

Autodeleveraging: Impossibilities and Optimization

Tarun Chitra
Gauntlet
tarun@gauntlet.xyz

December 9, 2025

Abstract

Autodeleveraging (ADL) is a last-resort loss socialization mechanism for perpetual futures venues. It is triggered when solvency-preserving liquidations fail. Despite the dominance of perpetual futures in the crypto derivatives market, with over \$60 trillion of volume in 2024, there has been no formal study of ADL. In this paper, we provide the first rigorous model of ADL. We prove that ADL mechanisms face a fundamental *trilemma*: no policy can simultaneously satisfy exchange *solvency*, *revenue*, and *fairness* to traders. This impossibility theorem implies that as participation scales, a novel form of *moral hazard* grows asymptotically, rendering ‘zero-loss’ socialization impossible. On the positive side, we show that three classes of ADL mechanisms can optimally navigate this trilemma to provide fairness, robustness to price shocks, and maximal exchange revenue. We analyze these mechanisms on the Hyperliquid dataset from October 10, 2025, when ADL was used repeatedly to close \$2.1 billion of positions in 12 minutes. By comparing our ADL mechanisms to the one used in practice on Hyperliquid and Binance, we demonstrate empirically that Hyperliquid’s production queue overutilized ADL by $\approx 28\times$ relative to our optimal policy, imposing roughly \$653 million of unnecessary haircuts on winning traders. This comparison also suggests that Binance overutilized ADL far more than Hyperliquid. Our results both theoretically and empirically demonstrate that optimized ADL mechanisms can dramatically reduce losses of trader profitability while maintaining exchange solvency.

1 Introduction

Perpetual futures (or simply, perpetuals) are by far the most popular crypto derivatives contract. These contracts allow for duration-independent hedging of cryptocurrencies, such as Bitcoin, Ethereum or Solana. That is, unlike traditional future contracts, perpetuals are not subject to a fixed expiration date. Instead, perpetuals are similar to *contracts for difference*, where market participants repeatedly send each other payments depending on how the spot price trades relative to the futures price. For instance, if the spot price is higher than the futures price, then traders with short futures positions pay the traders with long futures until the prices are equalized. We note that while Robert Shiller’s 1993 proposal for perpetual futures provides the genesis of this idea, the BitMEX exchange’s 2016 launch of the XBTUSD perpetual swap became the first live implementation [64, 103, 104].

Perpetual futures have a higher notional trading volume than spot cryptocurrency volume. This is akin to index futures for US stocks, where there is far more volume traded in the futures contracts than the underlying spot assets, with the S&P 500 futures index future trading roughly ten times more volume than the spot asset [34]. In 2024, centralized crypto exchanges such as Binance and Bybit processed nearly \$58.5 trillion (notional) in perpetuals trades [36]. On the other

hand, centralized exchanges processed nearly \$17.4 trillion of spot cryptocurrency trades [37], which is a 3.3x ratio of perpetuals to spot volume.

The main reason for the elevated usage of perpetual futures, much like their index futures cousins, is due to the high level of leverage they offer. In particular, the capital requirements for trading perpetual futures are much lower than spot trading. Moreover, the continuous nature of perpetual futures allows for easier margin management than duration-based futures. This is illustrated via the large gap between the maximum leverage offered by perpetuals exchanges (up to 125x [21]) versus the index futures on the Chicago Mercantile Exchange (around 10-15x leverage for index futures [35]).

Centralized vs. Decentralized Exchanges. Historically, the majority of perpetuals trading volume has been concentrated on centralized exchanges (CEXes). These venues, which include Binance, Bybit, and Coinbase, require users to deposit collateral that is custodied by the exchange. The exchange is responsible for ensuring that users' positions are solvent and that their collateral is not used to cover losses from other traders. In practice, this involves traders having to trust the exchange to not use their collateral for other purposes. In 2022, centralized exchange FTX was found to have utilized customer collateral for other purposes, leading to multi-billion-dollar losses [55].

In response, there has been an increasing trend in the usage of decentralized exchanges (DEXes) for trading perpetual futures. These exchanges provide far more transparency and auditability into the mechanics of the exchange. In particular, user's can always verify what their collateral is being used for and what positions it is collateralizing. Moreover, these exchanges are permissionless and pseudonymous, meaning anyone can trade on these exchanges without explicitly revealing their identity.

The transparent nature of decentralized exchanges does have, however, a downside. Most positions in decentralized exchanges need to be fully overcollateralized, which generally means traders can take less leverage than they would on a centralized exchange. Generally speaking, this is because the exchange doesn't know the identity of a trader and can only enforce global collateral invariants (*e.g.* assets > liabilities) by forcing users to post more collateral than if the exchange knew their identity. Early decentralized perpetuals exchanges, such as perpetual protocol, were unable to offer much more than 5x leverage.

However, in recent years, hybrid models of decentralized exchanges where withdrawals and deposits are permissionless, but certain exchange operations are centralized, have become popular. The most popular exchange of this form is Hyperliquid, which has helped increase the market share of perpetuals volume on DEXs from roughly 1% in 2024 to over 8% in 2025. These hybrid models allow for the exchange to use cryptographic commitments to enforce collateral rights while also allowing for oracles, markets, and solvency to be controlled by a smaller set of participants. This allows the exchange to avoid issues that fully open decentralized exchanges have, where weak collateral can be added to an exchange that cannot be sold or liquidated successfully during a market crash to ensure solvency.

More concretely, there have been recent incidents on perpetual DEXs that highlight how novel market structures can be exploited or stressed in practice. These include but are not limited to: Hyperliquid's XPL pre-market price spike and the JellyJelly attack [45, 86, 87]. These episodes led to emergency safeguards and community post-mortems, underscoring the need for robust oracle design, listing policies, and circuit breakers to mitigate manipulation. Such emergency responses would have been nearly impossible in the fully decentralized setting, demonstrating the value of hybrid models of decentralization in perpetual futures.

Liquidations. Despite the empirical prominence of perpetuals markets, there has been little formal study of the stability and robustness of these markets relative to the large body of research into spot cryptocurrency trading. For instance, the main price stability mechanism used for perpetuals is the *funding rate*, a continuous payment stream between the long and short positions. There has been some study of how to replicate funding rates via other financial instruments [2, 9, 66], but little to no empirical study of how well such approximations work. Even more surprising is the lack of formal study of how robust perpetual exchanges are under high leverage conditions, despite their usage being driven by high leverage traders.

A crucial component to ensuring the stability of perpetual futures markets with high leverage is *liquidation*. This is the process of closing a user’s position that is worth less than the cash collateral they have posted to the exchange. Perpetuals exchanges rely on liquidation mechanisms to ensure that they stay solvent — that is, the assets they hold are greater than the liabilities they owe to traders. To the author’s knowledge, there has been no formal study of liquidations in perpetuals exchanges, despite there being a large body of work studying liquidations in overcollateralized cryptocurrency lending (e.g. [72, 90, 95, 108]).

Liquidations are executed by taking the collateral of a trader and selling it in a manner that minimizes losses that the exchange realizes. In centralized exchanges, liquidations are usually executed by the exchange itself in its own spot markets. For instance, suppose that a user uses \$1,000 worth of collateral to open a long position with 10x leverage. This implies that if the price falls from the initial price when the contract was opened, p_0 , by roughly 10%, then the user’s position will have a value of -\$1,000, so that user’s net position is \$0. If the price decreases beyond this point, then the user’s position will have a negative value, implying that the user owes the exchange (and hence, other traders) capital. We say that a trader’s position is *insolvent* if such a negative net position is realized.

Exchanges utilize liquidations to remove trader positions before they are insolvent. For instance, in our example, an exchange might deem a trader’s position to be liquidatable if the price falls by 9%, such that the trader’s position is closed before insolvency. Provided that the exchange can execute this position (inclusive of transaction fees and market impact costs) at a price in between $0.91p_0$ and $0.9p_0$, the position can be closed profitably. This profit is usually distributed to other traders in the exchange, added to an insurance fund for future losses, or realized as a profit for the exchange.

Centralized exchanges usually execute liquidations in their own spot markets. This is because they can effectively guarantee atomicity of liquidations, ensuring that the liquidation transaction does not get front run and have a worse execution price than expected. In decentralized exchanges, liquidations are usually executed by third-party actors known as liquidators. Liquidators can be viewed as traders who have the ability to warehouse the inventory of a collateral position and exit profitably. From an economic perspective, liquidations need to be profitable for either the exchange or the liquidator in order to help with exchange solvency.

Autodeleveraging. In severe market dislocations, it is sometimes possible for liquidations to be so deeply unprofitable that they are unable to be executed by any party. When this happens, the exchange can reach a state of insolvency, where the assets held are less than the liabilities owed to its traders. This excess in liabilities is known as the *shortfall*. In such scenarios, a natural strategy to reduce the risk of total insolvency is to socialize losses by haircutting profitable (or *winning*) traders. In the worst case, a winning user’s position is completely closed to zero, leaving them with a potentially large opportunity cost. This process is known as *autodeleveraging* (ADL) and is a last-resort measure used by exchanges.

Autodeleveraging is unique in that it algorithmically socializes losses on winners based on a public rule or mechanism that an exchange posts. This means that all of the position closures and implicit liquidations of insolvent positions are performed in an irreversible and atomic manner. The closest analogue to autodeleveraging in traditional finance is a central counterparty (CCP) such as a derivatives clearing house. In traditional finance, CCPs employ default waterfalls and loss mutualization (including variation-margin gains haircutting (VMGH)) that share the same core idea of allocating residual losses across non-defaulting members [43, 44, 48, 74, 93, 112]. CCPs generally have manual, non-algorithmic socialization and disbursement of insurance funds during such events with disputes reconciled via the legal system. We note, furthermore, that only clearing members bare the losses in CCPs, which leads to a different principal-agent dynamic than perpetuals markets as smaller traders need not bare the final losses.

Autodeleveraging, to the best of the author’s knowledge, was first introduced by the crypto exchange Huobi in 2015 [105]. Prior to the implementation of ADL, exchanges would manually socialize losses, with the exchange serving as the sole CCP. By 2016, the BitMEX exchange implemented ADL via a formula that ranked trader positions for ADL by their profit times their leverage [24]. This formula persists to today on virtually every centralized exchange (including Binance, who implemented the formula in 2019 [20]) and on the decentralized exchange Hyperliquid [68]. We note that there are some decentralized exchanges such as Drift [46] and Paradex [88] that use a different ADL mechanism (the *pro-rata* mechanism defined in §2.4). However, well over 95% of perpetuals trading volume is on exchanges that utilize the BitMEX ranking model.

In practice, ADL is reserved for extreme tail events where routine liquidation mechanisms are insufficient to maintain venue solvency. There are two primary scenarios that can trigger ADL activation. First, a sudden and large price movement can push positions far beyond their liquidation thresholds before the exchange’s liquidation system can execute the necessary trades. Second, a temporary but severe loss of executable liquidity in the order book can prevent the exchange from closing positions at prices near their liquidation triggers, even when the price movement itself is not extreme.

