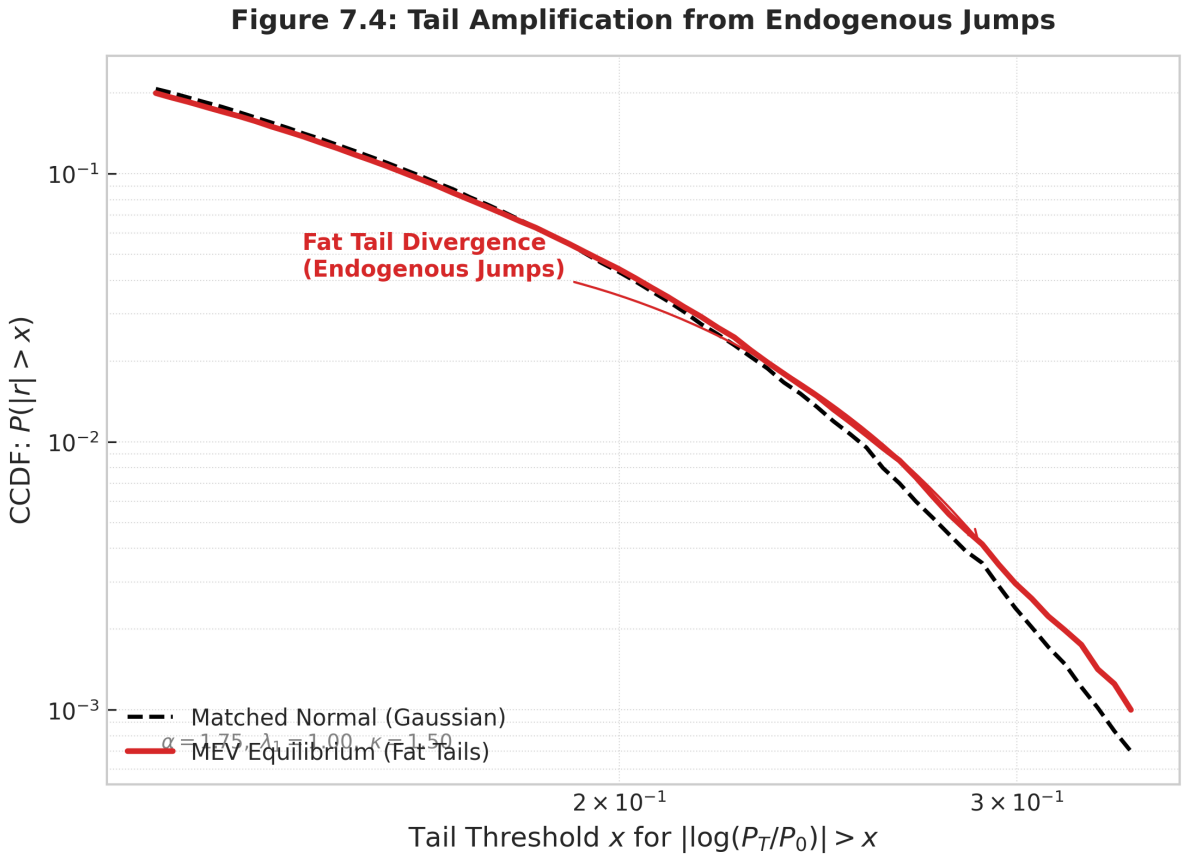


Figure 8: Tail Amplification from Endogenous Jumps. Endogenous Leptokurtosis and Fat Tails.

$z = 0$. This should be read as a joint consequence of the endogenous jump-mixture structure and the equilibrium compression of trading around small mispricing states. Rare predation jumps increase overall kurtosis, while the safe-harbor mechanism shifts more paths into low-impact continuation states. The center peak is therefore not a stand-alone identification of the safe harbor; rather, the safe-harbor mechanism affects the decomposition through the equilibrium policy $\theta^\star$.

Fatter tails (predation jump signature). In the tails, the MEV curve lies above the matched normal on the log-density scale. The inset makes this stark: for $|z| > q_{0.98}$, the MEV jump-diffusion assigns substantially higher probability to extreme outcomes. This tail amplification reflects (i) endogenous jump arrivals whose intensity rises with equilibrium behavior through $\lambda(\theta)$, and (ii) the pricing-measure amplification $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$, which increases the effective frequency of predation events under $\mathbb{Q}$.

7.5 Summary of Numerical Mechanisms

Section 7 provides a transparent numerical illustration of the equilibrium mechanisms implied by the coupled insider HJB and bot impulse-control (QVI) structure. The goal is not calibration or estimation, but to **visualize equilibrium policy functions**, **document comparative statics in economically interpretable levers**, and **connect strategic predation to distributional implications for returns**.

(i) One mechanism, four views: Safe Harbor → Mechanism → Bot → Market. Across the four main figures, a single economic object reappears in different guises: **predation risk behaves like a state-dependent "shadow tax" on execution** that suppresses trading where informational rents are modest, while simultaneously creating an endogenous rare-event channel in prices.

(ii) What the numerics add beyond the theory. The theory establishes well-posedness and equilibrium existence under broad shape restrictions; the numerics show **how large and how structured** the implied regions and wedges can be on a transparent baseline. In particular, the figures make three qualitative implications hard to miss: predation risk produces a **state-contingent inaction set** (the safe harbor), that region is

monotonically widened by higher fear and stronger detection technology, and market-design levers that raise predation costs or reduce detectability can **expand safe-harbor protection**, with potential tradeoffs in information revelation.

(iii) Limitations and how to read the illustrations. Because the exercise is intentionally not calibrated, the numeric magnitudes should be interpreted as **illustrative** rather than empirical claims. The computational domain is truncated and controls are discretized, consistent with standard viscosity/impulse-control numerics; robustness is assessed by refinement and truncation checks rather than statistical fit. The key takeaways are therefore **qualitative patterns and mechanisms**, not point estimates.

8 Conclusion: MEV as a Private Deterrence Technology

8.1 Core Thesis and Big-Picture Contribution

This paper's central claim is institutional but deliberately narrower than "code replaces law": transparent DeFi microstructure can implement a **private deterrence technology**. Sandwich-capable MEV searchers impose immediate, state-dependent execution losses on aggressive informational trading, so deterrence arises from transaction ordering and predation exposure rather than from delayed public sanction.

We formalize this thesis by embedding predation risk into a continuous-time execution problem and coupling it to a strategic predator's impulse-control problem. On the trader side, the HJB is written in terms of projected predation intensity and projected jump/loss objects under the insider's information set. On the bot side, the QVI disciplines the structural hazard, intervention region, and comparative statics of predation exposure. Together, these objects deliver a tractable coupled equilibrium without relying on the inverse of an expected boundary-hitting time as the compensator.

8.2 Mechanism: The Shadow Tax and the Safe-Harbor Region as "Deterrence"

The model delivers a sharp mechanism: **MEV predation suppresses informed execution by imposing a composite marginal predation cost that scales exponentially with the jump-loss object**. In the insider's HJB, the key term is the CARA-induced exponential penalty $e^{\alpha\Psi(\theta,p)}$, where $\Psi(\theta,p) = T(\theta)+J(t,p)-J(t,p+\gamma(p,\theta))$. This generates a generalized "MEV shadow tax" that decomposes into (i) visibility/detection (via $\lambda'(\theta)$), (ii) sandwich geometry (via $T'(\theta)$), and (iii) endogenous state collapse (via the non-local value erosion term).

The outcome is a **state-contingent safe-harbor region** around small mispricing in which optimal execution is negligible—precisely the economic signature of deterrence when expected predation losses dominate marginal informational rents. The transfer-neutral benchmark in Section 5.7 shows that this wedge can also generate a genuine welfare distortion rather than a pure redistribution. This deterrence logic is uniquely DeFi-native: aggressive trades are more detectable in transparent transaction pools, and that detectability raises projected predation exposure through $\lambda(\theta)$.

