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

Ruoran Lai Tingyi Lin

Abstract

Can code deter insider trading more effectively than law? In decentralized markets lacking public enforcement, we identify a novel mechanism of "privatized regulation": predatory algorithms (MEV bots) that impose an immediate, state-dependent penalty on toxic order flow. By solving a coupled HJB-QVI equilibrium between a risk-averse insider and a strategic predator, we demonstrate that the threat of "sandwich attacks" functions as an endogenous "shadow tax" on information. This tax generates a structural "Liquidity Trap"—a safe-harbor region where algorithmic deterrence compels insiders into inaction, effectively replicating a regulatory blackout period without legal intervention. However, we show that this "governance by code" comes at a systemic cost: the strategic interaction amplifies jump risk premiums, generating endogenous fat tails and excess crash risk in asset prices. Our findings suggest that MEV is not merely a transaction cost but a market-based substitute for legal deterrence, replacing the delayed penalties of the SEC with the immediate microstructure costs of volatility.

JEL Classification: G14, D47, C73, G12

Keywords: MEV; sandwich attacks; algorithmic predation; privatized enforcement; insider 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.The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Any remaining errors are our own.

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 (Daian et al. 2019; 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 an endogenous, market-based enforcement technology that disciplines informed trading—albeit at the price of new tail risks and welfare distortions (Budish, Cramton, and Shim 2015; Bahrani, Garimidi, and Roughgarden 2024).

1.1 A private-enforcement 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 replaces much of this public enforcement with code and competition: the public visibility of order flow and the ability to re-order transactions generate immediate, state-contingent losses for traders who reveal valuable information or urgency. In this sense, MEV searchers act as private enforcers—they do not merely trade against informed orders; they can punish them mechanically through execution externalities (Daian et al. 2019; Cong and He 2019).

This perspective reframes a long-standing concern in market microstructure: informa-

tion 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. 2019).

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 (Kyle 1985). Back (1992) provides a continuous-time formalization, delivering tractability and characterizing equilibrium pricing rules under broad distributions (Back 1992). 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 (Daian et al. 2019; Weintraub et al. 2022). 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. 2019). 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 (Budish, Cramton, and Shim 2015). DeFi order flow exhibits a parallel tension: continuous-time mempool visibility and sequential execution create exploitable priority opportunities (Daian et al. 2019). 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 (Daian et al. 2019). 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 enforcement 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 (Daian et al. 2019; 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. This "liquidity trap" is not imposed by regulation; it arises endogenously from the threat of private predation (Daian et al. 2019).

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

The "private enforcement" 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 predatory behavior "cheap" and immediate, thereby curbing informed exploitation of uninformed liquidity (Daian et al. 2019). 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 insider predation by taxing informational rents. 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; Cong and He 2019).

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

Private enforcement 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 fat-tailed returns and elevated crash 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 enforcement can reduce informed trading but increases 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 (Daian et al. 2019; Weintraub et al. 2022). Third, we complement market-design and blockchain-economics perspectives: continuous-time priority creates rents and arms races (Budish, Cramton, and Shim 2015), 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 deliver three contributions. (i) *Mechanism:* MEV acts as a private enforcement technology that imposes an endogenous shadow tax on informed execution and generates a safe-harbor/liquidity-trap region (Becker 1968; Daian et al. 2019). (ii) *Method:* We provide a tractable coupled equilibrium linking a risk-averse insider's execution problem to a predator's partial-observation impulse-control problem (Kyle 1985; Bahrani, Garimidi, and Roughgarden 2024). (iii) *Pricing:* Strategic predation endogenizes jump risk and, through a jump-risk premium channel, produces fat tails and elevated crash risk (Merton 1976; Liu and Tsyvinski 2021; Cong, Li, and Wang 2021).

1.9 Roadmap

The remainder of the paper is organized as follows. Section 2 describes the trading environment and the insider's execution problem. Section 3 derives the insider's HJB and characterizes the shadow-tax wedge and safe-harbor region. Section 4 models the MEV bot's information and impulse-control problem and derives the QVI and optimal intervention boundary. Section 5 characterizes equilibrium as a coupled HJB–QVI fixed point and presents comparative statics. Section 6 studies the induced return dynamics and pricing implications under jump risk. Section 7 provides numerical illustrations, and the Appendix contains proofs.

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.

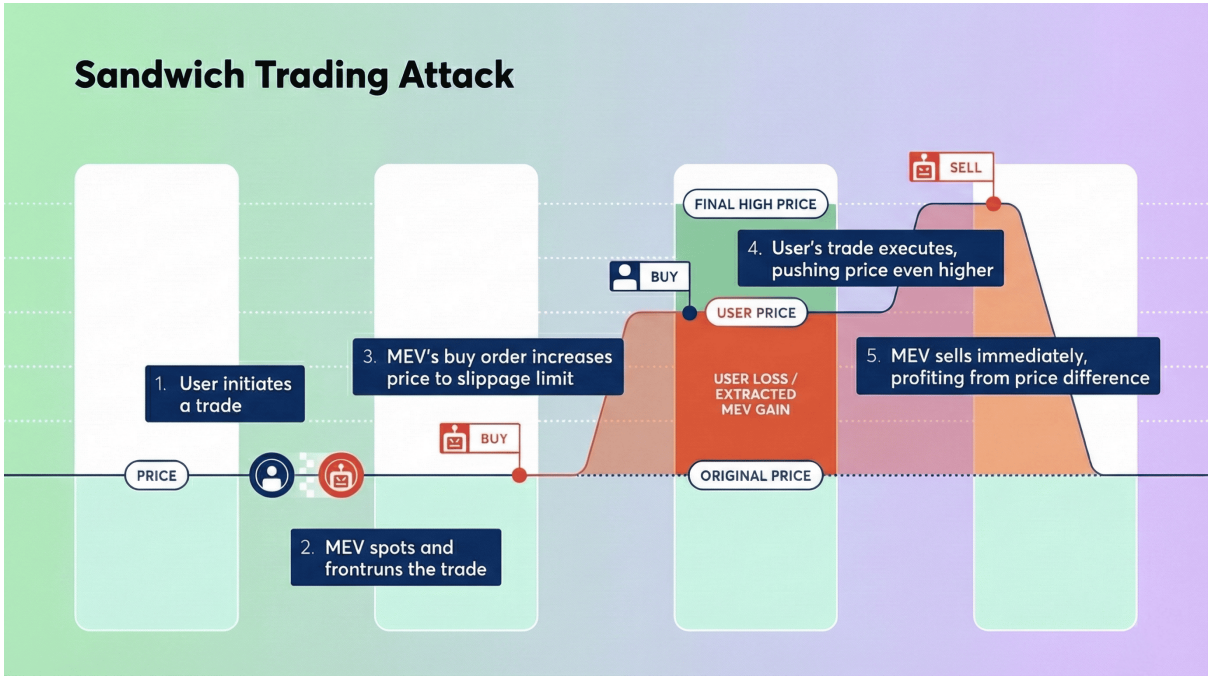


Figure 1: **Timeline of a Sandwich Attack Mechanism.**
Note: This figure illustrates the discrete lifecycle of a predatory MEV extraction, tracing how a searcher exploits mempool visibility to strategically order front-running and back-running transactions around a victim, thereby endogenously determining the execution price.

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 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 "privatized enforcement": 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 punishment that is immediate and mechanically implemented through execution outcomes rather than through delayed legal sanctions.

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 "privatized enforcement" 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. 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 2.2 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 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.

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.