To illustrate how such scenarios can lead to insolvency, consider a concrete example. Suppose that a trader posts \$1,000 of collateral to take a 10x leveraged long position with a notional value of \$10,000, where the position becomes liquidatable when the price declines by 9% from its initial value p_0 . Under normal market conditions, the exchange would execute a liquidation when the price reaches $0.91 p_0$, closing the position and recovering the collateral. However, if an abrupt 12% price drop occurs between price updates or during a period of low liquidity, the position may be closed at a significantly worse execution price, such as $0.88 p_0$ rather than the intended $0.91 p_0$. In this scenario, the account’s equity becomes negative, representing approximately 2% of its initial collateral value. When many traders simultaneously experience such execution slippage and the exchange’s insurance fund reserves cannot absorb the aggregate resulting losses, the venue triggers ADL. Under the exchange’s predefined ADL rules, the system algorithmically reduces or closes profitable opposing positions until the shortfall is covered and solvency is restored.

Inherently, ADL creates *moral hazard* — that is, losing traders have an incentive to take on more risk in anticipation of being socialized by winning traders. The goal of an exchange is to balance the risk of insolvency against haircutting winning trader profits. Moral hazard is measured by this trade-off — if it is possible to reduce the risk of insolvency to zero without haircutting winning traders, then the system has no moral hazard. On the other hand, if it is not possible to reduce the risk of insolvency by any appreciable amount without severely haircutting winning trades, the system has excessive moral hazard. The job of an exchange operator is to choose an ADL mechanism that tries to induce a low level of moral hazard under most tail event scenarios. This connects to principal-agent models of moral hazard and linear/robust contracting [31, 49, 67], with

key differences: many agents act simultaneously, externalities operate through a common solvency constraint, and tail-risk mitigation substitutes for standard effort observability.

Despite the rarity of ADL, it is certainly not hypothetical: there have been dozens of instances of ADL activations on major venues since 2018 [3, 20, 24, 38]. The largest episode of ADL usage was on October 10-11, 2025 (UTC), when outsized liquidations exhausted insurance funds and multiple venues simultaneously invoked ADL [96]. However, other famous incidents such as the March 2020 “Black Thursday” and May 19, 2021 deleveraging event were other prominent days involving ADL usage. Most of these usages of ADL often involved controversies, in that highly profitable but lower risk (*e.g.* lower leverage) traders found their positions closed out by ADL mechanisms. These traders have often threatened to sue exchanges for unfair treatment, with rumors and public speculation surfacing during Oct. 10–11, 2025 (subsequently denied by Wintermute) [40, 111]¹.

This incidents demonstrate a natural trade-off that perpetuals exchanges have to make: they can either aggressively socialize losses to their winners and (potentially) lose the future revenue of these users while preserving solvency or hold the losses due to the insolvency on their balance sheet. In the former case, the exchange effectively chooses to prioritize short term solvency and potential liquidation profits over long term value and fees generated by users. The latter case, effectively creates a long term risk for the exchange’s balance sheet, especially if it does not have sufficient funds to cover insolvencies. This strongly suggests that one can view the choice of how an exchange implements ADL as a trade-off between solvency, moral hazard, and long term revenue for an exchange.

We note that open interest post October 10, 2025 that suggests that Hyperliquid has lost nearly 50% of open interest, whereas competitors such as Lighter and even Binance have recovered to pre-event levels.² Open interest is generally more expensive to manipulate than trading volume and can serve as a coarse measurement of expected future revenue for an exchange. Numerous parties have argued that this open interest loss for Hyperliquid is due to their overly aggressive ADL policy.³

This Paper. In this paper, we provide (to the author’s best knowledge) the first formal mathematical model of autodeleveraging. The main model of the paper is venue-agnostic: it abstracts away pricing (*e.g.* order book versus automated market maker microstructure) and focuses on how positions are opened and closed. As a warning to the reader: this paper is more verbose than necessary. The author aims to provide a self-contained, pedagogical introduction to perpetual futures and liquidations, as the current literature and open source documentary does not cover liquidations and ADL in any manner comprehensive enough for formal study. Throughout the paper, there will be numerous simple numerical examples to illustrate the basic concepts and mechanics of perpetual futures.

The first section of the paper (§2) provides definitions and examples of the basic objects — *positions* — and basic operations that one can perform on them. In the process of defining the set of operations on positions, we will naturally need to define the economics of positions using funding rates and arbitrage. These definitions will hopefully provide answers to the questions of why one opens a position, how much one has to spend on keeping a position open, and how positions on different exchanges achieve price synchronization. From these economic properties, we will naturally

¹See also prior complaints during May 2020 on Binance futures when profitable shorts were closed via ADL [39]

²Source: DeFiLlama perpetuals dashboard, <https://defillama.com/perps>, accessed Nov. 30, 2025.

³See, *e.g.* , public commentary from @fiddybps1, @0xReaz, @Eugene_Bulltime, and @0xLouisT on X (Nov. 2025): <https://x.com/fiddybps1/status/1978750722901762321>, <https://x.com/0xReaz/status/1986486213101166599>, https://x.com/Eugene_Bulltime/status/1994380900067582182, <https://x.com/0xLouisT/status/1990815109787058654>.

define solvency for both individual traders and of the exchange itself.

Liquidations and ADL then naturally follow as operations that exist to try to use economic means to enforce solvency constraints with high probability. We will define the set of ADL mechanisms in a sufficiently broad manner to allow for venue-specific ADL, this allowing us to compare Binance, Hyperliquid, Drift, and other exchanges. Our decomposition of the ADL strategy space into *severity* (*i.e.*, dollar amount impacted by ADL) and *haircuts* (*i.e.*, per-trader percentage liquidated) allows us to separate exchange solvency and trader profitability into two separate optimization problems. This will be key to simplifying our results and constructing practical algorithms.

Given a perpetuals exchange, defined in the combinatorial manner of §2, we next focus on notions of risk that are endemic to perpetuals exchanges in §3. We construct a set of risk measurements that, while specific to perpetuals exchanges, are connected to classical risk metrics such as Value-at-Risk. Subsequently, we describe four facets of perpetuals exchange risk that we will analyze: extreme value analysis, Schur ordering, fairness, and moral hazard.

Extreme value analysis is the classical actuarial estimation of tail probabilities and rare events. We provide a simple primer to make the text self-contained and our notation consistent. The Schur ordering provides a way of partially ordering risk measurements, even if they are not fully comparable. This allows us to compare ADL mechanisms across venues, despite the fact that the outcomes might not have explicitly comparable risks in every situation. Algorithmic fairness measures (axiomatic and optimization-based) are subsequently introduced to concretely answer the question of whether ADL is more fair to certain traders than others. Finally, we describe the general moral hazard problem and specialize it to the perpetuals exchange setting. In this setting, it corresponds to whether the exchange operator has an incentive to take on more risk than necessary because they can socialize worst-case losses on their best traders.

With these risk tools in tow, we are now prepared to describe the main results of the paper. We have five classes of results in this paper:

1. *Negative*: Moral hazard is strictly increasing in the size of the exchange.
2. *Fairness*: The Pro-Rata ADL mechanism is the most fair mechanism
3. *Robustness*: Given a risk model g for historical user behavior, there is a unique ADL mechanism maximally robust to price shocks
4. *ADL as Stackelberg Game*: Multiple price shocks causing ADL to be used repeatedly are a Stackelberg game that can be studied rigorously.
5. *Empirical*: Using Hyperliquid’s data from October 10, 2025, where ADL was used over 40 times in a 10 minute period, we evaluate the performance of difference ADL mechanisms

Negative Results. Our negative results focus on quantifying limits to how much ADL can actually reduce insolvency without inuring traders with disproportionate losses. We explicitly ask the question of how the loss of profit due to ADL to the trader with the highest PNL scales as the number of positions n grows. We quantify this by looking at ratios of the maximum winner’s profit to the total shortfall (over all traders) and to the worst trader’s shortfall. These ratios can be viewed as analogues of the Value-at-Risk and Expected Shortfall metrics in traditional risk modeling, as we show in Appendix C. The main difference is that while Value-at-Risk and Expected Shortfall measure the total expected loss (over all traders) in tail events, our metrics only measure the expected loss to the most winning trader. Moreover, Appendix C.5 shows that classical VaR/ES rankings echo the same conclusion: queue-based ADL always delivers the lowest top-winner survival even under these traditional risk lenses.

When these ratios, as constructed, are bounded below by a constant, then the winner can expect to always retain at least a constant fraction of their earnings after an ADL event. However, if these ratios decay to zero, then the winner has no guarantee of retaining any of their profits after an ADL event. We show that under some mild technical assumptions, these ratios decay as $O(\frac{b_n}{n})$, where b_n is the *extreme value scale* (EV scale) of profits (see §3). For most well-behaved probability distributions over the set of profits, $b_n = o(n)$, so the ratios decay to zero. This result establishes a fundamental limit of these markets: moral hazard scales with participation. It suggests that as crypto exchanges grow, ADL must become either more frequent or more severe for top traders; they cannot “grow out” of the problem.

Fairness Results. As the negative results show that no ADL mechanism can scale to preserve winner’s profit, a natural follow-up question is, what is the most fair way of socializing profitable users? We consider two notions of fairness in §6: axiomatic and optimization-based. They are both common in the algorithmic fairness literature within machine learning [4, 16, 50, 63, 73, 82]. We demonstrate that the pro-rata ADL rule (used by Drift [46] and Paradex [88]) is the unique fair ADL rule in *both* the axiomatic and optimization senses [71].

Axiomatic fairness involves defining rules or invariants that a particular ADL mechanism must satisfy to be deemed fair. The rules we analyze for fairness are monotonicity, scale invariance and Sybil resistance. Roughly speaking, monotonicity says that if you were the k th highest winner (ranked by PNL) before ADL, then you will be the k th highest winner after ADL. Scale invariance states that if everyone scaled up their positions by an equal factor, their loss due to ADL does not change. And finally, Sybil resistance states that splitting a position into multiple positions does not change the loss due to ADL.

We show that under mild smoothness conditions, the pro-rata ADL rule is the unique rule to satisfy these axioms. In stark contrast, we show that the queue-based mechanisms used by Binance and Hyperliquid are effectively “anti-fair”: in our model and data they produce extremely low Profit-to-Solvency Ratios (PTSR), imposing disproportionately large losses on the top winner. This finding contradicts the common intuition that queues are orderly; instead, they concentrate losses on the most successful traders.

Optimization-based fairness involves showing that for a large class of individual trader utility functions, a particular mechanism maximizes social welfare (*i.e.*, the total profit). This is more of a system-wide fairness notion, as it shows that the system as a whole is fair, even though there might be individuals who were punished more by an ADL mechanism. Unlike the axiomatic rules, which have to hold for every trader exactly, the optimization-based fairness rules can be thought of as ‘fuzzier’ and allowing for more slack. Despite this, we still show that the pro-rata ADL rule is the unique rule to maximize social welfare for a large class of concave utility functions.