Finally, the same projected physical law supports a pricing extension in which predation-induced jumps contribute to jump-driven tail risk under an LP-selected pricing kernel. The paper's main message is therefore twofold: MEV can deter aggressive informed trading through a shadow-tax mechanism, and the same predation channel can leave priced distributional consequences in equilibrium returns.

References

Adams, Austin, Benjamin Y. Chan, Sarit Markovich, and Xin Wan. 2025. "Don't Let MEV Slip: The Costs of Swapping on the Uniswap Protocol." In *Financial Cryptography and Data Security: 28th International Conference, FC 2024*, 172–91. Lecture Notes in Computer Science. Cham: Springer.

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *The Quarterly Journal of Economics* 84 (3): 488–500.

Back, Kerry. 1992. "Insider Trading in Continuous Time." *The Review of Financial*

Studies 5 (3): 387–409.

Bahrani, Maryam, Pranav Garimidi, and Tim Roughgarden. 2024. "Transaction Fee Mechanism Design in a Post-MEV World." In *6th Conference on Advances in Financial Technologies (AFT 2024)*, Leibniz International Proceedings in Informatics (LIPIcs), vol. 316, 29:1–29:24. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

Bartoletti, Massimo, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. 2022. "A Theory of Automated Market Makers in DeFi." *Logical Methods in Computer Science* 18 (4).

Becker, Gary S. 1968. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76 (2): 169–217.

Biais, Bruno, Christophe Bisière, Matthieu Bouvard, and Catherine Casamatta. 2019. "The Blockchain Folk Theorem." *The Review of Financial Studies* 32 (5): 1662–1715.

Budish, Eric. 2025. "Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains." *The Quarterly Journal of Economics* 140 (1): 1–62.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *The Quarterly Journal of Economics* 130 (4): 1547–1621.

Buterin, Vitalik, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. 2019. "EIP-1559: Fee Market Change for ETH 1.0 Chain." *Ethereum Improvement Proposals*.

Catalini, Christian, and Joshua S. Gans. 2018. "Initial Coin Offerings and the Value of Crypto Tokens." *NBER Working Paper No. 24418*.

Chiu, Jonathan, Charles M. Kahn, and Thorsten V. Koeppl. 2022. "Grasping De(centralized) Fi(nance) Through the Lens of Economic Theory." *Bank of Canada Staff Working Paper* 2022-43.

Cong, Lin William, and Zhiguo He. 2019. "Blockchain Disruption and Smart Contracts." *The Review of Financial Studies* 32 (5): 1754–1797.

Cong, Lin William, Ye Li, and Neng Wang. 2021. "Tokenomics: Dynamic Adoption and Valuation." *The Review of Financial Studies* 34 (3): 1105–1155.

Cong, Lin William, Ke Tang, Yanxin Wang, and Xi Zhao. 2023. "Inclusion and Democratization Through Web3 and DeFi? Initial Evidence from the Ethereum Ecosystem." *NBER Working Paper No. 30949*.

Crandall, Michael G., Hitoshi Ishii, and Pierre-Louis Lions. 1992. "User's Guide to Viscosity Solutions of Second Order Partial Differential Equations." *Bulletin of the American Mathematical Society* 27 (1): 1–67.

Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. "Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability." In *2020 IEEE Symposium on Security and Privacy (SP)*, 910–927.

Fleming, Wendell H., and H. Mete Soner. 2006. *Controlled Markov Processes and Viscosity Solutions*. 2nd ed. New York: Springer.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53 (6): 1315–1335.

Lehar, Alfred, and Christine A. Parlour. 2025. "Decentralized Exchange: The Uniswap Automated Market Maker." *The Journal of Finance* 80 (1): 321–374.

Liu, Yukun, and Aleh Tsyvinski. 2021. "Risks and Returns of Cryptocurrency." *The Review of Financial Studies* 34 (6): 2689–2727.

Merton, Robert C. 1976. "Option Pricing When Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics* 3 (1–2): 125–144.

Milionis, Jason, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. "Automated Market Making and Loss-Versus-Rebalancing." *arXiv preprint*, arXiv:2208.06046.

Mohan, Vijay. 2022. "Automated Market Makers and Decentralized Exchanges: A DeFi Primer." *Financial Innovation* 8 (20): 1–20.

Øksendal, Bernt, and Agnès Sulem. 2019. *Applied Stochastic Control of Jump Diffusions*. 3rd ed. Cham: Springer.

Rondelet, Antoine, and Quintus Kilbourn. 2023. "Mempool Privacy: An Economic Perspective." *arXiv preprint*, arXiv:2307.10878.

Shleifer, Andrei, and Robert W. Vishny. 1997. "The Limits of Arbitrage." *The Journal of Finance* 52 (1): 35–55.

Weintraub, Ben, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. 2022. "A Flash(bot) in the Pan: Measuring Maximal Extractable Value in Private Pools." In *Proceedings of the ACM Internet Measurement Conference (IMC '22)*.

Appendix A. Existence of Markovian Equilibrium (Proof of Proposition 2)

Notation. To avoid conflict with the loss notation used elsewhere in the paper, we denote the fixed-point operator by $\mathscr{T}$. The reduced-form triple is $g = (\lambda, \gamma, T)$, where $\lambda(\theta)$ is the reduced-form predation intensity, $\gamma(p, \theta)$ is the reduced-form price displacement at an attack, and $T(\theta)$ is the reduced-form discrete wealth loss.

A.1 Compact truncation and admissible reduced-form class

To place the fixed-point argument on a compact convex set, we work on truncated compact domains for states and controls, consistent with compact-control assumptions and standard in viscosity/impulse-control existence proofs (e.g., [2, 3, 4]). Let the price domain be $\mathcal{P} = [\underline{p}, \overline{p}] \subset (0, \infty)$ and the insider control domain be $\Theta = [-\bar{\theta}, \bar{\theta}]$. For the bot's augmented state $\xi = (v, \Sigma, p, L)$, let $\mathcal{X} := \mathcal{V} \times \mathcal{S} \times \mathcal{P} \times \mathcal{L}$ be a compact product of intervals.

Definition A.1 (Admissible class $\mathcal{A}$). Fix constants $\bar{\lambda}, \bar{\gamma}, \bar{T} > 0$ and Lipschitz moduli $L_\lambda, L_\gamma, L_T > 0$. Let $\mathcal{A}$ be the set of all triples $g = (\lambda, \gamma, T)$ satisfying:

1. **(A1) Intensity regularity.** $\lambda : \Theta \to [0, \bar{\lambda}]$ is continuous and even, and its restriction to $[0, \bar{\theta}]$ is increasing and convex, and $\mathrm{Lip}(\lambda) \le L_\lambda$.

2. **(A2) Jump regularity.** $\gamma : \mathcal{P} \times \Theta \to \mathbb{R}$ is continuous, $\|\gamma\|_\infty \le \bar{\gamma}$, and uniformly Lipschitz:
$$\big|\gamma(p, \theta) - \gamma(p', \theta')\big| \le L_\gamma \Big(|p - p'| + |\theta - \theta'| \Big).$$

3. **(A3) Loss regularity.** $T : \Theta \to [0, \bar{T}]$ is continuous, convex, and $\mathrm{Lip}(T) \le L_T$.

4. **(A4) Primitive well-posedness.** Remaining primitives (e.g., $k(L)$, σ_P, and the bot's impulse primitives Γ and R) are continuous and bounded on truncated domains, and satisfy standard growth/regularity conditions ensuring dynamic programming and viscosity characterizations for the HJB/QVI below (see [3, 5, 4]).