0.5

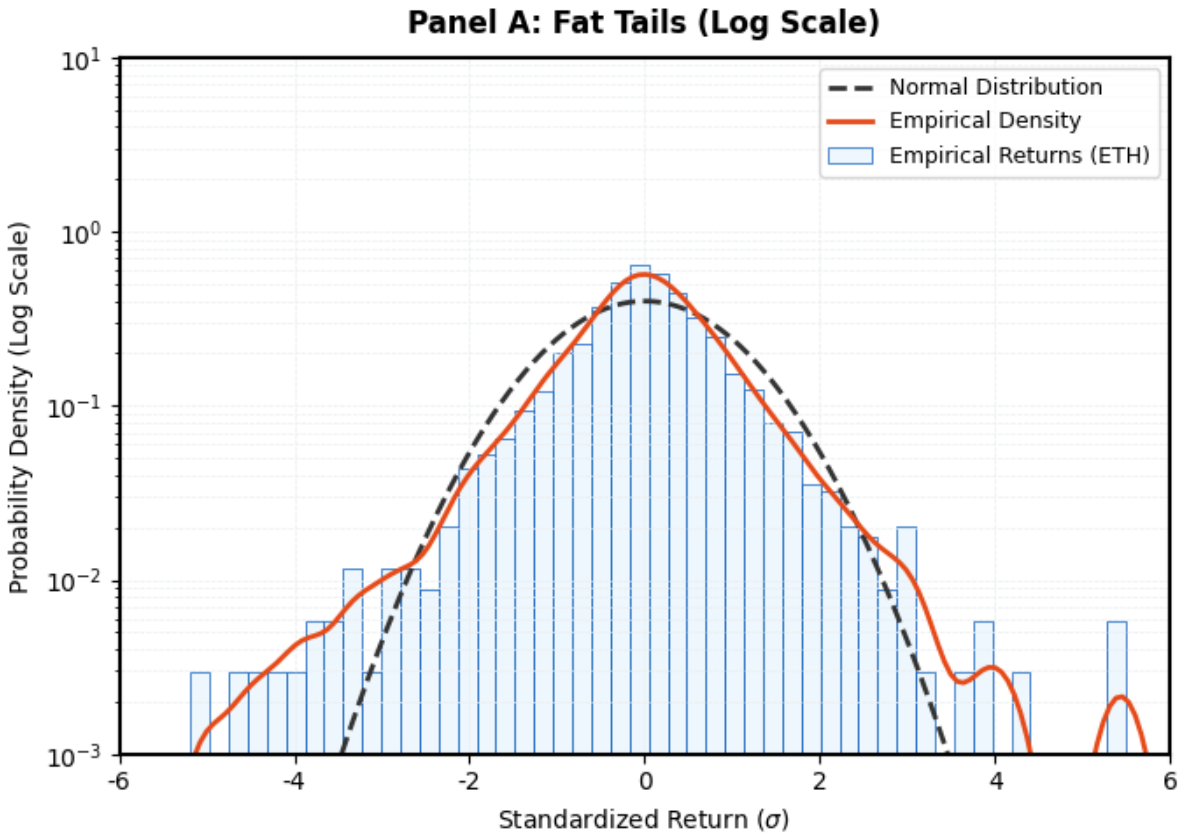
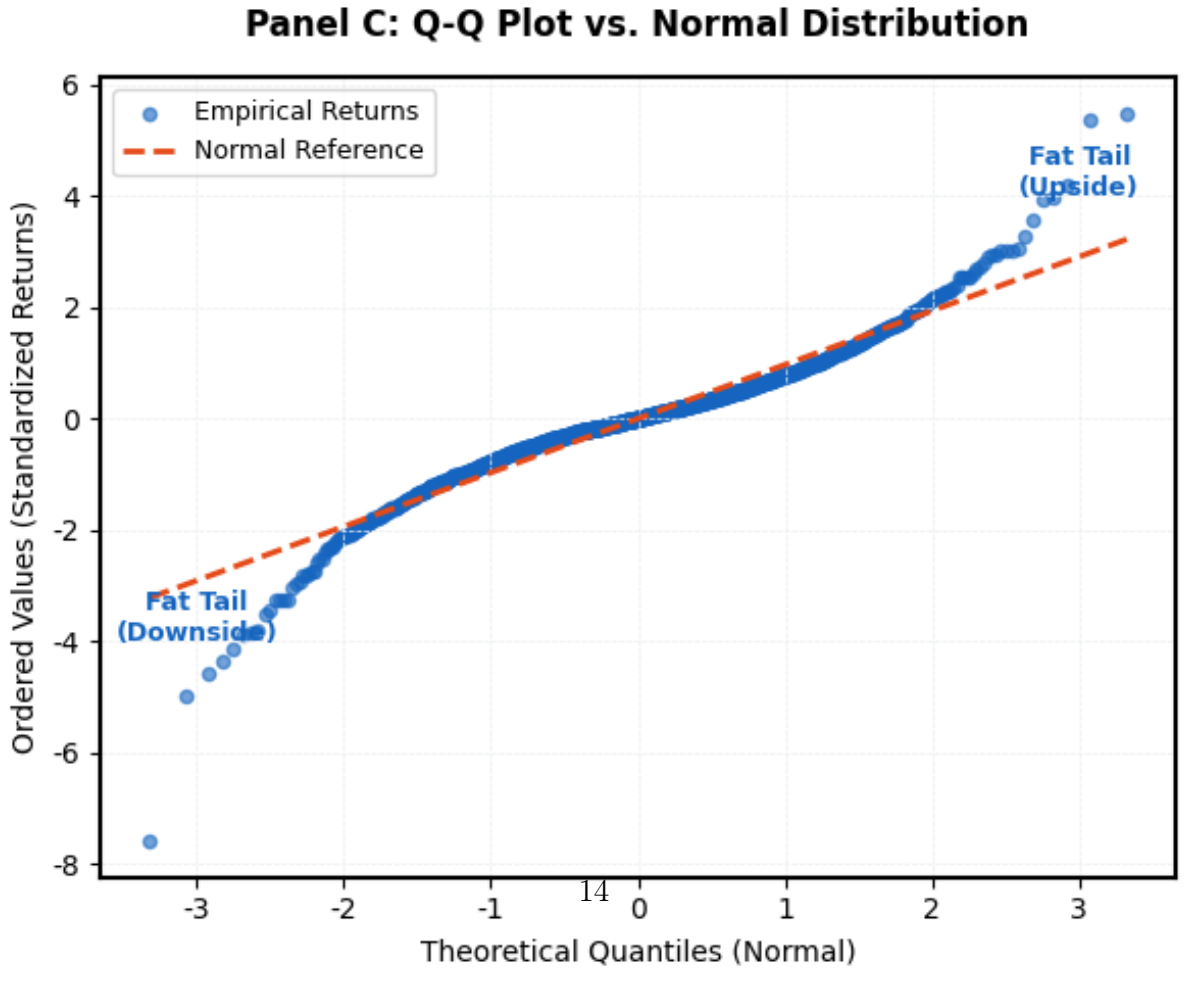


Figure 2: Density vs. Gaussian Rolling Kurtosis / Time Series

0.8



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 2.3 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.

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.