Robustness Results. A natural question, beyond fairness, is whether an ADL rule is *robust* to subsequent price shocks. The main question here is whether the choices of severities (*e.g.* total amount to socialize) and haircuts (*e.g.* per-account socialization) increases the chance of insolvency should another price shock occur. In §7, we consider a new threat model where an ADL mechanism is graded based on how well it reduces to total shortfall from two price shocks. The model involves a price shock occurring that triggers ADL, an ADL policy being executed, and then another price shock occurring.

We show that a natural way to model total shortfall is via a *risk model* g that maps a user’s amount of leverage used to a risk score. A higher risk score implies that a trader’s leverage level places a higher burden on the rest of the exchange’s solvency, especially in the repeated price shock

scenario. We first show that the risk model is effectively a way to quantify the total shortfall. Given a stochastic process for price shocks, one can construct an optimal risk model (using convex duality) g^* that maximizes the total solvency gained per dollar of winner position that is haircut by an ADL policy. With the optimal risk model in hand, we construct a linear-time algorithm to choose haircuts for each trader that are weighted by this model. We then show that the resulting ADL policy derived from these weights has lowest expected total shortfall and is the optimal ADL policy (in the Schur ordering sense).

ADL as a Stackelberg Game. A direct extension of the robustness threat model is to consider a sequential ADL problem, where for multiple rounds, a price shock is received, the exchange chooses an ADL policy, and then losses are realized. This model might at first glance seem to only have theoretical value since ADL is itself a rare event. However, on October 10, 2025 there were multiple ADL events that occurred over a short time window on both Hyperliquid and Binance. While the account-level data for Binance’s usage of ADL is not public, we do have this data for Hyperliquid. Starting at 10:21:00 UTC, there were at least 40 recorded ADL events on Hyperliquid in a 10 minute period. This strongly suggests that there is clustering or cascading of ADL events, making the multiple round model necessary to study.

Unlike in the single round models (used for fairness and robustness), the exchange is allowed to adjust their ADL policy after each ADL shock. One can view this as a sequential game and we argue that it is a *Stackelberg game*. Stackelberg games have a well-known follower-leader decomposition, where one user (the leader) plays a strategy first. Then the follower conditions their strategy on the leader’s action, before playing their own move. In the ADL setting, the exchange is the leader and the traders are the followers.

We first show that viewing ADL as a dynamic, multiple round Stackelberg game leads to sharply different equilibria. We prove a crucial “principal-agent” tension: any strategy an exchange executes that minimizes the time to recover solvency (*i.e.*, number of rounds) necessarily maximizes the loss of exchange long-term revenue. This trade-off is the economic heart of the paper: the exchange wants to be solvent in few rounds, but doing so destroys the best customers (*e.g.* the whales). Intuitively, this occurs because any strategy that quickly recovers solvency autodeleverages the largest winners more aggressively. These winners, however, are also often the most profitable users and generate high long-term fees for the exchange. By liquidating them aggressively, you’re removing the capital that they would use to later create positions and generate fees. We note that threats to sue exchanges for ADL losses in profit after October 10, 2025 provide indirect empirical evidence to support this theoretical result [40, 111].

This implies that any optimal ADL policy over multiple rounds also has to include incentive compatibility constraints that balance future exchange revenue, current exchange solvency, and current trader profits. We construct an explicit algorithm for optimizing a revenue and solvency based objective and show that a simple mirror descent algorithm achieves vanishing regret for this objective. We conclude with a simple example demonstrating that no static policy that dominates the dynamic policy, especially when exchange long-term revenue is a consideration.

Empirical Results: The \$653 Million Inefficiency. We conclude by analyzing the October 10, 2025 ADL events on Hyperliquid. Unlike centralized exchanges like Binance and Bybit, Hyperliquid publicly shares all position data, including position leverage, collateral, PNL, and ADL status. Using this data, we evaluate the efficiency of Hyperliquid’s ADL mechanism (which is a slight modification of the original 2015 Huobi queueing mechanism). We show that Hyperliquid autodeleveraged roughly $\sim \$660\text{m}$ (simulated) to $\sim \$2.1\text{bn}$ (realized) of winning trader PNL during the 12 minute cascade on

October 10, 2025.

Over the same period, there were at least 40 ADL events which accumulated a total of $\sim \$23.2\text{m}$ of true negative equity positions.⁴ Measured shock-by-shock, this implies the production queue generated about $\sim \$653\text{M}$ of overshoot, *i.e.*, haircuts beyond the deficits that actually materialized. This result highlights a massive capital inefficiency in current queue-based mechanisms: they are not just unfair, but they destroy hundreds of millions of dollars of trader value unnecessarily.

We then compare the pro-rata, risk-weighted pro-rata, and dynamic mechanisms to the live queueing mechanism used. We utilize these methods with different severity optimization algorithms and show that all of these algorithms were able to learn to autodeleverage substantially less than what Hyperliquid did in practice. Utilizing these algorithms on the historical data, we find that nearly 98% of the over autodeleveraging used by Hyperliquid could have been removed by a smarter ADL policy.

Notation.

- For a vector $x \in \mathbf{R}^n$, we define $(x)_+ = \max(x, 0)$, where this is done coordinate-wise and $(x)_- = \min(x, 0)$ (also coordinate-wise). Using our convention, we have $(x)_+ = -(-x)_-$
- We denote the set of integers $\{1, \dots, n\}$ by $[n]$.
- For a sequence $x_1, x_2, \dots$, we denote the subsequence from s to t as $x_{s:t} = (x_s, x_{s+1}, \dots, x_t)$.
- We use $\mathbf{1}_E$ or $\mathbf{1}\{E\}$ to denote the indicator function of an event E , which is 1 if E is true and 0 otherwise.
- We use the notation $f(n) \asymp g(n)$ to denote that $f(n) = \Theta(g(n))$, *i.e.*, there exist constants $c, C > 0$ such that $c|g(n)| \leq |f(n)| \leq C|g(n)|$ for sufficiently large n .
- We write $X_n \asymp_p Y_n$ for random variables X_n, Y_n if there exist constants $c, C > 0$ such that $c|Y_n| \leq |X_n| \leq C|Y_n|$ with high probability as $n \rightarrow \infty$.
- For vectors $x, y \in \mathbf{R}^n$, we write $x \prec y$ if x is majorized by y (see §3.2 for details).

2 Background

We will first introduce a mathematical model for solvency risks in a perpetuals exchange. Using this model, we define the key components for measuring the solvency of perpetuals exchanges: trader equity, exchange deficits, and autodeleveraging (ADL) policies. These components will allow us to formalize the notion of moral hazard in perpetuals exchanges and sets the notation up for our main results in §5, §6, §7, and §8.

2.1 Perpetuals Exchanges.

A *perpetuals futures exchange* (or perpetuals exchange) is a derivatives trading venue that allows traders to use leverage to speculate on the prices of cryptocurrency assets. The main asset traded is a perpetual future, which is a continuous, expiryless futures contract. Traders speculate on the contract in either the long (buy) or short (sell) direction. Traders are allowed to make leveraged

⁴Hyperliquid’s on-chain fill log for October 10, 2025 records 34,983 individual ADL executions across 19,337 distinct wallets and 162 tickers. We analyze the per-fill data but aggregate those executions into the 40 queue waves discussed in the text.

bets up to a maximum leverage cap $\ell^{\max}$. As is often common with levered positions, the trader has to post collateral, known as *margin*, to keep their positions *solvent*. We define a precise notion of solvency in this section, but for now think of solvency conditions as meaning that a trader's assets are greater than their liabilities.

The exchange incentivizes futures prices to approximately match spot prices by having long and short traders pay each other a continuous rate, known as the *funding rate*. The funding rate is determined by the spot price of the asset and the relative imbalance of long and short futures positions. Arbitrageurs keep the price of the futures contract in line with the spot price by opening futures positions that push the contract price towards the spot price. For simplicity, we assume that trades occur in discrete time $t \in \mathbf{N}$, but note that our results can easily be extended to continuous time $t \in \mathbf{R}_+$.

Spot Price Oracle. Perpetual futures exchanges require a spot price oracle that provides a spot price of the asset from a spot trading venue. Price oracles are assumed to be expensive or difficult to manipulate, which is dependent on the liquidity of the spot market. Most perpetuals exchanges utilize oracles that aggregate price reporting over multiple venues to increase the overall liquidity represented in the price and hence, the manipulation cost. We assume that there is a non-manipulable spot price $\hat{p}_t$ for an asset at all times $t \in \mathbf{R}_+$. Furthermore, we denote the futures price quoted by the exchange as p_t and similarly assume that it exists for all $t \in \mathbf{R}_+$.

Positions. Given a spot price oracle and futures price, a perpetuals exchange is able to open and close *positions* created by its traders. A position consists of collateral deposited by the trader, a leverage amount, and a direction (*e.g.* buy or sell). A trader's position is updated as the spot price changes, realizing a gain or a loss for the trader. The accumulation of gains and losses over the lifetime of the position is the trader's PNL (profit-and-loss). We define a perpetuals exchange with n traders as the set

$$\mathcal{P}_n = \{(q_i, c_i, t_i, b_i) \in \mathbf{R}_+^3 \times \{-1, 1\} : i \in [n]\} \quad (1)$$

where $q_i \geq 0$ is the notional quantity of futures held by the trader, $c_i \geq 0$ is the trader's collateral (cash) position (also known as the *initial margin*), $t_i \in \mathbf{N}$ is the timestamp that the position was created, and b_i is their side (-1 is sell, 1 is buy). We denote each position $\mathbf{p}_i = (q_i, c_i, t_i, b_i)$ and assume that without the loss of generality that there is a fixed, strictly increasing ordering of these orders such that we can refer to $\mathbf{p}_i$ unambiguously.

For each position $\mathbf{p}$, we define the *notional exposure* at time t is the gross dollar size

$$n_{i,t} = p_t q_i.$$

The notional exposure is the gross dollar size of the trader's position and is used to define leverage.

Example. In this section, we fix a canonical running example that we will reuse across definitions to illustrate each exchange function. Our simple test perpetuals exchange will have five positions, $\mathcal{P}_5$:

$$\begin{aligned} \mathcal{P}_5 = \{ & \mathbf{p}_A = (q_A, c_A, t_A, b_A) = (1, 2, 0, +1), \\ & \mathbf{p}_B = (q_B, c_B, t_B, b_B) = (1, \frac{2}{3}, 0, +1), \\ & \mathbf{p}_C = (q_C, c_C, t_C, b_C) = (4, \frac{8}{3}, 0, -1), \\ & \mathbf{p}_D = (q_D, c_D, t_D, b_D) = (1, \frac{2}{19}, 0, +1), \\ & \mathbf{p}_E = (q_E, c_E, t_E, b_E) = (1, \frac{10}{99}, 0, -1) \}. \end{aligned}$$

Moreover, if at $t = 0$ we have $p_0 = 1$ for $\mathcal{P}_5$, then the notional exposure are:

$$n_{A,0} = 1 \cdot 1 = 1, \quad n_{B,0} = 1 \cdot 1 = 1, \quad n_{C,0} = 1 \cdot 4 = 4, \quad n_{D,0} = 1 \cdot 1 = 1, \quad n_{E,0} = 1 \cdot 1 = 1.$$

Note this toy example is net short by 2 contracts: $\sum_i b_i q_i = -2$. For AMM-style venues, the pool holds the offsetting inventory $q_V = 2$ with $b_V = +1$, so across traders plus venue $\sum_i b_i q_i + b_V q_V = 0$. For order-book venues, interpret $\mathcal{P}_5$ as a subset of accounts. The complementary matched positions are omitted but implied.⁵ All zero-sum statements (*e.g.*, funding rates, which will be defined shortly) are taken over traders plus the venue inventory.