Lemma A.1 (Compact convexity of $\mathcal{A}$)**.** *$\mathcal{A}$ is nonempty, convex, and compact under the uniform topology induced by the sup norm on the compact domains.*

Proof. Nonemptiness is immediate by constructing any triple satisfying the bounds and monotonicity/convexity conditions. Convexity follows since continuity and Lipschitz bounds are preserved under convex combinations, and convexity is preserved by mixing. Compactness follows from uniform boundedness and equicontinuity (uniform Lipschitz moduli) on compact domains, hence relative compactness by Arzelà–Ascoli. Finally, $\mathcal{A}$ is closed under uniform convergence because continuity, boundedness, Lipschitz bounds, evenness, monotonicity in $|\theta|$, and convexity in $|\theta|$ are preserved under uniform limits on compact domains; thus $\mathcal{A}$ is compact. □

A.2 Insider best response given $g \in \mathcal{A}$

Fix any $g = (\lambda, \gamma, T) \in \mathcal{A}$. Let $J^g(t,p)$ denote the insider's certainty-equivalent value. Under the reduced-form jump representation, J^g solves the HJB equation (as stated in the paper) on $[0,T] \times \mathcal{P}$ with terminal condition $J^g(T,p) = 0$. Writing only the key nonlocal jump term, define

$$\Psi^g(\theta,p) \;:=\; T(\theta) \;+\; J^g(t,p) \;-\; J^g\big(t, p+\gamma(p,\theta)\big).$$

The intensity $\lambda(\theta)$ enters through the CARA jump term $-\lambda(\theta)\frac{1}{\alpha}\left(e^{\alpha\Psi^g(\theta,p)}-1\right)$.

Define the nonlocal CARA jump operator

$$\mathcal{I}^g[J](t,p,\theta) := -\frac{\lambda(\theta)}{\alpha}\left(e^{\alpha[T(\theta)+J(t,p)-J(t,p+\gamma(p,\theta))]}-1\right).$$

Lemma A.2 (Uniform exponential integrability and bounded certainty equivalent)**.** *Under (A1)–(A4), admissible controls on the truncated domain generate terminal wealth with uniformly finite exponential moments on* $[0,T]$*. Consequently, the certainty equivalent* J^g *is bounded on* $[0,T] \times \mathcal{P}$*.*

Proof. On the truncated domain, p, θ, $k(L)$, $\lambda(\theta)$, and $T(\theta)$ are uniformly bounded. Hence the continuous part of the wealth process is bounded in absolute value by a constant multiple of T, while the jump loss is bounded above by $\bar{T}N_T$. Because the attack intensity is bounded by $\bar{\lambda}$, the counting process has finite exponential moments on finite horizons; equivalently, for every $c > 0$,

$$\mathbb{E}\big[e^{cN_T}\big] < \infty.$$

Therefore X_T admits finite exponential moments uniformly over admissible controls. Since CARA utility is defined on all of $\mathbb{R}$, the value function is finite, and the associated certainty equivalent is bounded above and below by finite constants depending only on the truncated primitives. □

Lemma A.3 (Properness of the nonlocal CARA operator)**.** *Fix $g \in \mathcal{A}$ and $\theta \in \Theta$. If $u - v$ attains a global maximum at (t, p), then*

$$u(t,p) - u\big(t, p + \gamma(p,\theta)\big) \geq v(t,p) - v\big(t, p + \gamma(p,\theta)\big).$$

Consequently,

$$\mathcal{I}^g[u](t,p,\theta) \leq \mathcal{I}^g[v](t,p,\theta).$$

Proof. If $u - v$ attains a global maximum at (t, p), then

$$u(t,p) - v(t,p) \geq u\big(t, p + \gamma(p,\theta)\big) - v\big(t, p + \gamma(p,\theta)\big),$$

which rearranges to the stated inequality for the nonlocal increments. The map

$$x \mapsto -\frac{\lambda(\theta)}{\alpha}\big(e^{\alpha[T(\theta)+x]} - 1\big)$$

is decreasing in x because $\lambda(\theta) \geq 0$ and the exponential is increasing, so the order of the increments is reversed under $\mathcal{I}^g$. □

Lemma A.4 (HJB well-posedness)**.** *Under (A1)–(A4), the insider HJB admits a unique bounded continuous viscosity solution $J^g \in C([0,T] \times \mathcal{P})$.*

Justification. Lemma A.2 gives boundedness of the certainty equivalent. Lemma A.3 shows that the jump term is a proper continuous nonlocal operator. The HJB is therefore a bounded state-constraint integro-HJB with continuous coefficients on a compact domain. Existence follows from the dynamic programming principle, while comparison and uniqueness follow from standard viscosity theory for second-order integro-differential equations with proper nonlocal terms; see [1] and [5]. On the truncated state domain, we interpret the equation in the *state-constraint viscosity* sense, so no explicit Dirichlet/Neumann boundary data are required. □

Assumption 5 (Unique insider optimizer (sufficient strong concavity)). *For each (t,p), the insider Hamiltonian is* uniformly strictly concave *in θ on Θ: there exists a constant $m_I > 0$ such that for all admissible (t,p) and all $\theta_1, \theta_2 \in \Theta$,*

$$\mathcal{H}^g(t,p,\theta_2) - \mathcal{H}^g(t,p,\theta_1) - \partial_\theta \mathcal{H}^g(t,p,\theta_1)\,(\theta_2 - \theta_1) \le -\frac{m_I}{2}|\theta_2 - \theta_1|^2.$$

A sufficient (checkable) condition is that the quadratic execution cost dominates the curvature of the jump-penalty term on the truncated domain, i.e. the negative curvature lower bound of $-\frac{1}{2}k(L)p\theta^2$ exceeds a uniform upper bound on the second derivative of $-\lambda(\theta)\frac{1}{\alpha}\big(e^{\alpha\Psi^g(\theta,p)} - 1\big)$ over Θ.

Lemma A.5 (Unique insider best response). *Under Assumption 5, the maximizer is unique and defines a measurable feedback $\theta^g : [0,T] \times \mathcal{P} \to \Theta$.*

Proof. Continuity in θ and compactness of Θ imply existence of maximizers. Strict concavity yields uniqueness. Measurability follows under uniqueness (or via measurable selection arguments), see standard treatments in [3, 5]. □

Define the insider best-response operator $\mathrm{BR}_I(g) := \theta^g$.

A.3 Bot best response given θ^g: QVI, uniqueness, and No-Zeno

Belief linearization as a continuous summary operator. To implement the paper's strategic linearization (Assumption 2) without proving global linearity of θ^g, we extract a local slope from the active-trading region.

Assumption 6 (Slope-extraction window). *Fix a compact window $\mathcal{N}_B \subset [0,T] \times \mathcal{P}$ contained in the active-trading region and a strictly positive continuous weight function w on $\mathcal{N}_B$. Assume*

$$\inf_{(t,p)\in\mathcal{N}_B} |V - p| \ge \underline{m} > 0,$$

so the window stays away from the safe-harbor boundary where the local slope degenerates.

Definition 2 (Belief summary operator). *For an admissible feedback rule $\theta : [0,T]\times\mathcal{P} \to \Theta$, define*

$$\mathcal{B}(\theta) := \frac{\int_{\mathcal{N}_B} w(t,p)\,\theta(t,p)\,(V-p)\,dt\,dp}{\int_{\mathcal{N}_B} w(t,p)\,(V-p)^2\,dt\,dp}.$$

This is the weighted local least-squares slope of θ with respect to mispricing $V - p$ on the active window.

We therefore write

$$\beta^g \;:=\; \mathcal{B}(\theta^g),$$

and obtain continuity directly.