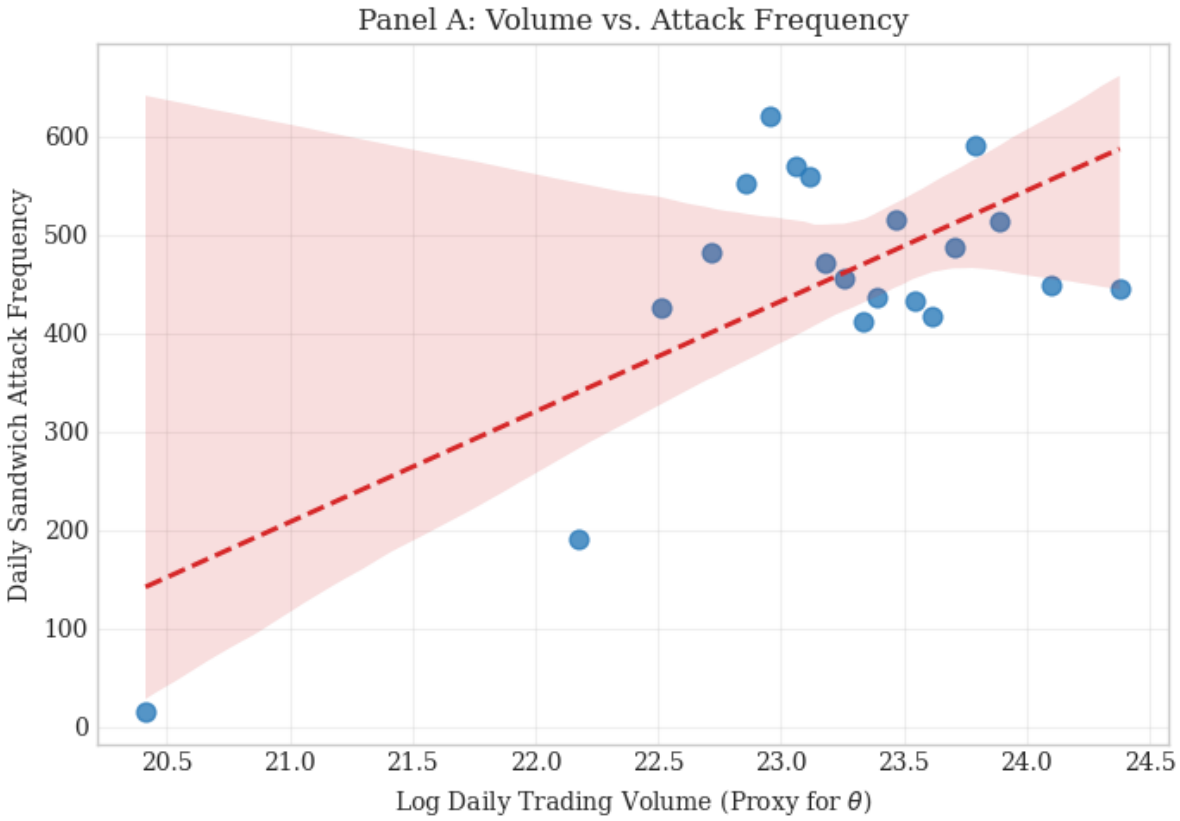


Figure 4: **Empirical Stylized Facts of Sandwich Attacks.** Panel A presents a binned scatter plot showing the positive correlation between log daily trading volume (proxy for execution-related order-flow intensity) and the frequency of sandwich attacks. Panel B illustrates an event study of the normalized attack frequency around high-volume shocks (defined as days exceeding the 90th percentile of volume). Shaded areas represent 95% confidence intervals.

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.

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 induced by the bot's optimal impulse intensity $a^\star$ in Chapter 4. Formally, one may start from a structural loss $T(\theta, a)$ and define

$$T(\theta) \equiv T(\theta, a^\star).$$

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 Chapter 4, the bot chooses an impulse intensity $a > 0$ generating a structural jump $\Gamma(p, \theta, a)$. The reduced-form jump in this chapter is

$$\gamma(p, \theta) \equiv \Gamma\big(p, \theta, a^\star(V, p, \Sigma, L)\big),$$

where $a^\star$ is the bot's optimal impulse control and V, Σ are its belief states.

Remark 1 (**From Microstructure to an Effective Intensity**). *In Chapter 4, predation arises endogenously from the MEV bot's impulse-control problem. Accordingly, the structural hazard is state-dependent, $\lambda^*(v, p, \Sigma, L)$, where (v, Σ) summarizes the bot's private belief state and L captures instantaneous liquidity conditions.*

In Chapter 3, the relevant object for the insider (and for pricing) is the predictable compensator *of predation arrivals under the insider's information set. Let $\mathcal{F}_t^I$ denote the insider's filtration (generated by publicly observed states such as $(P_s, L_s)_{s \le t}$ and the insider's own execution policy $(\theta_s)_{s \le t}$). The* projected *(insider-relevant) intensity is therefore*

$$\lambda_t^I \;\equiv\; \mathbb{E}\big[\lambda^*(V_t, P_t, \Sigma_t, L_t) \,\big|\, \mathcal{F}_t^I\big]\,. \tag{5}$$

To obtain the reduced-form representation used in Chapter 3, we invoke a time-scale separation. MEV strategies evolve at the block level, while the insider's execution problem operates on a slower horizon. Under this separation, the fast-moving belief state (V_t, Σ_t) *is well-approximated by a quasi-stationary distribution conditional on locally constant execution aggressiveness* θ_t*, so that*

$$\lambda_t^I \;\approx\; \lambda(\theta_t), \qquad \lambda(\theta) := \mathbb{E}[\lambda^*(V_t, P_t, \Sigma_t, L_t) \,|\, \theta_t = \theta]\,, \tag{6}$$

where $\lambda(\theta)$ *captures the systematic component of predation risk relevant for the insider's HJB and equilibrium pricing. For notational simplicity we suppress any residual dependence on* (P_t, L_t) *and write* $\lambda(\theta_t)$*. This motivates the maintained reduced-form restriction that* $\lambda(\theta)$ *is increasing and convex in* $|\theta|$*.*

The MEV Hazard Rate (Reduced-Form Controlled Intensity)

Predation arrivals are modeled by a point process N_t with 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 collapse ("liquidity trap"):*** $\lambda(\theta)\, e^{\alpha\Psi}\big(-J_p(t,p+\gamma)\big)\gamma_\theta(p,\theta)$.

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 cannot directly observe the fundamental value V or the insider's trading intent. Instead, it monitors the public mempool for the total order flow Y_t:

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

where B_t is a standard Brownian motion representing uninformed noise trading, assumed independent of W_t, and $\sigma_Y > 0$ is the volatility of noise flow.

Assumption 2 (Strategic Linearization)**.** *The bot assumes the insider employs a linear execution strategy:*

$$\theta_t \approx \beta(V - P_t), \qquad \beta > 0 \text{ constant over the bot's short window.}$$

Thus,

$$dY_t = \beta(V - P_t)\, dt + \sigma_Y\, dB_t.$$

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. Define the normalized innovation:

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

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

Posterior mean dynamics.

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

Equivalently,

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

Riccati equation. Since V is static, uncertainty decreases deterministically:

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

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 $\theta \approx \beta(v - p)$.

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,

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

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

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

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

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

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:

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

4.5 Optimal Strategy, Regions, and the Effective Intensity

The QVI partitions the state space into:

$$\mathcal{C}:\ \Phi > \mathcal{M}\Phi \quad \text{(wait and filter)}, \qquad \mathcal{A}:\ \Phi = \mathcal{M}\Phi \quad \text{(attack immediately)}.$$

Let the first hitting time of the attack region be

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

Bridge to Chapter 3 (reduced-form intensity). Define the effective intensity as the inverse conditional expected time to reach the attack region:

$$\lambda^\star(v, p, \Sigma, L) := \mathbb{E}\big[\tau_{\partial\mathcal{A}} \mid (v_0, p_0, \Sigma_0, L_0) = (v, p, \Sigma, L)\big]^{-1}.$$

Finally, the reduced-form objects entering Chapter 3 are induced by the bot's optimal intervention:

$$\gamma(p,\theta) \equiv \Gamma(p,\theta,a^\star(v,p,\Sigma,L)), \qquad T(\theta) \equiv T(\theta,a^\star(v,p,\Sigma,L)),$$

and the associated reduced-form intensity is summarized by $\lambda(\theta)$, increasing and convex in $|\theta|$.

5 Equilibrium via a Coupled HJB–QVI Fixed Point

This section characterizes a Markovian equilibrium of the strategic interaction between the insider (Chapter 3) and the MEV bot (Chapter 4). The equilibrium is expressed as a fixed point between: (i) the insider's optimal continuous execution policy $\theta^\star$ under predation risk, and (ii) the bot's optimal impulse policy $a^\star$ and its induced effective predation threat.

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 (from filtering)

The bot's belief state is

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

and the combined augmented state is

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

From impulse policy to reduced-form "arrival intensity"

In Chapter 4, the bot attacks when Ξ_t enters the attack region $\mathcal{A}$:

$$\tau_{\partial\mathcal{A}} = \inf\{t \geq 0 : \ \Xi_t \in \mathcal{A}\}.$$

Summarize the threat by an effective intensity proxy:

$$\lambda^\star(\xi) := \mathbb{E}[\tau_{\partial\mathcal{A}} \mid \Xi_0 = \xi]^{-1}, \qquad \xi = (v, \Sigma, p, L).$$

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

Under $\mathbb{P}$, the AMM price is driven by the insider's continuous execution flow and the bot's discrete attacks:

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

where N_t is a counting process of attacks with predictable intensity approximated by:

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}) \approx \lambda^\star(\Xi_t)\, dt.$$

Reduced-form jump objects induced by the bot:

$$\gamma(p, \theta) \equiv \Gamma(p, \theta, a^\star(v, \Sigma, p, L)), \qquad T(\theta) \equiv T(\theta, a^\star(v, \Sigma, p, L)).$$

5.3 Insider Optimality (HJB, Chapter 3 Object)

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

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

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

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

5.5 Markovian Equilibrium as a Fixed Point

Definition 1 (Markovian equilibrium)**.** *A Markovian equilibrium is a collection of objects* $\Theta^\star,\ \Phi^\star,\ a^\star,\ \lambda^\star,\ \gamma,\ T,\ \lambda(\theta)$ *such that:*

- ***Insider optimality:*** $\theta_t^\star = \Theta^\star(t,P_t;V,L_t)$ *solves the insider's HJB.*
- ***Bot optimality:*** $\Phi^\star$ *solves the QVI, yielding* $a^\star$ *and attack boundary* $\partial\mathcal{A}$*.*
- ***Consistency:*** $\gamma, T, \lambda(\theta)$ *are induced by* $a^\star$ *and rationalized by* $\lambda^\star(\Xi_t)$*.*

5.6 Existence

Assumption 3. *Controls are compact. Regularity: continuity, concavity/convexity conditions. No-Zeno:* $C_{gas} > 0$*. Well-posedness: unique viscosity solutions for HJB and QVI.*

Proposition 2. *Under Assumption 3, there exists a Markovian equilibrium.*

Proof strategy. Define a best-response mapping and apply Schauder's fixed-point theorem.