Position Creation. When an exchange opens a new position, it performs two tasks:

1. Selecting a price p_t for the position and constructing $\mathbf{p} = (q, c, t, b)$.
2. Construction of an equal but opposite position, $\bar{\mathbf{p}} = (q, c, t, -b)$ is created

The latter condition is needed because the perpetuals exchange is meant to be a neutral matching layer for traders. This means that the exchange should not have a net exposure to the spot asset, as it increases the risk of default for traders. Specifically, the position $\bar{\mathbf{p}}$ is created to ensure that the total notional exposure of the exchange to the spot asset is zero.⁶ We define

$$\bar{\mathcal{P}}_n = \mathcal{P}_n \cup \{\bar{\mathbf{p}} : \mathbf{p} \in \mathcal{P}_n\}$$

as the full set of positions held by the exchange (including matching, complementary positions).

While all venues construct $\bar{\mathbf{p}}$ in the same manner, the price p_t is chosen in a venue-specific manner. Each venue has a *liquidity model* for determining how prices p_t are chosen. The two most popular liquidity models are order books and automated market makers (AMMs).

Order book venues maintain a limit order book of bids and asks for perpetual futures contracts at discrete price levels, and have been extensively studied in the market microstructure literature [7, 26, 57, 76]. Traders may submit limit orders, which rest in the book until they are matched, or market orders, which execute immediately against resting liquidity. When a trader opens a new position as a taker, the execution price p_t is the volume-weighted average of the resting orders they consume, *i.e.*, the price they would obtain by submitting a market order of size q into the book at time t . The largest order book venues in the perpetual futures market include centralized exchanges such as Binance and Bybit and decentralized exchanges such as Hyperliquid and Lighter.⁷

Automated market maker (AMM) venues maintain a pool of collateral and inventory for the perpetual, and quote prices according to a deterministic pricing rule or bonding curve that depends on the pool state, typically implemented via constant-function market makers [8, 10]. Traders interact with the pool via swap-like trades: submitting an order of size q moves the pool along its pricing curve, and the execution price p_t is the average price paid over this path. When a trader opens a new position against the AMM, the pool itself holds the offsetting inventory $\bar{\mathbf{p}}$, so that the

⁵We note that in centralized exchanges, the exchange itself creates the complementary positions whereas in decentralized exchanges with order books (like Hyperliquid and Lighter), a liquidity provider like HLP or LLP owns the position; see Hyperliquid’s HLP documentation [69] and the Lighter whitepaper [79].

⁶In centralized venues such as Binance Futures, the matching engine and clearing layer are designed to operate a matched book over user positions, as reflected in their open interest endpoints for USD\$- and COIN-margined futures [23]. In decentralized order-book venues like Hyperliquid, the protocol core similarly avoids taking directional exposure. Offsetting inventory is held in separate HLP vaults on behalf of liquidity providers [69].

⁷From a risk-warehousing perspective, Hyperliquid’s HLP vaults cause the protocol to function as a hybrid between a pure order-book venue and an inventory-taking pool: user orders are matched via an order book, but directional exposure is warehoused collectively in HLP vaults on behalf of liquidity providers [69, 86, 96].

trader's exposure is matched by an equal and opposite position in the pool, as in inventory-taking perpetual venues such as GMX, Drift, and Perpetual Protocol v2 [46, 59, 91].

In the remainder of this paper, we will ignore the liquidity model of a perpetuals exchange. This is because the default conditions that we will study for autodeleveraging only depend on the positions $\mathcal{P}_n$ and not the process by which they were created.

Example. Consider an order-book venue at time t with three resting sell (ask) limit orders

$$(p^1, q^1) = (1.0, 1), \quad (p^2, q^2) = (1.1, 3), \quad (p^3, q^3) = (1.2, 10).$$

A trader submits a market buy order for quantity $q = 5$. The exchange matches this order against the resting asks in price-time priority: it first consumes $q_1 = 1$ at $p_1 = 1.0$, then $q_2 = 3$ at $p_2 = 1.1$, and finally 1 unit from the third order at $p_3 = 1.2$. The total notional paid is $1 \cdot 1.0 + 3 \cdot 1.1 + 1 \cdot 1.2 = 5.5$, so the execution price is the volume-weighted average $p_t = \frac{5.5}{5} = 1.1$. Ignoring fees, the taker acquires a new long position $\mathbf{p} = (q, c, t, b) = (5, c, t, +1)$ with notional exposure $q p_t = 5.5$. Aggregating across the three maker accounts, the exchange simultaneously creates an equal and opposite short position $\bar{\mathbf{p}}$ of size 5, so that total long and short contract quantities (over all accounts) remain matched.

Open Interest. Note that the notion of leverage is encoded into the quantity q_i of a position, as a user who can buy $p_t q$ units of spot assets with c_i units of collateral can open a position with $q_i \leq \ell_{\max} q$. This implies that the correct notion of the notional trades outstanding, known as *interest*, on a perpetuals exchange is simply the sum of price-weighted quantities. We divide the exchange $\mathcal{P}_n$ into long and short positions:

$$\begin{aligned} \mathcal{L}(\mathcal{P}_n) &= \{(q_i, c_i, t_i, b_i) \in \mathcal{P}_n : b_i = 1\} \\ \mathcal{S}(\mathcal{P}_n) &= \{(q_i, c_i, t_i, b_i) \in \mathcal{P}_n : b_i = -1\} \end{aligned}$$

We define the total set of long and short positions as:

$$\begin{aligned} \bar{\mathcal{L}}(\mathcal{P}_n) &= \mathcal{L}(\mathcal{P}_n) \cup \{\bar{\mathbf{p}} : \mathbf{p} \in \mathcal{L}(\mathcal{P}_n)\} \\ \bar{\mathcal{S}}(\mathcal{P}_n) &= \mathcal{S}(\mathcal{P}_n) \cup \{\bar{\mathbf{p}} : \mathbf{p} \in \mathcal{S}(\mathcal{P}_n)\} \end{aligned}$$

This is the set of all long and short positions, including matching positions.

Futures exchanges define the *open interest* of the exchange as the total outstanding notional value of the positions held by the exchange. Given the futures price p_t , the long open interest $L(\mathcal{P}_n, p_t)$ and short open interest $S(\mathcal{P}_n, p_t)$ are defined as

$$\begin{aligned} L(\mathcal{P}_n, p_t) &= \sum_{(q_i, c_i, t_i, b_i) \in \mathcal{L}(\mathcal{P}_n)} q_i p_t \\ S(\mathcal{P}_n, p_t) &= \sum_{(q_i, c_i, t_i, b_i) \in \mathcal{S}(\mathcal{P}_n)} q_i p_t \end{aligned}$$

Given the long and short open interest, we define the exchange's open interest simply as the sum, $OI(\mathcal{P}_n, p_t) = L(\mathcal{P}_n, p_t) + S(\mathcal{P}_n, p_t)$ ⁸ To lighten notation, we will write $L_t = L(\mathcal{P}_n, p_t)$, $S_t = S(\mathcal{P}_n, p_t)$, $OI_t = OI(\mathcal{P}_n, p_t)$ when the exchange $\mathcal{P}_n$ is already defined.

⁸We note that exchange's will report open interest over $\bar{\mathcal{P}}_n$, which would be twice the open interest reported here. Since we assume that $\mathcal{P}_n$ is the set of unmatched trader positions, one can view our definition of open interest (which differs from Binance's definition by a factor of two) as 'trader open interest' vs. 'exchange open interest'.

Example. For $\mathcal{P}_5$ at $t = 0$ with $p_0 = 1$, longs are $\{A, B, D\}$ and shorts are $\{C, E\}$, so

$$L_0 = \sum_{b_i=+1} q_i p_0 = (1 + 1 + 1) \cdot 1 = 3, \quad S_0 = \sum_{b_i=-1} q_i p_0 = (4 + 1) \cdot 1 = 5, \quad OI_0 = 8.$$

At $t = 1$ with $p_1 = 1.4$ (same quantities), $L_1 = 3 \cdot 1.4 = 4.2$, $S_1 = 5 \cdot 1.4 = 7.0$, $OI_1 = 11.2$.

Leverage and Margin. The *leverage* of a position is the ratio of the notional exposure to the collateral value. If the leverage is greater than 1, it means that the user has effectively borrowed money from the exchange to have an exposure higher than their cash balance. When the leverage is less than 1, the user's position is *overcollateralized*.

As the price of the asset changes, the leverage of the position naturally changes due to its dependence on notional exposure. When a position is opened, there is an *initial leverage* ratio that the position is created with. We define the initial leverage of a position as

$$\ell_i = \ell(\mathbf{p}_i) = \frac{n_{i,t_i}}{c} = \frac{p_{t_i} q_i}{c_i}$$

The larger the initial leverage, the most exposure the trader has to the asset and implicitly, the more risk the exchange is taking on lending to the user.

For risk management purposes, exchanges define leverage limits in terms of initial and maintenance *margins*. Margins are bounds on the ratio of notional exposure to collateral that the exchange enforces as an invariant. The *initial margin ratio*, $m_I \in (0, 1)$ is defined such that for any valid position, we have:

$$m_I p_{t_i} q_i \leq c_i \quad (2)$$

This invariant implies that the maximum leverage that a position can have, $\ell^{\max}$ is $\ell^{\max} = \frac{1}{m_I}$.