Lemma A.6 (Continuity of the belief summary operator)**.** *Under Assumption 6, $\mathcal{B}$ is continuous with respect to the sup norm on admissible feedback rules.*

Proof. The denominator in Definition 2 is strictly positive by Assumption 6. For two admissible feedback rules θ_1, θ_2,

$$|\mathcal{B}(\theta_1) - \mathcal{B}(\theta_2)| \le \frac{\int_{\mathcal{N}_B} w(t,p)\,|V-p|\,dt\,dp}{\int_{\mathcal{N}_B} w(t,p)\,(V-p)^2\,dt\,dp}\|\theta_1 - \theta_2\|_\infty,$$

so $\mathcal{B}$ is Lipschitz, hence continuous, in the sup norm. □

Given β^g, the bot solves the stationary impulse-control problem whose value $\Phi^g(\xi)$ satisfies the QVI in the paper:

$$\min\{r\Phi - \mathcal{L}^{\beta^g}_{\text{bot}}\Phi,\;\; \Phi - \mathcal{M}^{\beta^g}\Phi\} = 0,$$

with intervention operator

$$\mathcal{M}^{\beta^g}\Phi(\xi) = \sup_{a\in\mathcal{U}}\Big\{\Phi(\xi^a) + R(a, v-p) - C_{\text{gas}}\Big\},$$

where ξ^a is the post-impulse state induced by the structural displacement Γ. The per-impulse fixed cost satisfies $C_{\text{gas}} > 0$.

Assumption 7 (Unique bot impulse)**.** *The maximizer over $a \in \mathcal{U}$ in the intervention operator is unique for each ξ (e.g., $R(\cdot,\cdot)$ strictly concave in a and $\mathcal{U}$ compact).*

Lemma A.7 (QVI well-posedness and unique impulse)**.** *Under (A4), $C_{\text{gas}} > 0$, compact $\mathcal{U}$, continuity/growth bounds on (R, Γ), and Assumption 7, the QVI admits a unique bounded continuous viscosity solution Φ^g, and there exists a unique optimal impulse rule $a^g(\xi)$ and an associated intervention region $\mathcal{A}^g = \{\xi \in \mathcal{X} : \Phi^g(\xi) = \mathcal{M}^{\beta^g}\Phi^g(\xi)\}$.*

Justification. These are standard results in impulse control of jump diffusions: the value solves the QVI in the viscosity sense; comparison yields uniqueness; verification yields existence of an optimal impulse policy. On truncated domains, we interpret the QVI in the *state-constraint viscosity* sense, which is standard for technical compactifications (see [4] and the control/viscosity discussion in [3]). Uniqueness of the impulse follows under Assumption 7. □

Lemma A.8 (No-Zeno)**.** *If $C_{\text{gas}} > 0$, any optimal impulse strategy has almost surely finitely many interventions on any finite horizon (no Zeno accumulation).*

Proof. If infinitely many impulses occur before time T, the total discounted cost is at least $\sum_i e^{-rT} C_{\text{gas}} = +\infty$, contradicting boundedness of the value on the compact domain. This is the standard no-accumulation argument for fixed-cost impulse control (see [4]). □

Define the bot best-response operator $\mathrm{BR}_B(\beta^g) := (a^g, \mathcal{A}^g)$.

A.4 Induced reduced-form objects and definition of $\mathscr{T}$

Given $(a^g, \mathcal{A}^g)$, define the induced reduced-form jump and loss by

$$\gamma'(p,\theta) \;:=\; \mathbb{E}\big[\Gamma\big(p,\theta,a^g(\Xi_t)\big) \mid \theta_t=\theta,\; P_t=p\big]\,, \qquad T'(\theta) \;:=\; \mathbb{E}\big[T_{\text{struct}}\big(\theta,a^g(\Xi_t)\big) \mid \theta_t=\theta\big]\,,$$

where T_{struct} denotes the structural loss function in the bot problem (as in the main text).

Structural hazard and projection. The mark-free QVI above is the baseline special case. To derive the structural hazard kernel explicitly, we temporarily enrich the attack technology with an exogenous mark u capturing block-level execution conditions; taking the mark distribution to be degenerate collapses the construction back to the baseline operator.

Assumption 8 (Block-level opportunity clock and marks)**.** *Conditional on the bot state Ξ_t, candidate attack opportunities arrive according to a marked point process $\mathcal{N}^g(dt,du)$ with compensator*

$$\nu^g(\Xi_t)F^g(du \mid \Xi_t)\,dt,$$

where ν^g is a bounded continuous opportunity intensity and $F^g(\cdot \mid \xi)$ is a mark distribution describing builder access, within-block position, and competing-flow conditions. At a

candidate opportunity with mark u, the bot compares continuation value $\Phi^g(\xi)$ with the mark-specific intervention value

$$\mathcal{M}_u^{\beta^g}\Phi^g(\xi) := \sup_{a\in\mathcal{U}}\Big\{\Phi^g(\xi^{a,u}) + R(a, v-p, u) - C_{\text{gas}}\Big\},$$

where $\xi^{a,u}$ is the post-attack state under mark u. The mark-free operator in the main text is recovered when u is degenerate and $R(a, v-p, u) = R(a, v-p)$. The bot attacks iff $\mathcal{M}_u^{\beta^g}\Phi^g(\xi) \geq \Phi^g(\xi)$, using the unique maximizing impulse $a^{g,\star}(\xi, u)$.

Lemma A.9 (Hazard kernel induced by the QVI)**.** *Under Assumption 8, the attack counting process can be written as*

$$N_t = \int_0^t \int \mathbf{1}\big\{\mathcal{M}_u^{\beta^g}\Phi^g(\Xi_{s-}) \geq \Phi^g(\Xi_{s-})\big\}\,\mathcal{N}^g(ds, du),$$

and therefore has bot-filtration compensator

$$\int_0^t h_g^B(\Xi_s)\,ds, \qquad h_g^B(\xi) = \nu^g(\xi)\int \mathbf{1}\big\{\mathcal{M}_u^{\beta^g}\Phi^g(\xi) \geq \Phi^g(\xi)\big\}\,F^g(du \mid \xi).$$

In the deterministic-mark special case, $h_g^B(\xi) = \nu^g(\xi)\mathbf{1}_{\{\xi\in\mathcal{A}^g\}}$.

Proof sketch. This is the predictable compensator of the thinned marked opportunity clock: a candidate event occurs with intensity $\nu^g(\xi)$, and conditional on the current state it is converted into an actual attack exactly on the mark set for which the QVI intervention payoff weakly dominates continuation. The resulting attack rate is therefore the baseline opportunity rate times the conditional attack probability implied by the intervention rule. □

Assumption 9 (Structural hazard kernel and stability)**.** *For each $g \in \mathcal{A}$, the bot best response induces a bounded continuous structural hazard kernel*

$$h_g^B : \mathcal{X} \to [0, \bar{\lambda}].$$

Moreover, if $g_n \to g$ uniformly in $\mathcal{A}$, then $h_{g_n}^B \to h_g^B$ uniformly on $\mathcal{X}$.

Assumption 10 (Range compatibility for the projection)**.** *The truncation domains are chosen so that for every β in the attainable range of β^g and every $\theta \in \Theta$, there exists*

$(v,p) \in \mathcal{V} \times \mathcal{P}$ *such that* $|\beta\,(v-p)| = |\theta|$. *Equivalently,* $\Gamma^g_{|\theta|} := \{\xi \in \mathcal{X} : |\beta^g(v-p)| = |\theta|\}$ *is nonempty for all* $\theta \in \Theta$.

To return to a θ-only intensity, first define the capped upper-envelope

$$\tilde{\lambda}(\theta) \;:=\; \min\Big\{\bar{\lambda},\; \sup\big\{h^B_g(\xi) : \xi = (v,\Sigma,p,L) \in \mathcal{X},\; |\beta^g(v-p)| = |\theta|\big\}\Big\}.$$

To ensure that the intensity lies in the admissible shape-restricted class (A1), introduce a shape-preserving projection.