6 The Pricing Measure and Endogenous Excess Kurtosis

This section derives an equivalent pricing measure $\mathbb{Q}$ for the equilibrium price process and shows that MEV predation generates *endogenous leptokurtosis* (excess kurtosis) in risk-neutral returns.

6.1 Physical Dynamics and the Cox Approximation for Predation Arrivals

Under $\mathbb{P}$, the AMM price follows the reduced-form jump-diffusion implied by equilibrium:

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

Here W_t is a $\mathbb{P}$-Brownian motion, $\theta_t^\star = \Theta^\star(t, P_t; V, L_t)$, and

$$\gamma(p, \theta) \equiv \Gamma\big(p, \theta, a^\star(V_t, \Sigma_t, p, L_t)\big).$$

Cox (effective hazard) approximation. Represent attack arrivals by a Cox process with effective intensity:

$$\mathbb{P}(dN_t = 1 \mid \mathcal{F}_{t-}) \approx \lambda_t^{\mathbb{P}}\, dt, \qquad \lambda_t^{\mathbb{P}} \equiv \lambda^\star(\Xi_t), \qquad \Xi_t = (V_t, \Sigma_t, P_t, L_t).$$

Remark 2 (Interpretation)**.** *The Cox representation is a reduced-form device that preserves the key comparative statics: higher informativeness of insider flow accelerates learning and boundary-hitting, thereby increasing* $\lambda_t^{\mathbb{P}}$.

6.2 Pricing Kernel and Girsanov Transform for Point Processes

We construct an equivalent measure $\mathbb{Q} \sim \mathbb{P}$ under which the discounted price is a $\mathbb{Q}$-local martingale. Since jumps imply incompleteness, $\mathbb{Q}$ is selected by market risk preferences (e.g., LPs). Let $B_t = e^{rt}$ be the money-market account (constant short rate r). Define the Radon–Nikodym density process

$$\Lambda_t = \frac{d\mathbb{Q}}{d\mathbb{P}}\Big|_{\mathcal{F}_t}.$$

(We use Λ_t to avoid conflict with liquidity state L_t.)

Density process (explicit closed form)

Let ξ_t be the market price of diffusive risk and $\kappa_t > 0$ be the jump risk premium ratio (predictable). The Doléans–Dade exponential yields:

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

Equivalently, using the compensated process $\tilde{N}_t := N_t - \int_0^t \lambda_s^{\mathbb{P}} \, ds$,

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

Intensity transform under $\mathbb{Q}$. By Girsanov's theorem for point processes, the compensator transforms as:

$$\lambda_t^{\mathbb{Q}} = \kappa_t \, \lambda_t^{\mathbb{P}} = \kappa_t \, \lambda^\star(\Xi_t).$$

Remark 3. *If $\kappa_t > 1$, the market prices derivatives as if predation events were* more frequent *under $\mathbb{Q}$ than physically under $\mathbb{P}$.*

6.3 Determining the Jump Risk Premium κ_t: The LP Channel

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

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

6.4 Endogenous Excess Kurtosis under $\mathbb{Q}$

Let log return be

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

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

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

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

Cumulant dominance by jumps (short-horizon / high-jump-risk regime). A convenient approximation is:

$$\mathrm{Kurt}^{\mathbb{Q}}(R_T) - 3 \approx \frac{\mathbb{E}^{\mathbb{Q}}\left[\int_0^T \lambda_t^{\mathbb{Q}} \, \mathbb{E}^{\mathbb{Q}}[J_t^4 \mid \mathcal{F}_{t-}] \, dt \right]}{\mathrm{Var}^{\mathbb{Q}}(R_T)^2}.$$

Using $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^{\mathbb{P}} = \kappa_t \lambda^\star(\Xi_t)$,

$$\mathrm{Kurt}^{\mathbb{Q}}(R_T) - 3 \text{ increases with } \int_0^T \underbrace{\kappa_t}_{\text{pricing amplification}} \cdot \underbrace{\lambda^\star(\Xi_t)}_{\text{endogenous physical hazard}} \cdot \underbrace{\mathbb{E}[J_t^4 \mid \mathcal{F}_{t-}]}_{\text{endogenous jump size}} \, dt.$$

Small-jump approximation . If $\gamma_t \ll P_t$, then $J_t \approx \gamma_t / P_t$, and the numerator behaves like $\int \lambda_t^{\mathbb{Q}} \mathbb{E}[(\gamma_t/P_t)^4] \, dt$.

Proposition 4 (Endogenous leptokurtosis: structural multiplier effect)**.** *In equilibrium, excess kurtosis under $\mathbb{Q}$ is amplified by two multiplicative channels:*

- ***Physical endogeneity:*** *$\theta_t^\star$ increases informativeness and boundary-hitting frequency, raising $\lambda_t^{\mathbb{P}} = \lambda^\star(\Xi_t)$.*
- ***Pricing amplification:*** *LP risk aversion implies $\kappa_t > 1$, so $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^{\mathbb{P}}$.*

Therefore, in regimes where diffusive volatility is low relative to jump risk and/or $\kappa_t \lambda_t^{\mathbb{P}}$ is high, the jump-driven fourth moment grows faster than the variance squared, generating pronounced leptokurtosis.

Remark 4 (What we do not claim)**.** *This result establishes thicker tails / higher kurtosis under $\mathbb{Q}$ via higher jump intensity and endogenous jump sizes. It does not require (nor*

automatically imply) a change in the exponential tail rate *of jump sizes unless the measure change also tilts the jump-size distribution (e.g., via κ_t depending on marks).*

6.5 Conclusion of Section 6

By explicitly constructing the jump-diffusion measure change, MEV acts as an *endogenous rare-event generator*:

- **Mechanism:** Strategic interaction creates physical predation jumps with state-dependent hazard $\lambda_t^{\mathbb{P}} = \lambda^\star(\Xi_t)$ and state-dependent jump sizes $\gamma(P_{t-}, \theta_t^\star)$.
- **Pricing:** LP aversion to toxic flow implies $\kappa_t > 1$, amplifying jump arrivals under the pricing measure:
$$\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^{\mathbb{P}}.$$
- **Outcome:** Risk-neutral returns exhibit elevated kurtosis and thicker tails because jump intensity and jump magnitude are both endogenously magnified.

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 "liquidity trap/safe-harbor" 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 enforcement; 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 liquidity trap, an intervention boundary for the bot, and priced tail risk consistent with a 'privatized enforcer' 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 $g = (\lambda, \gamma, \mathcal{T})$. 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,

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

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 and sensitivity ranges in Appendix Table. 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 $g = (\lambda, \gamma, \mathcal{T})$, 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:

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

where $T(\cdot)$ denotes the best-response pipeline defined by the coupled HJB–QVI system. 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 7.1 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 trading region ("liquidity trap/safe harbor"), while eroding continuation value. This baseline visualization motivates the comparative statics in Section 7.2.

Economically, the Liquidity Trap in Figure 7.1 represents the **deterrence effect** of privatized enforcement. Unlike SEC investigations which occur with a lag, the MEV bot acts as an **always-on, algorithmic police force**. The 'Value Erosion' (Panel B) is effectively the **expected fine** levied by the market. When the insider stays within the Safe Harbor, they are effectively complying with the 'implied law' of the mempool to avoid the penalty.