In order to continually satisfy (2) as spot and future prices change, a trader often needs to 'top up' or add more collateral their position. The amount of collateral needed is related to the exchange's *maintenance margin ratio* m_μ , which will be defined in the sequel. As the user adjusts their collateral position dynamically, we let $c_{i,t}$ be the total quantity of collateral (margin) placed by the user at time t . This also means that the state of position is dynamic, $\mathbf{p}_{i,t} = (q_i, c_{i,t}, t_i, b_i)$. We define the *leverage at time t* as $\ell_{i,t} = \mathbf{1}_{t \geq t_i} \frac{p_t q_i}{c_{i,t}}$

Example. These five examples of $\mathcal{P}_5$ represent one highly overcollateralized and under leveraged position ($\mathbf{p}_A$), two somewhat leveraged positions with medium leverage ($\mathbf{p}_B, \mathbf{p}_C$) and two highly leveraged positions ($\mathbf{p}_D, \mathbf{p}_E$). We will initialize the example with $p_0 = 1$ and take a maximum leverage $\ell^{\max} = 10$ (e.g., $m_I = 0.10$). This gives the following opening leverages:

$$\begin{aligned} \ell_{A,0} &= \frac{1 \cdot 1}{2} = 0.5, & \ell_{B,0} &= \frac{1 \cdot 1}{2/3} = 1.5, & \ell_{C,0} &= \frac{1 \cdot 4}{8/3} = 1.5, \\ \ell_{D,0} &= \frac{1 \cdot 1}{2/19} = 9.5, & \ell_{E,0} &= \frac{1 \cdot 1}{10/99} = 9.9. \end{aligned}$$

If the mark moves to $p_1 = 1.4$ and cash is unchanged ($c_{i,1} = c_i$), then time-1 leverages are

$$\begin{aligned} \ell_{A,1} &= \frac{1.4 \cdot 1}{2} = 0.7, & \ell_{B,1} &= \frac{1.4 \cdot 1}{2/3} = 2.1, & \ell_{C,1} &= \frac{1.4 \cdot 4}{8/3} = 2.1, \\ \ell_{D,1} &= \frac{1.4 \cdot 1}{2/19} = 13.3, & \ell_{E,1} &= \frac{1.4 \cdot 1}{10/99} = 13.86. \end{aligned}$$

In practice, if $\ell_{i,1} > \ell^{\max}$ (e.g., $\mathbf{p}_D$ and $\mathbf{p}_E$), then the venue requires collateral additions from the trader to restore the constraint $\ell_{i,1} \leq \ell^{\max}$.

Funding Rate. A perpetual future is only useful as a hedging instrument if the futures price p_t and spot price $\hat{p}_t$ are sufficiently “close.” To incentivize this, the majority of perpetuals futures markets use a continuous payment stream between the long and short positions called a *funding rate*. If the price of the future is lower than spot, $p_t < \hat{p}_t$, traders holding short futures positions pay a payment to traders holding long futures positions. This encourages traders who want to earn these payments to open long positions, which will decrease the over all rate. Conversely, if $p_t > \hat{p}_t$, traders holding long positions pay those holding short positions.

Most funding rates are simple linear functions of the relative imbalance between the long and short positions. A common funding rate [33] formula takes the form:

$$\gamma(\mathcal{P}_n, p_t, \hat{p}_t) = \kappa \left(\frac{L(\mathcal{P}_n, p_t)}{S(\mathcal{P}_n, p_t)} - \frac{p_t}{\hat{p}_t} \right) \quad (3)$$

for a constant $\kappa > 0$. For brevity, write $\gamma_t = \gamma(\mathcal{P}_n, p_t, \hat{p}_t)$. Note that we have defined γ_t such that if $\gamma_t > 0$, then short traders pay long traders $\gamma_t\%$ of their holdings. This convention is opposite to what many exchanges use (*i.e.*, $\gamma_t < 0$ represents a payment from the short side to the long side). However, our convention leads to needing to check fewer negative signs in our formulas, as we will illustrate with our examples.

Given funding rates γ_t for $t \geq 0$, a position $\mathbf{p} = (q, c, t, b)$ linearly accumulates funding positions between time t and the current time T , if the position is solvent. That is, if the position is solvent at time T , the user receives a payments of the form

$$\Gamma((q, c, t, b), t, T) = \sum_{s=t+1}^T (bq)\gamma_s p_s$$

If $\Gamma > 0$, then over the interval $[t, T]$, the user receives a payment from the exchange whereas if $\Gamma < 0$, the user pays the exchange. When the user pays the exchange, their profits are distributed by the exchange to the other side of the market. For brevity, we use the notation $\Gamma_{i,t,T} = \Gamma(\mathbf{p}_i, t, T)$.

We also note that the funding rate, by construction, is zero-sum when one considers entire set of exchange positions $\overline{\mathcal{P}}_n$. To see, this note the following:

$$\sum_{\mathbf{p} \in \overline{\mathcal{P}}_n} \gamma_t p_t (bq) = \sum_{\mathbf{p} \in \overline{\mathcal{L}}(\mathcal{P}_n)} \gamma_t p_t (q) - \sum_{\mathbf{p} \in \overline{\mathcal{S}}(\mathcal{P}_n)} \gamma_t p_t (q) = 0 \quad (4)$$

where the last equality follows from the fact that $\mathbf{p} \in \overline{\mathcal{L}}(\mathcal{P}_n) \iff \bar{\mathbf{p}} \in \overline{\mathcal{S}}(\mathcal{P}_n)$.

Example. To make the linear funding rule concrete, take $\kappa = 1$ and suppose the oracle and mark follow $(\hat{p}_0, p_0) = (1, 1)$, $(\hat{p}_1, p_1) = (1.5, 1.4)$, $(\hat{p}_2, p_2) = (1.25, 1.3)$, with quantities unchanged (so $\frac{L}{S} = \frac{3}{5}$). Then with $\kappa = 1$, $\gamma_1 = \frac{L}{S} - \frac{p_1}{\hat{p}_1} = 0.6 - \frac{1.4}{1.5} \approx -0.3333$ and $\gamma_2 = 0.6 - \frac{1.3}{1.25} = -0.44$ (longs pay shorts at both steps). Per-step funding cash to each position is $\gamma_t p_t (b_i q_i)$ (positive means received, negative means paid):

t	$\mathbf{p}_A$	$\mathbf{p}_B$	$\mathbf{p}_C$	$\mathbf{p}_D$	$\mathbf{p}_E$
0	0.0000	0.0000	0.0000	0.0000	0.0000
1	-0.4667	-0.4667	+1.8667	-0.4667	+0.4667
2	-0.5720	-0.5720	+2.2880	-0.5720	+0.5720

Values follow Eq. (3) using $(\hat{p}_t, p_t)$ above.

Profit and Loss. Given funding rates, one can write the explicit profit and/or loss that a trader faces during the lifetime of their position. We define the profit-and-loss (PNL) of duration $T > 0$, $\text{PNL}_{s:T} : (\mathbf{R}_+^4 \times \{-1, 1\}) \times \mathbf{R}_+^T \times \mathbf{R}_+^T \rightarrow \mathbf{R}$ as:

$$\text{PNL}_{s:T}((q, c, \ell, t, b), p_{1:T}, \hat{p}_{1:T}) = \mathbf{1}_{s \leq t < T} \left(bq(p_{\hat{T}} - p_t) + \Gamma((q, c, t, b), t, \hat{T}) \right) \quad (5)$$

where $\hat{T} = \min(T, \tau)$ is the last time the position is solvent. We will use the shorthand $\text{PNL}_T = \text{PNL}_{0:T}$. We will formally define the quantity τ shortly when we discuss liquidations, but for now think of it as a martingale stopping time for position solvency. This condition states that the total position of the user is their collateral plus the net change from funding costs. A natural consequence of the zero-sum nature of funding is that the total PNL is also zero-sum:

$$\sum_{\mathbf{p} \in \bar{\mathcal{P}}_n} \text{PNL}_T(\mathbf{p}, p_{1:T}, \hat{p}_{1:T}) = \sum_{\mathbf{p} \in \bar{\mathcal{P}}_n} bq(p_{\hat{T}} - p_0) + \Gamma(\mathbf{p}, t, \hat{T}) = 0$$

The first term sums to zero, because there will be offsetting positions with $b = -1, b = 1$ and the second term is zero via (4).

Example. Under the additive funding convention in Eq. (5) with the prices, open interest, and funding above, PNL over $[0, 1]$ and $[0, 2]$ is:

$$\begin{aligned} \text{PNL}_{0:1}(\mathbf{p}_A) &= -0.0667, & \text{PNL}_{0:2}(\mathbf{p}_A) &= -0.7387, \\ \text{PNL}_{0:1}(\mathbf{p}_B) &= -0.0667, & \text{PNL}_{0:2}(\mathbf{p}_B) &= -0.7387, \\ \text{PNL}_{0:1}(\mathbf{p}_C) &= +0.2667, & \text{PNL}_{0:2}(\mathbf{p}_C) &= +2.9547, \\ \text{PNL}_{0:1}(\mathbf{p}_D) &= -0.0667, & \text{PNL}_{0:2}(\mathbf{p}_D) &= -0.7387, \\ \text{PNL}_{0:1}(\mathbf{p}_E) &= +0.0667, & \text{PNL}_{0:2}(\mathbf{p}_E) &= +0.7387. \end{aligned}$$

Equity. For each position $\mathbf{p}_{i,t}$, one can view the assets of the position as $c_{i,t}$ (*e.g.* the margin posted by the user) and the liabilities of the users as $-\text{PNL}_T(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T})$. The liabilities are the negative of PNL because if the user has a loss, then they have positive liabilities to the exchange (*e.g.* they owe the exchange money). On the other hand, if the user has a gain, then the exchange owes the user money (and is a negative liability). The *equity* of a position $\mathbf{p}_{i,t}$, $e(\mathbf{p}_{i,t})$ is simply the assets of the position minus the liabilities:

$$e(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}) = c_{i,t} + \text{PNL}_T(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}) \quad (6)$$

For notational convenience, we will write the terminal equity shorthand

$$e_T(\mathbf{p}_{i,t}) = e(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}).$$

A position is said to be *totally insolvent* if $e_T(\mathbf{p}_{i,t}) < 0$. Traders can update their collateral $c_{i,t}$ or close positions to avoid fully insolvency. We define the random variable $\tau(\mathbf{p}_{i,t}, p_{1:t})$ as the first time that a position is insolvent, *i.e.*,

$$\tau(\mathbf{p}_{i,t}, p_{1:t}) = \min\{s \leq t : e(\mathbf{p}_{i,s}, p_{1:s}) \leq \mu p_s q_i\}$$

For any time index t , we will refer to the *winner* and *loser* index sets as

$$\mathcal{W}_t = \{i : e(\mathbf{p}_{i,t}) > 0\}, \quad \mathcal{L}_t = \{i : e(\mathbf{p}_{i,t}) < 0\}.$$

Write their sizes as $k_t = |\mathcal{W}_t|$ and $m_t = |\mathcal{L}_t|$. We will use $\mathcal{W}_T, \mathcal{L}_T$ for the terminal sets at horizon T .

Example. To illustrate the maintenance test, let $\mu = 0.10$ and reuse the $[0, 1]$ price and funding path above. Equities at $t = 1$ (using $c_{i,1} = c_i$) are

$$\begin{aligned} e(\mathbf{p}_A) &= 2 - 0.0667 \approx 1.9333, & e(\mathbf{p}_B) &= \frac{2}{3} - 0.0667 \approx 0.6000, & e(\mathbf{p}_C) &= \frac{8}{3} + 0.2667 \approx 2.9334, \\ e(\mathbf{p}_D) &= \frac{2}{19} - 0.0667 \approx 0.0386, & e(\mathbf{p}_E) &= \frac{10}{99} + 0.0667 \approx 0.1677. \end{aligned}$$

At $t = 2$ (using the same path and $c_{i,2} = c_i$), equities are

$$\begin{aligned} e(\mathbf{p}_A) &\approx 1.2613, & e(\mathbf{p}_B) &\approx -0.0720, & e(\mathbf{p}_C) &\approx 5.6214, \\ e(\mathbf{p}_D) &\approx -0.6334, & e(\mathbf{p}_E) &\approx 0.8397. \end{aligned}$$

Note that under this price move, there are now negative equity positions that are liquidatable.