Assumption 11 (Shape-preserving projection for intensities)**.** *There exists a continuous operator* $\Pi_\lambda : C(\Theta) \to C(\Theta)$ *such that for any* $f \in C(\Theta)$, $\Pi_\lambda(f)$ *satisfies (A1) (continuous, even, increasing and convex in* $|\theta|$*, and* $\mathrm{Lip} \le L_\lambda$*), and* $\Pi_\lambda(f) = f$ *whenever* f *already satisfies (A1).*

Define the reduced-form intensity as

$$\lambda'(\theta) \;:=\; \Pi_\lambda\big(\tilde{\lambda}\big)(\theta).$$

Lemma A.10 (Continuity of the projection)**.** *Under Assumptions 10 and 9, and with* $\beta^g \neq 0$*, the map* $\tilde{\lambda}(\cdot)$ *is continuous on* Θ*, and hence* $\lambda'(\cdot) = \Pi_\lambda(\tilde{\lambda})(\cdot)$ *is continuous on* Θ.

Proof. Define the correspondence $\Gamma^g_{|\theta|} := \{\xi \in \mathcal{X} : |\beta^g(v-p)| = |\theta|\}$. By Assumption 10, $\Gamma^g_{|\theta|}$ is nonempty; it is compact-valued with closed graph on compact domains. By Assumption 9, h^B_g is continuous on $\mathcal{X}$. Therefore Berge's maximum theorem implies continuity of the value function $\theta \mapsto \sup_{\xi \in \Gamma^g_{|\theta|}} h^B_g(\xi)$ (see [6]). The outer truncation by $\min\{\bar{\lambda}, \cdot\}$ preserves continuity, so $\tilde{\lambda}$ is continuous. Finally, continuity of Π_λ in Assumption 11 yields continuity of $\lambda' = \Pi_\lambda(\tilde{\lambda})$. □

Definition (Best-response mapping). For any $g \in \mathcal{A}$, define $\mathscr{T}(g) = (\lambda', \gamma', T')$ by the pipeline:

1. solve insider HJB under g to obtain $\theta^g = \mathrm{BR}_I(g)$;
2. set $\beta^g = \mathcal{B}(\theta^g)$;
3. solve bot QVI under β^g to obtain $(a^g, \mathcal{A}^g) = \mathrm{BR}_B(\beta^g)$;
4. induce (λ', γ', T') as above.

A.5 Invariance and continuity of $\mathscr{T}$; existence by Schauder

Lemma A.11 (Self-mapping). *For all $g \in \mathcal{A}$, $\mathscr{T}(g) \in \mathcal{A}$.*

Proof sketch. Boundedness and Lipschitz properties of γ' and T' follow from boundedness/Lipschitz of structural primitives on compact domains and compactness of $\mathcal{U}$ (see [4]). Continuity and boundedness of λ' follow from Lemma A.10 and the cap by $\bar{\lambda}$. The intensity λ' satisfies (A1) by construction since $\lambda' = \Pi_\lambda(\tilde{\lambda})$ and Assumption 11 guarantees that Π_λ maps any continuous input back into the admissible shape-restricted class. Hence $\mathscr{T}(g)$ satisfies (A1)–(A4), i.e., $\mathscr{T}(\mathcal{A}) \subseteq \mathcal{A}$. □

Lemma A.12 (Insider stability). *Let $g_n \to g$ uniformly in $\mathcal{A}$. Then $J^{g_n} \to J^g$ uniformly on $[0,T] \times \mathcal{P}$, and under Assumption 5, $\theta^{g_n} \to \theta^g$ uniformly.*

Justification. Viscosity solutions are stable under uniform perturbations of coefficients and nonlocal terms when comparison holds, using half-relaxed limits (see [2]). Uniqueness of the maximizer yields continuous dependence of the argmax on the parameters (maximum theorem logic; [6]). □

Lemma A.13 (Continuity of belief summary). *If $g_n \to g$, then $\beta^{g_n} \to \beta^g$.*

Proof. Immediate from Lemma A.6 and Lemma A.12. □

Lemma A.14 (Bot stability). *Let $g_n \to g$. Then $\Phi^{g_n} \to \Phi^g$ uniformly on $\mathcal{X}$, and under Assumption 7, $a^{g_n} \to a^g$ in the policy topology; hence $\gamma'_n \to \gamma'$ and $T'_n \to T'$ uniformly.*

Justification. Impulse-control QVIs admit viscosity comparison and stability properties analogous to HJBs under coefficient perturbations; uniqueness ensures a continuous optimizer selection (see [4] and the viscosity framework in [2]). □

Lemma A.15 (Continuity of λ'). *Let $g_n \to g$ uniformly in $\mathcal{A}$. Under Assumptions 9, 10, and 11, we have $\lambda'_n \to \lambda'$ uniformly on Θ.*

Justification. By Assumption 9, $h^B_{g_n} \to h^B_g$ uniformly on $\mathcal{X}$. Together with $\beta^{g_n} \to \beta^g$ from Lemma A.13, Berge's maximum theorem yields $\tilde{\lambda}_n \to \tilde{\lambda}$ uniformly on Θ (with nonemptiness ensured by Assumption 10), and continuity of Π_λ in Assumption 11 implies $\lambda'_n = \Pi_\lambda(\tilde{\lambda}_n) \to \Pi_\lambda(\tilde{\lambda}) = \lambda'$ uniformly. □

Proposition 7 (Continuity of $\mathscr{T}$). *$\mathscr{T} : \mathcal{A} \to \mathcal{A}$ is continuous.*

Proof. Combine Lemmas A.12–A.15 componentwise. □

Theorem A.16 (Schauder fixed point)**.** *There exists $g^\star \in \mathcal{A}$ such that $\mathscr{T}(g^\star) = g^\star$.*

Proof. By Lemma A.1, $\mathcal{A}$ is nonempty, compact, and convex. By Proposition 7 and Lemma A.11, $\mathscr{T} : \mathcal{A} \to \mathcal{A}$ is continuous. Hence Schauder's fixed point theorem implies existence of a fixed point (see [7]). □

Completion of the proof of Proposition 2. Let $g^{eq} = (\lambda^{eq}, \gamma^{eq}, T^{eq})$ be a fixed point of $\mathscr{T}$. Define $\theta^\star = \mathrm{BR}_I(g^{eq})$, $\beta^\star = \mathcal{B}(\theta^\star)$, and $(a^\star, \mathcal{A}^\star) = \mathrm{BR}_B(\beta^\star)$. By construction of $\mathscr{T}$, the induced reduced-form triple from $(a^\star, \mathcal{A}^\star)$ coincides with g^{eq}. Thus insider optimality, bot optimality, and consistency (Definition 1 in the main text) all hold, establishing existence of a Markovian equilibrium. ■

Appendix B. Verification for the HJB and QVI (Sketch)

Boundary convention on truncated domains. All PDE/QVI statements below are understood in the *state-constraint viscosity* sense on the truncated compact domains. This avoids the need to impose artificial boundary conditions while preserving comparison/uniqueness on compact sets.