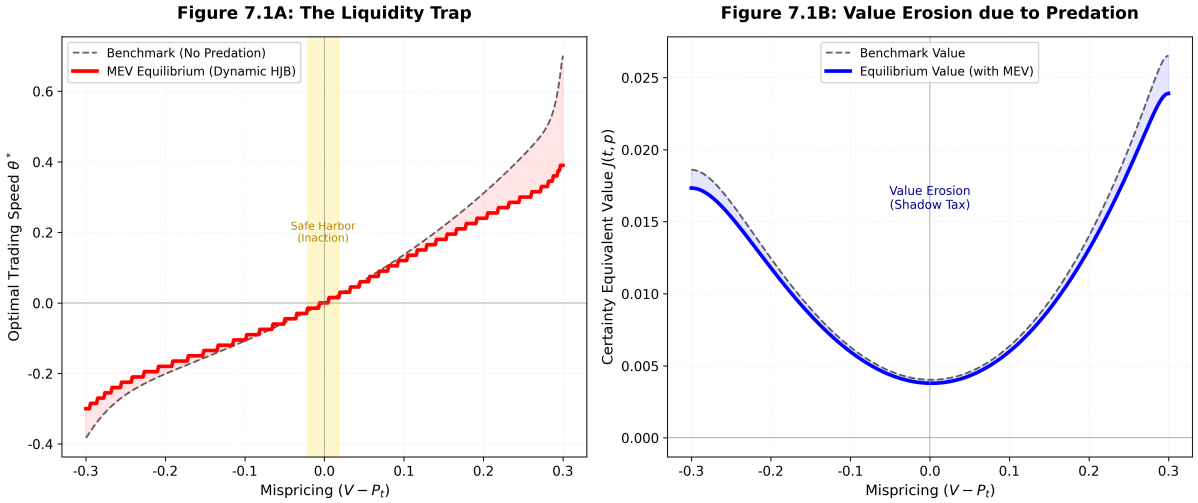


Figure 5: Baseline Equilibrium. Panel A shows the optimal trading speed θ^* versus mispricing, illustrating the safe harbor (inaction) region. Panel B shows the certainty equivalent value $J(t, p)$, highlighting value erosion (shadow tax) relative to the benchmark.

7.1.6 Robustness and interpretation

Because the numerical analysis is intended to illustrate mechanisms rather than fit data, we emphasize qualitative robustness. Appendix D 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 7.1–7.4 remain: predation risk compresses the insider's execution rule into a liquidity-trap region; higher fear (larger α) and stronger detection technology (larger λ_1) deepen the trap; 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 7.1 establishes the baseline mechanism: predation risk compresses the insider's equilibrium execution rule and creates a near-zero trading region around $m = V - p = 0$ ("liquidity trap/safe harbor"), while reducing continuation value $J(0, p)$. Building on this baseline, this subsection quantifies how the trap deepens 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 $\Psi(\theta, p) = \mathcal{T}(\theta) + J(t, p) - J(t, p + \gamma(p, \theta))$.

7.2.1 Measuring the Liquidity Trap

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 liquidity trap 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 a liquidity trap 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 Appendix Table D. 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 7.2A reports the heatmap of $W(\alpha,\lambda_1)$ across the parameter grid, and Figure 7.2B plots cross-sections of W as a function of α for representative values of λ_1.

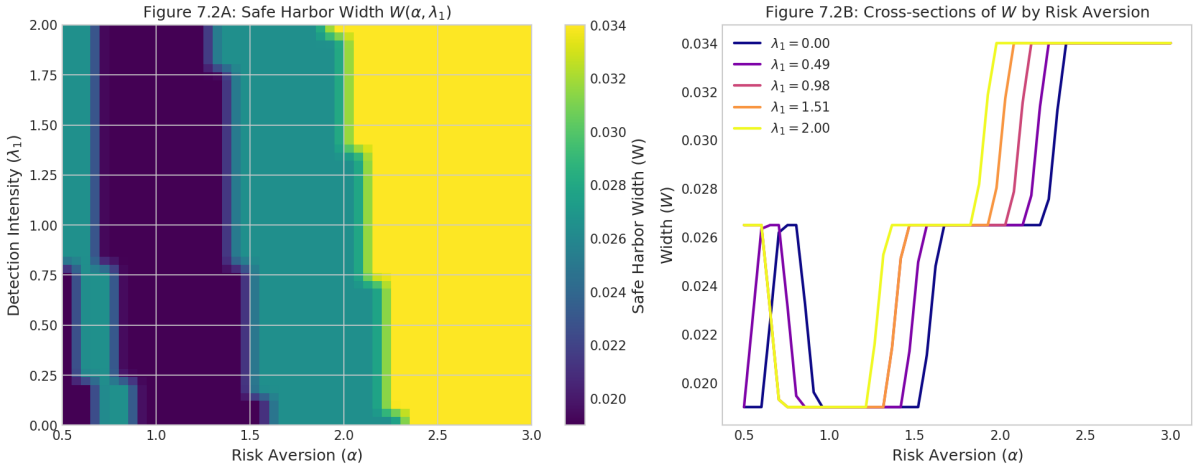


Figure 6: Safe Harbor Width. Panel A plots a heatmap of the safe-harbor width $W(\alpha,\lambda_1)$. Panel B reports cross-sections of W as a function of α. Higher risk aversion α and stronger detection technology λ_1 expand W, indicating a wider region of mispricing in which optimal execution is negligible.

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 liquidity trap 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 D.

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 7.3 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.

To keep the illustration transparent and computationally light, we summarize the QVI-implied intervention condition by a reduced-form boundary proxy (C_{gas}, σ_Y) that preserves the model's monotone comparative statics.

Connection to trader safe harbor (link to Figure 7.2). 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 7.2. 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