Maintenance Margin. As spot and future prices change, an exchange needs to enforce dynamic collateral requirements to avoid insolvencies. The *maintenance margin ratio*, m_μ , represents the ratio of collateral to notional exposure that a trader must maintain under adverse price moves. With $m_\mu \in (0, 1)$ and equity from eq. (6), a position is maintenance-insolvent (and liquidatable) when

$$e(\mathbf{p}_{i,t}) \leq m_\mu p_t |q_i|. \quad (7)$$

Note that maintenance margin depends on the equity of a position (*e.g.* includes the position's PNL) whereas initial margin only depends on the initial cash position. If a user violates (7) due to a price move, they can add collateral within a specified time interval to keep their position open. Otherwise, the venue liquidates the position.

Example. Take $m_\mu = 0.10$ and reuse the $[0, 1]$ path above with $p_1 = 1.4$ and the computed equities

$$e(\mathbf{p}_A) \approx 1.9333, \quad e(\mathbf{p}_B) \approx 0.6000, \quad e(\mathbf{p}_C) \approx 2.9334, \quad e(\mathbf{p}_D) \approx 0.0386, \quad e(\mathbf{p}_E) \approx 0.1677.$$

The maintenance thresholds are $m_\mu p_1 |q| = (0.14, 0.14, 0.56, 0.14, 0.14)$ for (A, B, C, D, E) . Here $e(\mathbf{p}_D) \leq m_\mu p_1 |q|$ while others remain above the threshold. Position D is maintenance-insolvent and would need to add collateral or be liquidated.

Leverage Mass. The final quantity we will define aims to measure how much leverage the winning (positive equity) and losing (negative equity) sides have on the exchange. This will serve as a coarse measure of an exchange's risk exposure and will be used in §5. We define a natural aggregate leverage quantity for both the winning and losing sides that normalizes leverage by equity, allowing us to compare traders with very different margin ratios on a common “true asset composition” scale. For $i \in \mathcal{W}_t$ define the *winner effective leverage* $\lambda_{i,t}^+ = \frac{n_{i,t}}{e(\mathbf{p}_{i,t})}$, and for $i \in \mathcal{L}_t$ the *loser effective leverage* $\lambda_{i,t}^- = \frac{n_{i,t}}{|e(\mathbf{p}_{i,t})|}$. The *winner* and *loser leverage masses* are

$$\ell_t^+ = \sum_{i \in \mathcal{W}_t} \lambda_{i,t}^+, \quad \ell_t^- = \sum_{i \in \mathcal{L}_t} \lambda_{i,t}^-, \quad \bar{\ell}_t^\pm = \frac{\ell_t^\pm}{n} \text{ (per-trader)}. \quad (8)$$

Note that unlike the leverage $\ell_{i,t}$, we divide by the equity (which can be larger or smaller than leverage, depending on the trader's PNL).

Example. Consider $t = 1$ with $p_1 = 1.40$. Then $n_{i,1} = p_1 q_i = 1.4$ for $i \in \{A, B, D, E\}$ and 5.6 for C . Using the equities at $t = 1$ above, all five are winners and

$$\begin{aligned}\lambda_{A,1}^+ &= \frac{1.4}{1.9333} \approx 0.7246, & \lambda_{B,1}^+ &= \frac{1.4}{0.6000} = 2.3333, & \lambda_{C,1}^+ &= \frac{5.6}{2.9334} \approx 1.9090, \\ \lambda_{D,1}^+ &= \frac{1.4}{0.0386} \approx 36.269, & \lambda_{E,1}^+ &= \frac{1.4}{0.1677} \approx 8.353.\end{aligned}$$

Hence $\ell_1^+ \approx 49.59$; the large $\lambda_{D,1}^+$ reflects D 's very small equity at $t = 1$. Recall D 's opening leverage was only $\ell_{D,0} \approx 9.5$, so effective leverage can greatly exceed raw leverage when equity has been eroded by losses. At $t = 2$ with $p_2 = 1.30$, $n_{i,2} = p_2 |q_i|$ equals 1.3 for $i \in \{A, B, D, E\}$ and 5.2 for C . From the equity example above, the winner/loser sets are $\mathcal{W}_2 = \{A, C, E\}$ with $e_A \approx 1.2613$, $e_C \approx 5.6214$, $e_E \approx 0.8397$ and $\mathcal{L}_2 = \{B, D\}$ with $e_B \approx -0.0720$, $e_D \approx -0.6334$. Winner effective leverages are

$$\lambda_{A,2}^+ = \frac{1.3}{1.2613} \approx 1.031, \quad \lambda_{C,2}^+ = \frac{5.2}{5.6214} \approx 0.925, \quad \lambda_{E,2}^+ = \frac{1.3}{0.8397} \approx 1.548,$$

so $\ell_2^+ \approx 3.504$. Loser effective leverages are

$$\lambda_{B,2}^- = \frac{1.3}{|-0.0720|} \approx 18.05, \quad \lambda_{D,2}^- = \frac{1.3}{|-0.6334|} \approx 2.052,$$

so $\ell_2^- \approx 20.102$. The ratio of leverage masses $\ell_2^- / \ell_2^+ \approx 5.7$ shows that, at this horizon, most effective leverage sits on the losing side of the book. Here B and D have the same notional $n_{i,2}$ but very different effective leverages, illustrating how $\lambda_{i,t}^\pm$ can diverge from the raw leverage $\ell_{i,t}$ when equities differ.

2.2 Liquidations

When a user's position is insolvent, *i.e.*, (7) holds, the exchange must *liquidate* the position — that is, remove some or all of the position from $\mathcal{P}_n$. Liquidations are used to ensure that the exchange itself remains solvent, which ensures that traders can withdraw their funds and profits as expected. The process of liquidating a position is probabilistic, however, and can fail in a number of ways.

We will illustrate the high-level process of liquidation using bankruptcy and liquidation prices. Then we will describe how exchanges can lose solvency when liquidations fail, leading to the usage of autodeleveraging mechanisms. Our description of liquidations is minimal, avoiding formal mathematical description unless needed, as we are concerned with liquidation failure versus liquidation mechanics. We refer the interested reader to [9, 90, 95, 104] for detailed accounts of liquidation mechanisms and modeling on centralized and decentralized venues.

2.2.1 Liquidation Prices

Before defining how liquidations work, we need to define criteria for when a position is eligible for liquidation. These criteria depend on when certain prices are reached where a trader position has low or negative equity. We focus on defining three prices: bankruptcy price (zero equity), liquidation price (low equity), and execution price (actual realized price of a liquidation). In practice, liquidation mechanisms often use more price-based variables to decide on execution. However, all mechanisms define these three prices, allowing us to abstract liquidations to mechanisms involving these prices.

Bankruptcy Price. The *bankruptcy price* of a position $\mathbf{p}_{i,t}$, $p^{bk}(\mathbf{p}_{i,t})$, is the highest price at which the position is totally insolvent. The threshold condition for a position to be totally insolvent is $e_T(\mathbf{p}_{i,t}) = 0$:

$$0 = e_T(\mathbf{p}_{i,t}) = c_{i,t} + \text{PNL}_T(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}) = c_{i,t} + b_i q_i (p^{bk}(\mathbf{p}_{i,t}) - p_t) + \Gamma(\mathbf{p}_{i,t}, t, \hat{T})$$

Rearranging this gives:

$$p^{bk}(\mathbf{p}_{i,t}) = \max \left(p_t - \frac{c_{i,t} + \Gamma(\mathbf{p}_{i,t}, t, \hat{T})}{b_i q_i}, 0 \right) \quad (9)$$

One can view the bankruptcy price as the worst-case liquidation price, in the sense that if the position is liquidated at any $p < p^{bk}(\mathbf{p}_{i,t})$, then other traders of the exchange will realize a loss. Such a loss, which is either borne by the exchange or by other traders, is known as *bad debt* (see §2.3).

The bankruptcy price also constrains the maximum leverage that a position can have. Suppose that the funding is zero, *e.g.* $\Gamma(\mathbf{p}_{i,t}, t, \hat{T}) = 0$. The bankruptcy price then simplifies to:

$$p^{bk}(\mathbf{p}_{i,t}) = \max \left(p_t - \frac{c_{i,t}}{b_i q_i}, 0 \right) = \max \left(p_t - \frac{b_i}{\ell_{i,t}} p_t, 0 \right) = p_t \max \left(1 - \frac{b_i}{\ell_{i,t}}, 0 \right) \quad (10)$$

This formula represents the common maxim that a perpetuals position with $\ell_i \geq 1$ times leverage will be liquidated when the price moves by $\frac{1}{\ell_i}\%$ in the direction opposite to b_i ; see, for example, [21, 66]. For instance, a maximal 10x leveraged position will be liquidated when the price moves by 10% from the initial price. Detailed numerical examples of bankruptcy prices are provided in Appendix B.

Liquidation Price. To avoid bad debt, an exchange defines a *liquidation price* $\hat{p}^{liq}(\mathbf{p}_{i,t}) \geq p^{bk}(\mathbf{p}_{i,t})$, when the position is liquidatable at time t . Generally, the liquidation price is a spot price derived from the oracle versus a futures price quoted by the exchange. This makes market manipulation — creating and executing non-economically rational orders to force a liquidation — costly to perform. When the maintenance margin condition (7) is satisfied, the exchange attempts to partially liquidate the position. The μ -liquidation price, $p^{liq}(\mathbf{p}_{i,t}, \mu)$ is defined as the maximal price where (7) holds. To avoid ambiguity about the base price used in p^{bk} , we write $\hat{p}^{liq}$ directly in terms of the entry price p_{t_i} (independent of bankruptcy):

$$\hat{p}^{liq}(\mathbf{p}_{i,t}, \mu) = \begin{cases} \max \left(\frac{p_{t_i} - (c_{i,t} + \Gamma(\mathbf{p}_{i,t}, t, \hat{T}))/q_i}{1 - \mu}, 0 \right), & b_i = +1, \\ \frac{p_{t_i} + (c_{i,t} + \Gamma(\mathbf{p}_{i,t}, t, \hat{T}))/|q_i|}{1 + \mu}, & b_i = -1. \end{cases}$$

Detailed numerical examples of liquidation prices are provided in Appendix B.