B.1 Insider problem: HJB verification

Definition 3 (Admissible controls (insider))**.** *A control $\theta = \{\theta_t\}_{t\in[0,T]}$ is admissible if it is $\{\mathcal{F}_t\}$-predictable, takes values in the compact set Θ, and the controlled state dynamics admit a (weak/strong) solution on $[0,T]$ with finite objective value. On truncated compact domains with bounded coefficients, compact controls, and standard Lipschitz/linear-growth conditions, any bounded feedback $\theta(t,p) \in \Theta$ is admissible.*

Lemma B.17 (HJB verification (finite horizon, CARA case))**.** *Fix a reduced-form triple g in the admissible reduced-form class (as in Appendix 8.2). Let $V^g(t,p)$ denote the insider value (certainty equivalent) and let $J^g \in C([0,T]\times\mathcal{P})$ be the unique bounded viscosity solution of the insider HJB with terminal condition $J^g(T,p) = 0$. Assume the comparison principle holds in the state-constraint class and that the Hamiltonian admits a measurable maximizer $\theta^\star(t,p) \in \arg\max_{\theta\in\Theta}\{\cdots\}$ (e.g. by compactness and a uniqueness/strict concavity condition). Then:*

1. $J^g = V^g$ *on* $[0,T] \times \mathcal{P}$*;*

2. *the feedback control* $\theta^\star$ *is optimal.*

Proof sketch. (i) *DPP* $\Rightarrow$ *viscosity property.* By the dynamic programming principle for controlled jump-diffusions on compact domains, the value function V^g is a viscosity solution of the HJB in the state-constraint sense.

(ii) *Comparison* $\Rightarrow$ *uniqueness.* The assumed comparison principle yields uniqueness of bounded viscosity solutions, hence $V^g \equiv J^g$.

(iii) *Verification / optimality of* $\theta^\star$*.* Since $\theta^\star$ attains the HJB Hamiltonian supremum pointwise, standard verification arguments imply that the corresponding controlled performance achieves the value. See, e.g., [3, 5]. □

B.2 Bot problem: QVI verification (impulse control)

Definition 4 (Admissible impulse controls (bot))**.** *An impulse control is a sequence* $(\tau_i, a_i)_{i\geq 1}$ *of* $\{\mathcal{F}_t\}$*-stopping times and impulses* $a_i \in \mathcal{U}$ *such that the post-impulse state remains in the truncated domain and the discounted total payoff is well-defined. Under a fixed per-impulse cost* $C_{\text{gas}} > 0$*, optimal strategies are non-explosive (No-Zeno) on any finite time interval.*

Lemma B.18 (QVI verification and region decomposition)**.** *Fix* β^g *and consider the bot impulse-control problem on the compact state space* $\mathcal{X}$*. Let* $\Phi^g \in C(\mathcal{X})$ *be the unique bounded viscosity solution of the stationary QVI*

$$\min\{r\Phi - \mathcal{L}^{\beta^g}_{\text{bot}}\Phi,\ \Phi - \mathcal{M}^{\beta^g}\Phi\} = 0,$$

and assume $\mathcal{M}^{\beta^g}\Phi$ *is continuous (e.g. by compact* $\mathcal{U}$ *and continuity of primitives). Then:*

1. Φ^g *coincides with the bot value function;*

2. *the* attack region $\mathcal{A}^g := \{\xi \in \mathcal{X} : \Phi^g(\xi) = \mathcal{M}^{\beta^g}\Phi^g(\xi)\}$ *is closed and the* wait region $\mathcal{C}^g := \{\xi : \Phi^g(\xi) > \mathcal{M}^{\beta^g}\Phi^g(\xi)\}$ *is open;*

3. *if there exists a measurable selector* $a^\star(\xi) \in \arg\max_{a\in\mathcal{U}}\{\Phi^g(\xi^a) + R(a,\cdot) - C_{\text{gas}}\}$*, then the impulse rule "intervene when* $\xi_t \in \mathcal{A}^g$ *and choose* $a^\star(\xi_t)$*" is optimal.*

Proof sketch. (i) *DPP ⇒ QVI viscosity property.* The value of an impulse-control problem satisfies the associated QVI in viscosity sense.

(ii) *Comparison ⇒ uniqueness.* Uniqueness follows from comparison on compact domains.

(iii) *Verification / optimal impulse.* When the intervention operator is attained by a measurable maximizer and $C_{\text{gas}} > 0$ rules out Zeno, standard verification yields optimality. Closedness/openness follow because $\Phi^g - \mathcal{M}^{\beta g}\Phi^g$ is continuous. See, e.g., [4]. □

Remark 5 (Free boundary regularity)**.** *The existence/verification arguments above do not require C^1 smooth-fit regularity of the free boundary $\partial\mathcal{A}^g$. Continuity and comparison/uniqueness of viscosity solutions on truncated domains are sufficient for the fixed-point existence argument in Appendix 8.2.*

Appendix C. Measure Change and Intensity Transform (P to Q)

Setup. On $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$, let W be a $\mathbb{P}$-Brownian motion and N be a counting process with physical predictable compensator intensity λ_t^I, i.e. $\widetilde{N}_t := N_t - \int_0^t \lambda_s^I\, ds$ is a $\mathbb{P}$-local martingale. All processes are defined with respect to $\{\mathcal{F}_t\}$, and λ_t^I, ξ_t, κ_t are understood as $\mathcal{F}_{t-}$-predictable processes. In the main text, λ_t^I is the projected physical compensator under the insider/public filtration.

Assumption 12 (Sufficient conditions for equivalence)**.** *Fix a finite horizon $T < \infty$. Let ξ_t (diffusive market price of risk) and κ_t (jump risk premium ratio) be predictable processes. Assume:*

(C1) ***(Diffusive part)*** *$\int_0^T \xi_t^2\, dt < \infty$ a.s., and, for instance, ξ is bounded on $[0,T]$ (hence Novikov's condition holds on $[0,T]$).*

(C2) ***(Jump part positivity/boundedness)*** *$0 < \underline{\kappa} \leq \kappa_t \leq \bar{\kappa} < \infty$ a.s. on $[0,T]$ (so $\log \kappa_t$ is bounded).*

(C3) ***(Non-explosion / integrability)*** *$\int_0^T \lambda_t^I\, dt < \infty$ a.s. and $\int_0^T |\kappa_t - 1|\, \lambda_t^I\, dt < \infty$ a.s.*

Density process. Define the Doléans–Dade exponential (likelihood ratio) by

$$\Lambda_t := \exp\Big(-\int_0^t \xi_s\, dW_s - \frac{1}{2}\int_0^t \xi_s^2\, ds\Big)\ \exp\Big(\int_0^t \log\kappa_s\, dN_s - \int_0^t (\kappa_s - 1)\lambda_s^I\, ds\Big), \qquad 0 \le t \le T. \tag{18}$$

Equivalently, in differential form,

$$\frac{d\Lambda_t}{\Lambda_{t-}} = -\xi_t\, dW_t + (\kappa_t - 1)\, d\widetilde{N}_t, \qquad \Lambda_0 = 1, \tag{19}$$

where $\widetilde{N}_t = N_t - \int_0^t \lambda_s^I ds$.

Lemma C.19 (Well-defined equivalent measure and intensity transform)**.** *Under Assumption 12, Λ is a strictly positive $\mathbb{P}$-martingale on $[0,T]$ with $\mathbb{E}^P[\Lambda_T] = 1$. Therefore the probability measure $\mathbb{Q}$ defined by*

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_T} = \Lambda_T$$

is well-defined and satisfies $\mathbb{Q} \sim \mathbb{P}$ on $\mathcal{F}_T$. Moreover:

(i) The process $W_t^{\mathbb{Q}} := W_t + \int_0^t \xi_s\, ds$ is a $\mathbb{Q}$-Brownian motion on $[0,T]$.

(ii) The $\mathbb{Q}$-intensity of N is

$$\lambda_t^{\mathbb{Q}} = \kappa_t\, \lambda_t^I, \qquad 0 \le t \le T,$$

i.e. $\widetilde{N}_t^{\mathbb{Q}} := N_t - \int_0^t \lambda_s^{\mathbb{Q}}\, ds$ is a $\mathbb{Q}$-local martingale.