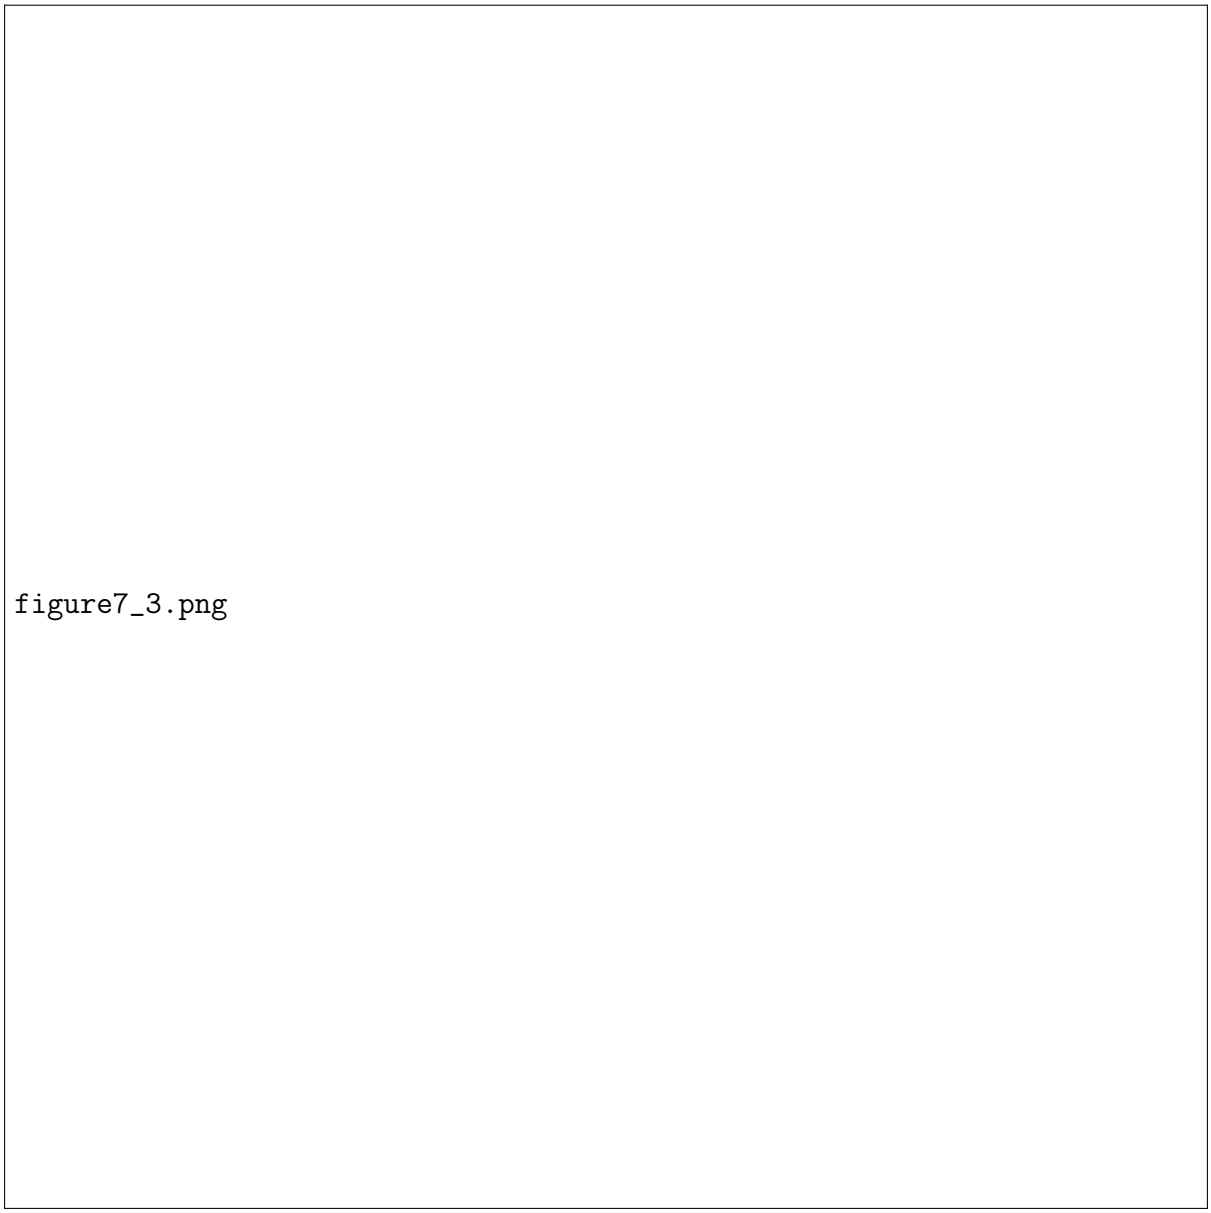


Figure 7: Predation Boundary. Higher fixed per-attack costs (C_{gas}) and higher observation noise (σ_Y) raise the threshold required for profitable intervention, shifting the predation boundary outward and expanding the safe region.

in which the bot rationally waits—though higher σ_Y may also imply slower information revelation.

7.4 The Market (Fat Tails): Liquidity-Trap Peak and MEV-Driven Tail Risk

This subsection connects the microstructure mechanisms highlighted in Sections 7.2–7.3 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 ("liquidity trap/safe harbor"), reducing typical price impact for most 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** 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 modeled with controlled intensity $\lambda(\theta)$, and jump sizes are governed by $\gamma(p, \theta)$. Under the pricing measure Q, the compensator transforms so that the effective jump intensity becomes $\lambda_t^Q = \kappa_t \lambda_t^P$, 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 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 liquidity-trap 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 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

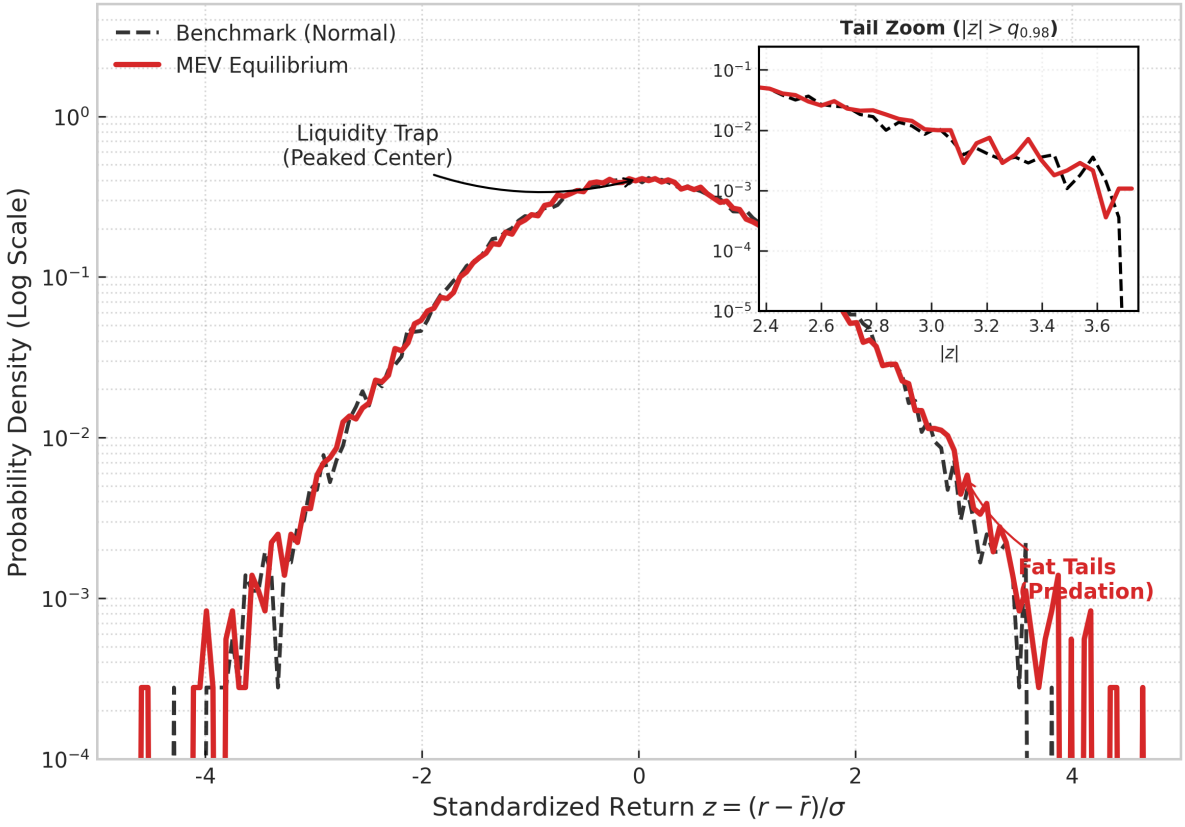


Figure 8: Money Plot: Liquidity-Trap Peak and MEV-Driven Fat Tails. Standardized log-density of terminal log returns $z = (r - \bar{r})/\sigma$, comparing a simulated equilibrium MEV jump-diffusion under the pricing measure Q (solid) to a matched normal distribution (dashed). The inset zooms into the tail region $|z| > q_{0.98}$.

7.4.3 Results: sharp center and fat tails

Figure 7.4 delivers a single unified visual message:

Sharper center (liquidity-trap signature). Relative to the matched normal benchmark, the MEV jump-diffusion exhibits **excess mass near** $z = 0$. This is the distribu-

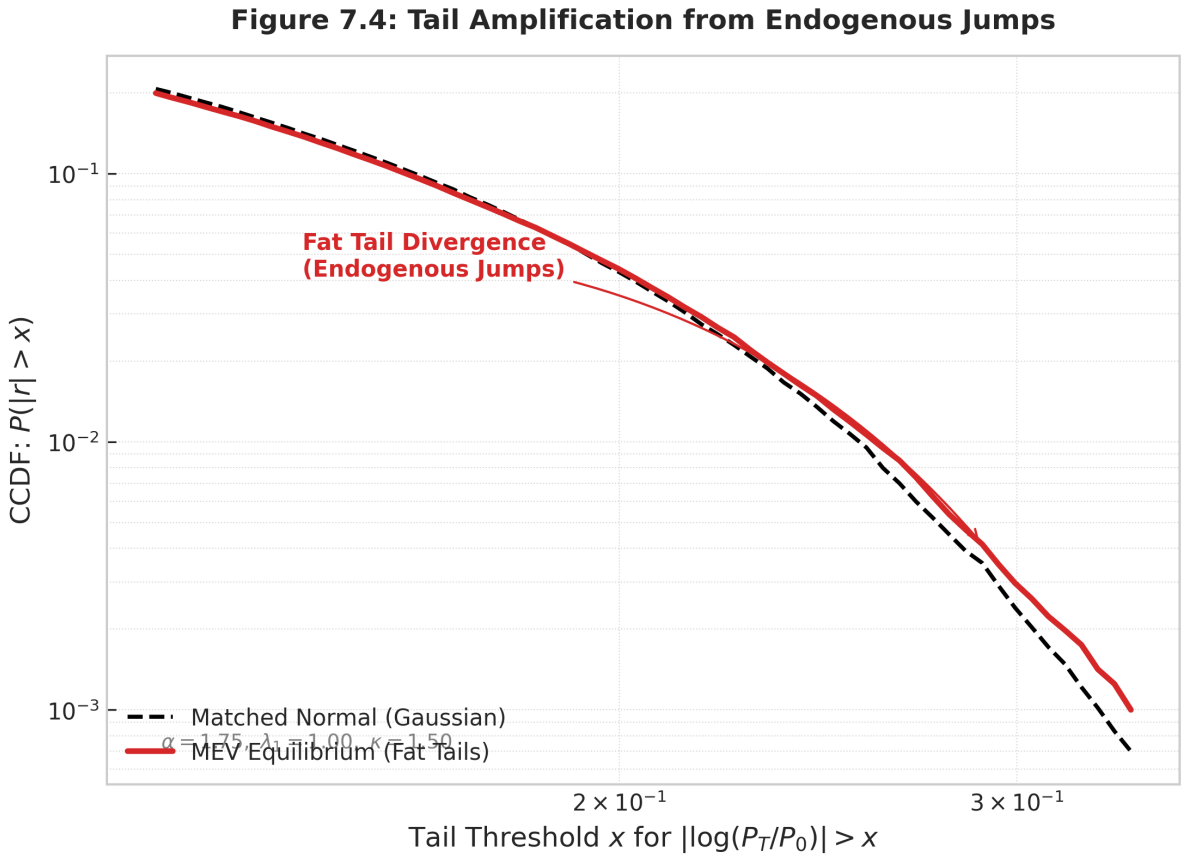


Figure 9: Tail Amplification from Endogenous Jumps. Endogenous Leptokurtosis and Fat Tails.