Execution Price. When $\hat{p}_t < p^{liq}(\mathbf{p}_{i,t}, \mu)$ a position is liquidatable. The venue then sells (if $b_i = +1$) or buys (if $b_i = -1$) a slice of size Δq , realizing an execution price $p^{exec}(\mathbf{p}_{i,t}, \mu, \Delta q)$. Whether a liquidation creates bad debt depends on the location of p^{exec} relative to the bankruptcy price p^{bk} :

- *Long* ($b_i = +1$): no shortfall if $p^{exec} \geq p^{bk}$; otherwise the realized shortfall is $D_t = (p^{bk} - p^{exec}) \Delta q$.
- *Short* ($b_i = -1$): no shortfall if $p^{exec} \leq p^{bk}$; otherwise the realized shortfall is $D_t = (p^{exec} - p^{bk}) \Delta q$.

Shortfalls are first absorbed by the insurance fund (§2.3) up to its balance. Any residual shortfall not covered by the insurance fund is socialized via ADL (§2.4). The choice of liquidation size Δq and impact parameter α together determines whether thin-equity positions can typically be closed without realizing bad debt. Detailed numerical examples of execution prices and shortfalls are provided in Appendix B.

2.2.2 Liquidation Mechanics

Given liquidation prices, a natural question is how liquidations are mechanically executed. At a high level, the exchange seizes a low equity position's cash and collateral and sells it to the market. Such a sale can realize a profit or a loss for the exchange and is best thought of as a trading strategy itself. We will provide a high-level description of liquidation mechanisms as trading strategies, with the caveat that there are many idiosyncrasies in practical implementations (see, *e.g.* [20, 24, 68] for venue documentation and differences between auctions, order-book liquidations, and RFQ-style closures).

Liquidation Costs. Most exchanges charge penalties to users who are liquidated as a means of disincentivizing liquidations and moral hazard. These fees come in three flavors: fixed charges, insurance fund fees, and liquidation incentives. The fixed fees correspond to reimbursement for gas and/or operational costs that an exchange realizes for performing a liquidation. The insurance fund fees are proportional to the liquidation price and allow for the exchange to cover bad debt. Finally, the liquidation incentives are used to encourage third-party actors known as liquidators to hold the risk of buying the position of size Δq from the exchange and exiting it profitably.

In decentralized exchanges, liquidators are usually any market participant with enough capital who can connect to the exchange. On the other hand, in centralized exchanges, liquidators are usually whitelisted parties approved by the exchange to perform liquidations. We denote the set of liquidation costs for a liquidation of size Δq at time t as $\tau_t(\Delta q) \in \mathbf{R}_+$, which represents the cash cost paid by the user upon liquidation. Detailed numerical examples of liquidation costs are provided in Appendix B.

Liquidation Strategies. The policy by which an exchange chooses the liquidation quantity Δq is known as the *liquidation strategy*, $L : \mathcal{P}_n \times \mathbf{R}_+^T \times \mathbf{R}_+^T \rightarrow \mathbf{R}$. The strategy $L(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T})$ outputs a quantity Δq to liquidate. Generally speaking, the strategy has some model of the environment (*e.g.* the price impact function of the spot and futures exchanges) and utilizes that to pick the quantity. Most exchanges utilize a simple a greedy liquidation strategy that chooses the minimal Δq to ensure that the equity does not satisfy (7). Such strategies choose Δq such that an equality of the form

$$e_T(\mathbf{p}_{i,t}) + b \Delta q (p_t^{exec} - p_t) - \tau_t(\Delta q) = \mu p_t (q - \Delta q) \quad (11)$$

approximately holds. For notational simplicity, we assume that $L(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}) = 0$ if the position $\mathbf{p}_{i,t}$ is not liquidatable.

As an explicit example, suppose that we have linear price impact, *i.e.*, $p^{exec} = p_t \mp \frac{\alpha}{2} \Delta q$. We receive a quadratic equation in Δq :

$$\frac{\alpha}{2} \Delta q^2 - \mu p_t \Delta q + (\mu p_t q - e_T + \tau) = 0,$$

The greedy liquidation strategy simply takes the smallest feasible root of this quadratic and uses it as the liquidation quantity Δq . We refer the reader to Appendix B for detailed numerical examples of this strategy and related liquidation procedures.

Bad Debt. A liquidation for position $\mathbf{p}_{i,t}$ generates *bad debt* if the post-liquidation equity (including execution price and fees) is non-positive. Formally, we define an adjusted terminal equity $\tilde{e}_T(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}, \Delta q)$ that incorporates liquidation costs:

$$\tilde{e}_T = e_T + b \Delta q (p^{exec} - p_t) - \tau_t(\Delta q) \quad (12)$$

Intuitively, relative to marking the entire position at p_t , realizing a slice Δq at p^{exec} changes equity by the slippage term $b \Delta q (p^{exec} - p_t)$, and fees reduce equity via $\tau_t(\Delta q)$. We say that a liquidation creates bad debt if $\tilde{e}_T(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}, \Delta q) < 0$, *i.e.*, the liquidation leaves a residual liability for the venue.

Note that in practice, a liquidation strategy might retry or reattempt to liquidate the position repeatedly. We elide formulating the details of such a liquidation strategy here for simplicity, but note that our model can be easily extended to account for this. We define the *total bad debt* or *shortfall* of an exchange, D_t given a liquidation strategy L is,

$$D_t(L) = \sum_{\mathbf{p} \in \mathcal{P}_n} \max(0, -\tilde{e}(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}, \Delta q_i)) = \sum_{\mathbf{p} \in \mathcal{P}_n} -\tilde{e}(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}, \Delta q_i)_- \quad (13)$$

where $\Delta q_i = L(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T})$. This represents the shortfall that the exchange must cover to be solvent (§2.3). Detailed numerical examples of bad debt are provided in Appendix B.

Anatomy of a Liquidation. Given the bankruptcy, liquidation, and execution prices, we can now describe the high-level algorithm that liquidations follow (see Appendix B). We note that many live liquidation systems will have much more complex liquidation algorithms. These complexities deal with the coordination costs of coordinating many parties (*e.g.* oracle provider, liquidators, spot order book liquidity) and precise models that exchanges use for their liquidation strategy. However, we effectively lump all of these complexities into the definition of the liquidation strategy. A detailed algorithmic description of the liquidation loop and a step-by-step example are provided in Appendix B.

2.3 Exchange Solvency


In the previous section, we focused on liquidations as a means to remove individual trader solvencies. However, there is also a ‘macroscopic’ notion of solvency at the exchange level. The goal of the exchange is to make sure that traders can realize their full profits and losses given the cash collateral held at the exchange. In rare cases, an exchange might not be able to successfully execute a liquidation strategy, leading to the exchange being unable to pay unrealized but earned profits to some users.

In this section, we define solvency for exchanges and how exchanges use *insurance funds* to try to ensure solvency. An exchange’s insurance fund is the first line of defense against insolvency. However, if the insurance fund has insufficient balance, the exchange will need to utilize a autodeleveraging mechanism to socialize losses. In this section, we will define solvency for exchanges and how insurance funds are constructed.


Exchange Solvency. The *solvency* of an exchange is simply the total equity of all positions:

$$\text{Solv}_T(\mathcal{P}_n) = \sum_{\mathbf{p} \in \mathcal{P}_n} e_T(\mathbf{p}) = \sum_{\mathbf{p} \in \mathcal{P}_n} c_p + \text{PNL}_T(\mathbf{p}, p_{1:T}, \hat{p}_{1:T}) \quad (14)$$


An exchange is said to be *insolvent* if the following condition holds $\text{Solv}_T(\mathcal{P}_n) \leq 0$. The insolvency condition corresponds to when the total cash held by the exchange on behalf of users is less than



(a) Baseline 10% liquidation buffer: mark price $p_t = 0.88p_0$ is 10% above bankruptcy $\hat{p}_t = 0.80p_0$.



(b) Price recovery: increasing the mark price by $\Delta p_t = +0.06p_0$ and the bankruptcy level by $\Delta \hat{p}_t = +0.13p_0$ moves them to $p_{t+\Delta} = 0.94p_0$ and $\hat{p}_{t+\Delta} = 0.93p_0$, leaving every account with positive equity and a residual mark-bankruptcy buffer of about 1%.



(c) Example of a negative-equity account inside the sorted stack.

Figure 1: Sorted equity profiles for stylized liquidation examples. Negative positions (red) appear on the left, positive positions (green) on the right. Dashed lines highlight the bankruptcy level and liquidation triggers.

the total accrued profits. This means that if this insolvency condition is hit, the exchange will not be able to payout all profitable traders when they withdraw from the exchange. Note that this is a global condition on $\mathcal{P}_n$ as opposed to a per user constraint.

Most exchanges have a solvency threshold, $\delta > 0$, such that if it takes a global loss of δ units of equity to become truly insolvent, then the exchange is deemed approximately insolvent. This notion of approximate insolvency is useful in practice, as it allows an exchange to have a risk parameter for tuning how aggressively it enters the ADL regime. Formally, we define an exchange to be δ -insolvent if the following condition holds $\text{Solv}_T(\mathcal{P}_n) \leq \delta$.

Comparison to overcollateralized lending. In overcollateralized lending, where a user uses one crypto asset as collateral to borrow another crypto asset, there is a similar notion of a user's equity position. However, in overcollateralized lending, the user's equity position is not global, but rather is local to the user. The goal of a lending protocol is to ensure that the equity $e_{\text{lend}}(p) \geq 0$ for every user's position p . Ensuring per account solvency condition inherently forces overcollateralized

lending protocols to offer far lower leverage than perpetuals exchanges. From a mathematical standpoint, perpetual solvency is easier to satisfy, allowing for higher leverage — there is only one global constraint that has to hold versus $O(n)$ local constraints for n users. However, in perpetuals exchanges, users who are solvent can face losses from users who are insolvent which cannot happen in isolated overcollateralized lending markets like Morpho.

Insurance Funds. Many exchanges maintain an insurance fund to cover bad debt that arises from failed liquidations. When a positive shortfall $D_t > 0$ is realized, the insurance fund is used to reduce D_t to as close to zero as possible. The insurance fund is typically funded using a portion of the exchange’s revenue from transaction, liquidation, and funding fees.

Insurance reserves in crypto funds were first popularized alongside perpetual futures in the mid-2010s (notably on BitMEX) [104]. A majority of perpetuals venues (centralized and on-chain) maintain such reserves (*e.g.* BitMEX [24], Binance [22], Drift [46], Bybit [30]). To illustrate the size of these funds, we note that mid-2025 estimates place OKX and Bybit insurance funds at roughly \$300M and \$150M, respectively [70].

We will first provide a simple model for how the insurance fund evolves over time. Let $\text{IF}_t \geq 0$ denote the insurancefund balance at time t . Furthermore, let $\mathcal{T}_t$ be the multiset of liquidations at time t with sizes $\{\Delta q_j\}_{j \in \mathcal{T}_t}$ and execution costs $\tau_t(\Delta q_j) \in \mathbf{R}_+$ (*cf.* eq. (12)). Let $V_t \geq 0$ denote the gross traded volume (absolute quantity) in the contract at time t . We introduce three parameters to model insurance fund growth: $\alpha, \beta, \eta \in [0, 1]$. The α, β, η parameters controls how much of liquidation, funding, and trading fee revenue is added IF_t , respectively.