Proof sketch (standard). Step 1 (Brownian part). By (C1), e.g. bounded ξ, Novikov's condition holds on $[0,T]$, hence the stochastic exponential $\exp(-\int_0^t \xi_s\, dW_s - \frac{1}{2}\int_0^t \xi_s^2 ds)$ is a true $\mathbb{P}$-martingale on $[0,T]$.

Step 2 (counting-process part). By (C2)–(C3), κ_t is $\mathcal{F}_{t-}$-predictable and strictly positive with bounded $\log\kappa_t$, and the compensator adjustment $\int_0^T (\kappa_t - 1)\lambda_t^I dt$ is finite a.s.; therefore the second exponential in (18) is a strictly positive $\mathbb{P}$-martingale. This follows from the standard Girsanov theorem / exponential-martingale change of measure for counting processes.

Step 3 (equivalence and intensity transform). Under the standard exponential-martingale conditions for the Brownian and point-process components, the Doléans–Dade exponen-

tial in (18) is a true martingale, so $\mathbb{E}^P[\Lambda_T] = 1$ and $\mathbb{Q}$ is well-defined with $\mathbb{Q} \sim \mathbb{P}$ on $\mathcal{F}_T$. Girsanov's theorem yields the drift shift for W and the compensator/intensity transform for N, giving $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$. □

Remark 6 (Interpretation in the main text)**.** *In the paper, λ_t^I is the projected physical compensator and $\kappa_t > 1$ is interpreted as a jump-risk premium amplification channel. Under Lemma C.19, jump arrivals are priced as if more frequent under $\mathbb{Q}$: $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^I$ on $[0, T]$.*

Appendix D. Numerical Implementation and Robustness

Purpose. This appendix documents the numerical procedures underlying Section 7. The main text provides mechanism images (policy functions, safe-harbor regions, and boundary diagrams); this appendix provides numerical credibility (domains, discretization choices, stability safeguards, convergence diagnostics, and the boundary proxy used for Figure 6). The numerical analysis is illustrative rather than calibrational.

D.1 Numerical primitives, normalizations, and domains

Normalizations. We adopt the normalizations in the main text: the horizon is normalized to $T = 1$ and the fundamental to $V = 1$. Thus intensities can be interpreted as expected event counts over the normalized horizon, and price objects are in percentage terms.

Domains. We solve the HJB on a truncated compact price domain centered around V,

$$p \in [p_{\min}, p_{\max}] = [0.7, 1.3],$$

and restrict the insider control to a compact set,

$$\theta \in [-\bar{\theta}, \bar{\theta}] = [-3, 3].$$

Truncation is a numerical device rather than an economic assumption.

Table D.1: Numerical settings used in Section 7.

Parameter	Symbol	Value	Note (Source)
Price Grid	N_p	81	$N_p = 81$
Control Grid	N_θ	81	$N_\theta = 81$
Time Steps	N_t	140	$N_t = 140$
Control Bound	$\bar{\theta}$	3.0	$\theta_{\max} = 3.0$
Clip Constant	C_{clip}	25.0	clip_val = 25.0
Boundary	–	Neumann	$J_0 = J_1$ (and symmetrically at upper end)

Note. Table D.1 reports the baseline computational grid, truncated domains, and numerical-stability settings used in Section 7; these choices are implementation devices for transparent mechanism illustration rather than calibration or estimation.

Reduced-form primitives. For numerical illustration we use transparent reduced-form predation primitives that satisfy maintained shape restrictions (monotonicity/convexity):

$$\lambda(\theta) = \lambda_0 + \lambda_1 |\theta|^\beta, \qquad \beta \geq 1,$$

$$\gamma(p, \theta) = \eta \, p \, \theta,$$

$$T(\theta) = \frac{\tau}{2}\theta^2.$$

This $T(\theta)$ is the same reduced-form loss object used in the main text, and $\mathscr{T}$ is reserved for the fixed-point operator in Appendix 8.2.

D.2 HJB solver details (discretization, boundaries, nonlocal term, stability)

Grids. The baseline discretization uses

$$(N_p, N_\theta, N_t) = (81, 81, 140),$$

i.e., 81 price nodes, an 81-point control grid (including $\theta = 0$), and 140 time steps on $t \in [0, T]$.

Backward iteration and terminal condition. We solve backward in time from the CARA-separation terminal condition:

$$J(T, p) = 0.$$

Control optimization. At each grid point (t_n, p_i), the optimal control is computed by grid search:

$$\theta^*(t_n, p_i) \in \arg\max_{\theta_k \in [-3,3]} \mathcal{H}\left(t_n, p_i, \theta_k; J\right),$$

where $\mathcal{H}$ denotes the discretized Hamiltonian consistent with the corrected HJB in the main text.

Nonlocal term and interpolation. The HJB contains a nonlocal continuation term through $J(t, p + \gamma(p, \theta))$. Since $p + \gamma(p, \theta)$ is generally off-grid, we evaluate $J(t, p + \gamma)$ using linear interpolation on the price grid. If the argument falls outside $[p_{\min}, p_{\max}]$, it is handled consistently with the boundary condition below.

Boundary conditions (Neumann / reflecting). At the truncated endpoints we impose a Neumann-type (zero-slope / reflecting) boundary condition implemented as

$$J_0 = J_1, \qquad J_{N_p-1} = J_{N_p-2}.$$

This avoids artificial anchoring from Dirichlet boundaries and mitigates boundary distortion of interior policy rules.

Exponential overflow protection (CARA clipping). The CARA jump term involves $\exp(\alpha\Psi)$ where

$$\Psi(\theta, p) = T(\theta) + J(t, p) - J\big(t, p + \gamma(p, \theta)\big).$$

To prevent overflow/underflow during intermediate iterations, we clip:

$$\alpha\Psi \leftarrow \max\big(\min(\alpha\Psi,\ C_{\text{clip}}),\ -C_{\text{clip}}\big), \qquad C_{\text{clip}} = 25.$$

This is a numerical stability safeguard, not an economic restriction.

D.3 Computing the safe-harbor width $W(\alpha, \lambda_1)$

Definition (restated). Let mispricing be $m \equiv V - p$. For each (α, λ_1) we compute the time-0 policy $\theta^*(0, p)$ and define the safe-harbor (near-zero trading) set:

$$S(\alpha, \lambda_1) = \{m \mid |\theta^*(0, V - m)| \leq \varepsilon\},$$

and its width:

$$W(\alpha, \lambda_1) = \mathrm{Leb}\big(S(\alpha, \lambda_1)\big).$$

Continuous interval from a discrete grid. To reduce "staircase" artifacts from discretization, we locate the boundary endpoints by linear interpolation at the crossings of $|\theta^*(0, p)| = \varepsilon$ on the p-grid, and then map endpoints back to $m = V - p$ to obtain $W(\alpha, \lambda_1)$.

D.4 Robustness and numerical credibility checks

Grid and truncation sensitivity (recommended reporting). Because the numerical exercise is illustrative rather than calibrational, we emphasize qualitative robustness of the main comparative statics to discretization and truncation choices. A concise robustness report typically includes: (i) grid refinement checks for (N_p, N_θ, N_t), (ii) stability of $W(\alpha, \lambda_1)$ under alternative ε values, and (iii) verification that the main patterns are not driven by boundary effects on $[p_{\min}, p_{\max}]$.