tional imprint of the liquidity-trap mechanism documented earlier: when predation risk suppresses optimal execution around small informational advantages, most price paths experience relatively muted incremental impact, concentrating realized returns around zero.

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 $\lambda_t^Q = \kappa_t \lambda_t^P$, which increases the effective frequency of predation events under 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: Trap → 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 trap), the trap is **monotonically deepened** 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 Privatized Enforcement

8.1 Core Thesis and Big-Picture Contribution

This paper's central claim is institutional: **in DeFi, where there is no SEC-style public enforcer, MEV searchers function as a privatized enforcer**. The transparent transaction-pool environment turns would-be informed trading into an "offense" that can be detected and punished by profit-seeking bots in real time, effectively substituting **code-based punishment** for legal sanction. In Becker-style language, deterrence arises from an endogenous expected penalty; the key difference is that the penalty is imposed by competitive private actors via transaction ordering rather than by a public regulator via fines or prosecution.

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, predation enters through a controlled hazard rate and jump losses, so the insider's certainty-equivalent HJB features an exponential jump penalty that acts as an **endogenous "shadow tax"** on execution. On the bot side, the searcher solves a stationary QVI that endogenously generates an attack/wait region and a predation boundary. Together, these objects constitute a micro-founded enforcement regime: enforcement intensity, punishment severity, and the boundary of intervention are all generated inside the equilibrium rather than assumed.

8.2 Mechanism: The Shadow Tax and the Liquidity Trap 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 $\Psi(\theta,p) = \mathcal{T}(\theta) + J(t,p) - J(t,p+\gamma(p,\theta))$. This generates a generalized "MEV shadow tax" that decomposes into (i) visibility/detection (via $\lambda'(\theta)$), (ii) sandwich geometry (via $\mathcal{T}'(\theta)$), and (iii) endogenous state collapse (via the non-local value erosion term).

The outcome is a **safe-harbor / liquidity-trap region** around small mispricing in which optimal execution is negligible—precisely the economic signature of deterrence when expected punishment dominates marginal informational rents. This deterrence logic is uniquely DeFi-native: the "probability of apprehension" is represented by a state- and action-dependent intensity $\lambda(\theta)$ that increases with execution aggressiveness, capturing the idea that aggressive trades are more detectable in transparent transaction pools.

Unlike traditional markets where front-running is illegal and enforcement is external to the trading mechanism, in DeFi the mempool's transparency and fee-based priority rules make predatory reordering technically feasible and economically scalable, enabling enforcement by bots rather than by courts. Finally, this "privatized enforcement" does not merely alter execution strategies; it leaves a distinct **distributional footprint** at the market level. Predation induces endogenous, state-dependent jump risk in the price process. Under the pricing measure, the jump intensity is further magnified by the jump risk premium, mapping micro-level "code enforcement" into a return distribution characterized by a sharper center and fatter tails.

Thus, the hard-core conclusion we invite readers to take away is this: within the institutional environment of DeFi, code serves not merely as a tool for matching and settlement, but as a regulatory function. **"Code > Law"** is not a slogan, but an equilibrium implication emerging endogenously from transparent mempools and strategic predation.

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." *IACR Cryptology ePrint Archive*, 2024/331.

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*.

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. 2019. "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges." *arXiv preprint*, arXiv:1904.05234.

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 function $T(\theta)$, we denote the fixed-point operator by $\mathcal{T}$. 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., (1; 2; 3)). 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, 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 (2; 4; 3)).

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, monotonicity (increasing), and convexity 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 $-\lambda(\theta)\frac{1}{\alpha}\left(e^{\alpha\Psi^g(\theta,p)} - 1\right)$.

Lemma A.2 (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. This is standard from the dynamic programming principle and viscosity-solution theory: existence follows from DPP, while comparison implies uniqueness on compact domains with continuous coefficients. On the truncated state domain, we interpret the PDE in the *state-constraint viscosity* sense (a standard device for technical truncations in control problems), so no explicit Dirichlet/Neumann boundary data are required (see (1; 2; 4)). □

Assumption 4 (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.3 (Unique insider best response)**.** *Under Assumption 4, 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 (2; 4). □

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 introduce a continuous summary operator $\mathcal{B}$ mapping feedback strategies to a scalar slope:

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

and assume:

Assumption 5 (Continuity of belief summary)**.** $\mathcal{B}$ *is continuous w.r.t. the sup norm on admissible feedback rules.*

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

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

with intervention operator

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

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

Assumption 6 (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.4 (QVI well-posedness and unique impulse)**.** *Under (A4), $C_{\text{gas}} > 0$, compact $\mathcal{U}$, continuity/growth bounds on (R, Γ), and Assumption 6, 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 (3) and the control/viscosity discussion in (2)). Uniqueness of the impulse follows under Assumption 6. □

Lemma A.5 (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 (3)). □

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

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

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

$$\gamma'(p,\theta) \;:=\; \Gamma(p,\theta,a^g(\cdot)), \qquad T'(\theta) \;:=\; T_{\text{struct}}(\theta, a^g(\cdot)),$$

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

Effective hazard proxy and projection.

Assumption 7 (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, $\Gamma_\theta := \{\xi \in \mathcal{X} : \beta^g(v-p) = \theta\}$ is nonempty for all $\theta \in \Theta$.*

Let $\tau_{\partial\mathcal{A}^g}$ denote the first hitting time of the intervention region $\mathcal{A}^g$ by the bot's state process Ξ_t. Define the expected hitting time functional

$$u^g(\xi) \;:=\; \mathbb{E}[\tau_{\partial\mathcal{A}^g} \mid \Xi_0 = \xi]\,, \qquad \xi \in \mathcal{X}.$$

Fix a small constant $\varepsilon > 0$ and define the (softened) state-dependent effective hazard proxy by

$$\lambda_g^\star(\xi) \;:=\; \frac{1}{\max\{u^g(\xi),\, \varepsilon\}}, \qquad \xi \in \mathcal{X}.$$

Assumption 8 (Regularity and stability of expected hitting times)**.** *For each $g \in \mathcal{A}$, the expected hitting time $u^g(\xi)$ is finite and continuous on $\mathcal{X}$. Moreover, if $g_n \to g$ uniformly in $\mathcal{A}$, then $u^{g_n} \to u^g$ uniformly on $\mathcal{X}$.*

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

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

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

Assumption 9 (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, increasing, convex, and* 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.6 (Continuity of the projection)**.** *Under Assumptions 7 and 8, 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 $\Gamma_\theta := \{\xi \in \mathcal{X} : \beta^g(v - p) = \theta\}$. By Assumption 7, Γ_θ is nonempty; it is compact-valued with closed graph on compact domains. By Assumption 8, u^g is continuous, hence $\lambda_g^\star(\xi) = 1/\max\{u^g(\xi), \varepsilon\}$ is continuous on $\mathcal{X}$. Therefore Berge's maximum theorem implies continuity of the value function $\theta \mapsto \sup_{\xi \in \Gamma_\theta} \lambda_g^\star(\xi)$ (see (5)). The outer truncation by $\min\{\bar{\lambda}, \cdot\}$ preserves continuity, so $\tilde{\lambda}$ is continuous. Finally, continuity of Π_λ in Assumption 9 yields continuity of $\lambda' = \Pi_\lambda(\tilde{\lambda})$. □