D.5 Convergence diagnostics (stopping rules)

HJB tolerance. We monitor sup-norm convergence and stop policy iteration when

$$\|J^{(n+1)} - J^{(n)}\|_\infty \leq tol_{\mathrm{HJB}}, \qquad tol_{\mathrm{HJB}} = 10^{-6}.$$

Coupled fixed-point tolerance (when applicable). When iterating coupled objects in a HJB–QVI fixed point (as described in the main text), we declare convergence when

$$\|g^{(n+1)} - g^{(n)}\|_\infty \leq tol_{\mathrm{FP}}, \qquad tol_{\mathrm{FP}} = 10^{-6}.$$

Outer-loop diagnostics. Because Schauder's theorem guarantees existence but not convergence of the damped fixed-point iteration, the coupled loop is treated as a numerical root-finding device. The relevant diagnostics are the damping parameter ω, the residual sequence

$$r_n := \|g^{(n+1)} - g^{(n)}\|_\infty,$$

and robustness across different initializations of g_0.

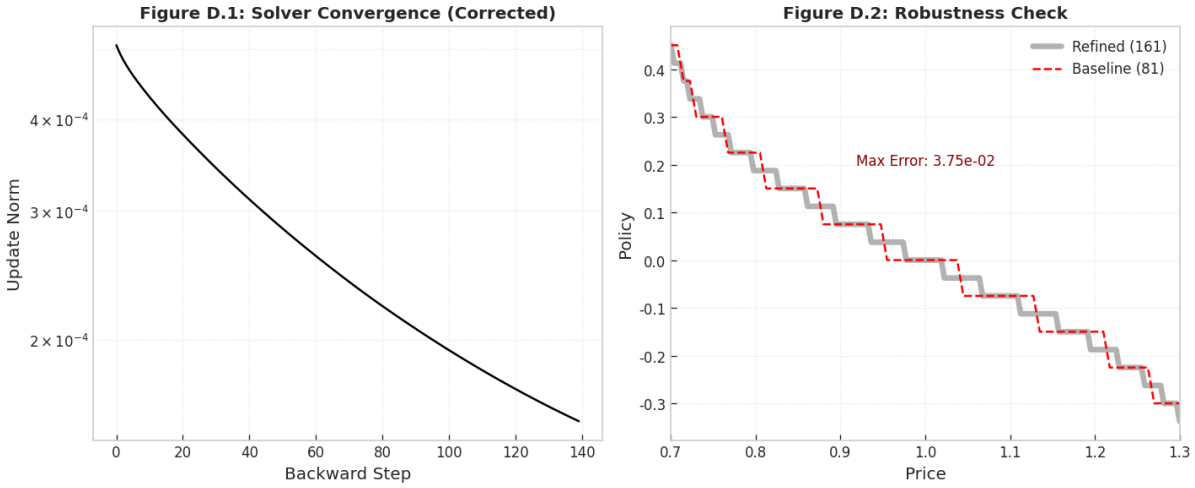


Figure 9: Solver Convergence and Robustness Check. Figure D.1 reports a stability diagnostic of the backward HJB solver: the update norm remains well-behaved and declines smoothly across backward steps, indicating numerically stable backward propagation on the truncated domain. Figure D.2 overlays the time-0 policy $\theta^*(0, p)$ computed on the baseline and refined price grids. The two policies are closely aligned; discretization mainly affects local step features, while the qualitative structure and comparative-statics implications are unchanged.

D.6 Figure 6 boundary proxy: definition and interpretation

Local-curvature expansion. Let $d := (v-p)/\sigma_Y$ denote the normalized informational distance to the insider's signal. Around the continuation/intervention boundary, suppose the mark-averaged intervention gain admits the second-order expansion

$$\Delta^\Phi(d; C_{\text{gas}}, \sigma_Y) = \frac{d^2}{2\chi\,\sigma_Y^2} - C_{\text{gas}} + o(d^2), \qquad \chi > 0,$$

where χ is the reciprocal local curvature scale of the value gain from intervening rather than waiting. Solving $\Delta^{\Phi} = 0$ at leading order yields the boundary proxy

$$b(C_{\text{gas}}, \sigma_Y) = \sqrt{2\chi\, C_{\text{gas}}} \cdot \sigma_Y.$$

Interpretation. The proxy preserves the intended monotone comparative statics: higher fixed per-attack costs C_{gas} and higher observation noise σ_Y raise the threshold b required for profitable intervention, implying an outward shift of the predation boundary and an expansion of the inaction (wait) region. The parameter χ is therefore not a free plotting knob; it is the local curvature coefficient of the intervention-gain expansion. The proxy is intended to preserve comparative statics rather than to replace a full QVI boundary surface.

Appendix E. Proof of the Conditional Welfare Proposition

Proof of Proposition 3. Fix a state $s = (p, L)$ and suppress the state subscript. Let θ^{FB} maximize the strictly concave function w, so $w'(\theta^{FB}) = 0$ and w' is strictly decreasing on the admissible interval.

Because G is convex and differentiable with $G(0) = 0$ and $\theta G'(\theta) \geq 0$ for all admissible θ, the point $\theta = 0$ is a global minimizer of G. Hence

$$G'(0) = 0, \qquad G'(\theta) \geq 0 \text{ for } \theta \geq 0, \qquad G'(\theta) \leq 0 \text{ for } \theta \leq 0.$$

Since $\pi = w - G$ is the sum of a strictly concave function and a concave function, it is strictly concave and therefore has a unique maximizer $\theta^{\star}$.

If $\theta^{FB} \geq 0$, strict concavity of w and the first-order condition $w'(\theta^{FB}) = 0$ imply $w'(0) \geq 0$. Therefore

$$\pi'(0) = w'(0) - G'(0) = w'(0) \geq 0.$$

At the transfer-neutral maximizer,

$$\pi'(\theta^{FB}) = w'(\theta^{FB}) - G'(\theta^{FB}) = -G'(\theta^{FB}) \leq 0.$$

Because π' is strictly decreasing, its unique zero lies in $[0, \theta^{FB}]$. Thus $0 \leq \theta^{\star} \leq \theta^{FB}$.

If $\theta^{FB} \leq 0$, strict concavity of w implies $w'(0) \leq 0$, so

$$\pi'(0) = w'(0) - G'(0) = w'(0) \leq 0.$$

At θ^{FB} one has

$$\pi'(\theta^{FB}) = -G'(\theta^{FB}) \geq 0,$$

because $G'(\theta^{FB}) \leq 0$ on the negative half-line. Again using that π' is strictly decreasing, its unique zero lies in $[\theta^{FB}, 0]$. Thus $\theta^{FB} \leq \theta^{\star} \leq 0$.

The two cases imply

$$|\theta^{\star}| \leq |\theta^{FB}|.$$

If $\theta^{FB} G'(\theta^{FB}) > 0$, then $\pi'(\theta^{FB})$ has the opposite sign from θ^{FB}, so the zero of π' lies strictly between 0 and θ^{FB}. Hence $|\theta^{\star}| < |\theta^{FB}|$.

Finally, if $\theta^{\star} \neq \theta^{FB}$, strict concavity of w and optimality of θ^{FB} imply

$$w(\theta^{\star}) < w(\theta^{FB}),$$

which is the claimed transfer-neutral welfare loss. □

References

[1] G. Barles and C. Imbert, "Second-order elliptic integro-differential equations: viscosity solutions' theory revisited," *Annales de l'Institut Henri Poincaré C, Analyse Non Linéaire*, vol. 25, no. 3, pp. 567–585, 2008.

[2] M. G. Crandall, H. Ishii, and P.-L. Lions, "User's guide to viscosity solutions of second order partial differential equations," *Bulletin of the American Mathematical Society*, vol. 27, no. 1, pp. 1–67, 1992.

[3] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, 2nd ed. New York: Springer, 2006.

[4] B. Øksendal and A. Sulem, *Applied Stochastic Control of Jump Diffusions*, 3rd ed. Cham: Springer, 2019.

[5] H. Pham, *Continuous-time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer, 2009.

[6] C. Berge, *Topological Spaces*. New York: Macmillan, 1963.

[7] H. Cornean, "Notes on Schauder's Fixed Point Theorem," Lecture Notes, Aalborg University.