Definition (Best-response mapping). For any $g \in \mathcal{A}$, define $\mathcal{T}(g) = (\lambda', \gamma', T')$ 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 $\mathcal{T}$; existence by Schauder

Lemma A.7 (Self-mapping)**.** *For all $g \in \mathcal{A}$, $\mathcal{T}(g) \in \mathcal{A}$.*

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 (3)). Continuity and boundedness of λ' follow from Lemma A.6 and the cap by $\bar{\lambda}$. The intensity λ' satisfies (A1) by construction since $\lambda' = \Pi_\lambda(\tilde{\lambda})$ and Assumption 9 guarantees that Π_λ maps any continuous input back into the admissible shape-restricted class. Hence $\mathcal{T}(g)$ satisfies (A1)–(A4), i.e., $\mathcal{T}(\mathcal{A}) \subseteq \mathcal{A}$. □

Lemma A.8 (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 4, $\theta^{g_n} \to \theta^g$ uniformly.*

Justification. Viscosity solutions are stable under uniform perturbations of coefficients and nonlocal terms when comparison holds, using half-relaxed limits (see (1)). Uniqueness of the maximizer yields continuous dependence of the argmax on the parameters (maximum theorem logic; (5)). □

Lemma A.9 (Continuity of belief summary)**.** *If $g_n \to g$, then $\beta^{g_n} \to \beta^g$.*

Proof. Immediate from Assumption 5 and Lemma A.8. □

Lemma A.10 (Bot stability)**.** *Let $g_n \to g$. Then $\Phi^{g_n} \to \Phi^g$ uniformly on $\mathcal{X}$, and under Assumption 6, $a^{g_n} \to a^g$ 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 (3) and the viscosity framework in (1)). □

Lemma A.11 (Continuity of λ')**.** *Let $g_n \to g$ uniformly in $\mathcal{A}$. Under Assumptions 8, 7, and 9, we have $\lambda'_n \to \lambda'$ uniformly on Θ.*

Justification. By Assumption 8, $u^{g_n} \to u^g$ uniformly on $\mathcal{X}$. Since the map $x \mapsto 1/\max\{x, \varepsilon\}$ is continuous (indeed Lipschitz) on $[0, \infty)$, we have $\lambda^\star_{g_n} \to \lambda^\star_g$ uniformly on $\mathcal{X}$. Then Berge's maximum theorem yields $\tilde{\lambda}_n \to \tilde{\lambda}$ uniformly on Θ (with nonemptiness ensured by Assumption 7), and continuity of Π_λ in Assumption 9 implies $\lambda'_n = \Pi_\lambda(\tilde{\lambda}_n) \to \Pi_\lambda(\tilde{\lambda}) = \lambda'$ uniformly. □

Proposition 5 (Continuity of $\mathcal{T}$)**.** *$\mathcal{T} : \mathcal{A} \to \mathcal{A}$ is continuous.*

Proof. Combine Lemmas A.8–A.11 componentwise. □

Theorem A.12 (Schauder fixed point)**.** *There exists $g^\star \in \mathcal{A}$ such that $\mathcal{T}(g^\star) = g^\star$.*

Proof. By Lemma A.1, $\mathcal{A}$ is nonempty, compact, and convex. By Proposition 5 and Lemma A.7, $\mathcal{T} : \mathcal{A} \to \mathcal{A}$ is continuous. Hence Schauder's fixed point theorem implies existence of a fixed point (see (6)). □

Completion of the proof of Proposition 2. Let $g^\star = (\lambda^\star, \gamma^\star, T^\star)$ be a fixed point of $\mathcal{T}$. Define $\theta^\star = \mathrm{BR}_I(g^\star)$, $\beta^\star = \mathcal{B}(\theta^\star)$, and $(a^\star, \mathcal{A}^\star) = \mathrm{BR}_B(\beta^\star)$. By construction of $\mathcal{T}$, the induced reduced-form triple from $(a^\star, \mathcal{A}^\star)$ coincides with $g^\star$. 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 2 (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.13 (HJB verification (finite horizon, CARA case))**.** *Fix a reduced-form triple g in the admissible reduced-form class (as in Appendix A). 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., (2; 4). □

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

Definition 3 (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.14 (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., (3). □

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 A.*

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 $\mathbb{P}$-intensity λ^P_t (predictable), i.e. $\tilde{N}_t := N_t - \int_0^t \lambda^P_s\, ds$ is a $\mathbb{P}$-local martingale. All processes are defined with respect to $\{\mathcal{F}_t\}$, and λ^P_t, ξ_t, κ_t are understood as $\mathcal{F}_{t-}$-predictable processes. In the main text we take $\lambda^P_t = \lambda^\star(\Xi_t)$ under the Cox/effective-hazard representation.

Assumption 10 (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)*** *$\int_0^T \lambda_t^P \, dt < \infty$ a.s. and $\int_0^T |\kappa_t - 1| \, \lambda_t^P \, dt < \infty$ a.s.*

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

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

Equivalently, in differential form,

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

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

Lemma C.15 (Well-defined equivalent measure and intensity transform)**.** *Under Assumption 10, 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

$$\lambda_t^{\mathbb{Q}} = \kappa_t \, \lambda_t^P, \qquad 0 \leq t \leq T,$$

i.e. $\tilde{N}_t^{\mathbb{Q}} := N_t - \int_0^t \lambda_s^{\mathbb{Q}} \, ds$ 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 $\int_0^T (\kappa_t - 1)\lambda_t^P dt$ 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). The product of two strictly positive martingales is a strictly positive 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 $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^P$. □

Remark 6 (Interpretation in the main text)**.** *In the paper we write $\lambda_t^P = \lambda^\star(\Xi_t)$ under the Cox/effective-hazard approximation, and interpret $\kappa_t > 1$ as a jump risk premium amplification channel. Under Lemma C.15, jump arrivals are priced as if more frequent under $\mathbb{Q}$: $\lambda_t^{\mathbb{Q}} = \kappa_t \lambda_t^P$ 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 7.3). 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.

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.

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.

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

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

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

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

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:

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

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 $[p_{\min}, p_{\max}]$.

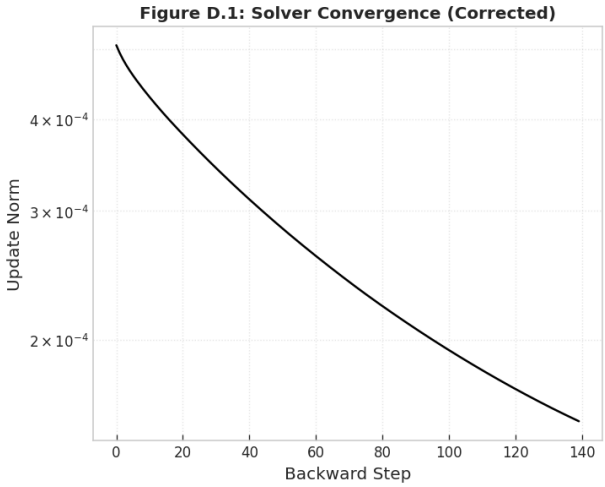
D.5 Convergence diagnostics (stopping rules)

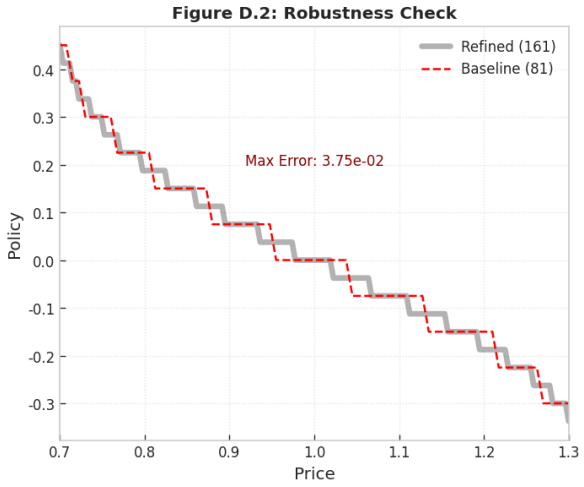
HJB tolerance. We monitor sup-norm convergence and stop policy iteration when

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

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

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





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 7.3 boundary proxy: definition and interpretation

Proxy definition. To keep Figure 7.3 transparent and computationally light, we summarize the QVI intervention boundary by the reduced-form 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 proxy is intended to preserve comparative statics rather than to replace a full QVI boundary surface.

References

(Appendix A-C)

[1] 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.

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

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

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

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

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