Given these parameters, the time evolution of the insurance fund is given by:

$$\text{IF}_{t+1} = \text{IF}_t + \alpha \sum_{j \in \mathcal{T}_t} \tau_t(\Delta q_j) + \eta p_t V_t + \beta \sum_{i=1}^n |\gamma_t| p_t q_i - \min\{\text{IF}_t, D_t\} \quad (15)$$

Note that the positive contributions represent the revenue sharing terms whereas the negative contribution represents the cover of a shortfall using the insurance fund.

The main quantity of interest, given IF_t , is the *residual shortfall*. This represents the amount of bad debt that cannot be covered by the insurance fund and needs to be socialized via ADL (§2.4). Formally, we define the residual shortfall as:

$$R_t = D_t - \min\{\text{IF}_t, D_t\} = (D_t - \text{IF}_t)_+ \quad (16)$$

We note that when $R_t > 0$, an ADL mechanism is triggered. We term times t such that $R_t > 0$ as *breach times*. The frequency of breach events is referred to as *breach frequency* or *breach rate*.

Example. Continuing from the previous liquidation example, the bad debt realized at $t = 4$ by $\mathbf{p}_E$ is

$$D_4 = (p_E^{exec} - p^{bk}(\mathbf{p}_{E,4})) \Delta q_E \approx 0.399.$$

Let IF_4 denote the pre-coverage fund balance. The coverage is $\min\{\text{IF}_4, D_4\}$ and the ADL residual is $R_4 = (D_4 - \text{IF}_4)_+$. We assume that the fund evolves via eq. (15) and model liquidation fees as affine in notional size:

$$\tau_t(\Delta q) = \tau^{fix} + \phi^{mark} p_t \Delta q + \phi^{exec} p_t^{exec} \Delta q,$$

where τ^{fix} is a fixed fee and ϕ^{mark}, ϕ^{exec} are proportional rates on mark and execution price, respectively. Using the Hyperliquid fee schedule $(\tau^{fix}, \phi^{mark}, \phi^{exec}) = (0, 20 \text{ bps}, 10 \text{ bps})$ (see §2.2.2), the fees are

$$\tau_4(\Delta q_E) = 0 + 20 \cdot 10^{-4} \cdot p_4 \Delta q_E + 10 \cdot 10^{-4} \cdot p_E^{exec} \Delta q_E \approx 0.00365.$$

Assuming $\gamma_4 = 0$ and neglecting trading volume V_4 for simplicity:

- *Sufficiently sized fund* ($R_4 = 0$). If $\text{IF}_4 \geq D_4$, the debt is fully covered. No ADL occurs. The balance updates to

$$\text{IF}_5 = \text{IF}_4 + \alpha \tau_4(\Delta q_E) - D_4.$$

- *Insufficiently sized fund* ($R_4 > 0$). If $\text{IF}_4 < D_4$, the fund is depleted to zero and pays IF_4 . The residual $R_4 = D_4 - \text{IF}_4$ is socialized via ADL (§2.4). The new balance is simply the inflows:

$$\text{IF}_5 = \alpha \tau_4(\Delta q_E).$$

Increasing α , β , or η builds IF_t faster, reducing the probability of $R_t > 0$.

Optimal Insurance Fund Size. A natural question to ask is, what is the optimal size for an insurance fund? While this question is nuanced in practice, we provide a simple mean-field model that provides some intuition. In particular, we determine the optimal fund size K^* by minimizing a total cost function that balances the opportunity cost of capital against the expected franchise damage from ADL events. Let $r > 0$ denote the opportunity cost of holding capital (*e.g.* the risk-free rate) and $\kappa > 0$ denote the reputation cost per unit of socialized loss (reflecting lost future volume). The exchange minimizes

$$\min_{K \geq 0} rK + \kappa \mathbf{E}[(D_T - K)_+].$$

This is a classic newsvendor-type problem [12]. Classical result show that the optimal size is the Value-at-Risk of the deficit distribution at quantile $1 - r/\kappa$:

$$K^* = \text{VaR}_{1-r/\kappa}(D_T) = \inf\{x : \mathbf{Prob}(D_T \leq x) \geq 1 - r/\kappa\}.$$

For completeness, we provide a derivation in Appendix B.8. Intuitively, as the reputation cost κ increases relative to the cost of capital r , the exchange holds a larger buffer to push ADL events further into the tail.

2.4 Autodeleveraging

Any mechanism for reducing or closing profitable user positions to reduce an exchange's insolvency is termed *autodeleveraging* (ADL). Mathematically, ADL mechanisms attempt to liquidating positions in order to reduce the residual R_t . This inherently means that users with positive equity (*e.g.* are profitable) will be forced to close their positions and not realize their full profit.

We first note that ADL mechanisms are not deterministic and foolproof. For example, if an exchange is already insolvent, $\text{Solv}_T(\mathcal{P}_n) \leq 0$, then no deterministic ADL mechanism can return solvency. To see this, we define four key quantifies, the total shortfall $D_T(\mathcal{P}_n)$, the maximum shortfall $\Delta_T(\mathcal{P}_n)$, the total winner's equity $W_T(\mathcal{P}_n)$, and the max-winner $\omega_T(\mathcal{P}_n)$:

$$\begin{aligned} D_T(\mathcal{P}_n) &= \sum_{\mathbf{p} \in \mathcal{P}_n} (-e_T(\mathbf{p}))_+ \geq 0 & \Delta_T(\mathcal{P}_n) &= \max_{\mathbf{p} \in \mathcal{P}_n} (-e_T(\mathbf{p}))_- \geq 0 \\ W_T(\mathcal{P}_n) &= \sum_{\mathbf{p} \in \mathcal{P}_n} e_T(\mathbf{p})_+ \geq 0 & \omega_T(\mathcal{P}_n) &= \max_{\mathbf{p} \in \mathcal{P}_n} e_T(\mathbf{p})_+ \geq 0 \end{aligned}$$

For notational simplicity, we will use Roman letters for total quantities and Greek letters for maximum or extreme value quantities. The shortfall represents the total amount of negative equity whereas the winners' equity represents the total amount of positive equity. By definition,

$\text{Solv}_T(\mathcal{P}_n) = W_T(\mathcal{P}_n) - D_T(\mathcal{P}_n)$, so if $\text{Solv}_T(\mathcal{P}_n) \leq 0$, then $D_T(\mathcal{P}_n) \geq W_T(\mathcal{P}_n)$, *i.e.*, no amount of forced liquidations of positive equity can cover the shortfall.

Instead, ADL mechanisms attempt to liquidate a fraction of the shortfall, then allow the market to react for some time (*e.g.* for users to close positions or post more margin and spot prices to update), before attempting a future deleveraging (if needed). This means that the ADL process should be viewed probabilistically, as a sequence of decisions that are dynamically adjusted based on market conditions. We will demonstrate that from this perspective, ADL can be formulated as a standard reinforcement learning problem.

ADL Policies. The fundamental object of reinforcement learning is the policy, which maps states of the world to actions. For ADL, the positions $\mathcal{P}_n$ serve as the state space, whereas the actions are the fraction of shortfall to socialize and which positions are socialized. While the majority of exchanges (*e.g.* Binance and Hyperliquid) use queue-based policies that greedily rank positions by PNL and leverage, we will define a formalism over a larger class of potential ADL policies. This larger class will enable us to find more efficient and fair ADL policies than queue-based models and includes pro-rata ADL policies used by smaller exchanges such as Drift and Paradex.

An *ADL policy* $\pi(\mathcal{P}_n)$ maps a perpetuals exchange to a fraction of shortfall to socialize, $\theta \in [0, 1]$, and a set of haircuts to apply to positions, $h \in [0, 1]^n$. We term the fraction θ the *severity* of the ADL policy. For notational convenience we define θ_π, h_π as the severity and haircuts induced by an ADL policy π . We say an ADL policy π is *valid* if the following three constraints always hold:

$$\sum_{\mathbf{p} \in \mathcal{P}_n} h_{\pi,i} e_T(\mathbf{p})_+ = \theta_\pi D_T(\mathcal{P}_n) = - \sum_{\mathbf{p} \in \mathcal{P}_n} h_{\pi,i} e_T(\mathbf{p})_- . \quad (17)$$

$$\theta_\pi D_T(\mathcal{P}_n) \leq W_T(\mathcal{P}_n). \quad (18)$$

where p_i is the i th trader's position. The first two conditions (17) ensure that when an ADL policy is executed, only a $\theta\%$ of the deficit is socialized from the winners to the losers (*i.e.*, a budget balance constraint). The third condition (18) ensures that θ_π is chosen such that one does not need to haircut more equity than what the winners have (*i.e.*, a feasibility constraint). We define the post-policy deficit D_T^π , winners W_T^π , and max-winner M_T^π as

$$\begin{aligned} D_T^\pi &= \theta_\pi D_T(\mathcal{P}_n) & \Delta_T^\pi &= \theta_\pi \Delta_T(\mathcal{P}_n) \\ W_T^\pi &= \sum_{\mathbf{p} \in \mathcal{P}_n} (1 - h_{\pi,i})(e_T(\mathbf{p}))_+ & \omega_T^\pi &= \max_{\mathbf{p} \in \mathcal{P}_n} (1 - h_{\pi,i})(e_T(\mathbf{p}))_+ \end{aligned}$$

Our goal is to construct policies that minimize an objective function that balances individual user profits with global exchange solvency.

Given an ADL policy outputting (θ_π, h_π) , the positive equity positions $\mathcal{W}_T = \{i : e_T(\mathbf{p}_i) > 0\}$ are haircut by $h_{\pi,i}$, leading to equities $e_T(\mathbf{p}_i) \leftarrow (1 - h_{\pi,i})e_T(\mathbf{p}_i)$ for $i \in \mathcal{W}_T$. For instance, if we have equities $e_T = (10, 5, -5)$ and the ADL policy outputs haircuts $(0.5, 0, -1)$, then the post-ADL equities are $e_T = (5, 5, 0)$. On the other hand, if the ADL policy outputs haircuts $h = (0.25, 0.5, -1)$, then the final equities after execution of the policy are $e_T = (7.5, 2.5, 0)$.

Examples of ADL Policies. We will first provide two canonical examples: *queueing* (or leverage ranking) and *pro-rata*. These examples represent the solvency policies of virtually all live perpetuals exchanges as of November 2025. For both of these policies, the choice of θ_π is independent of the choice of h_π .

Queue. The policy that the largest centralized exchange, Binance, and the largest decentralized exchange, Hyperliquid, use is the Queue policy. The implementation at these exchanges uses the PNL-leverage ranking. However, the Queue model is somewhat more generic, as we define it below.

At a high level, the Queue strategy is a greedy algorithm that ranks positions using a *score function*, s_T . Once positions are ordered by score, positions are closed in order of score from highest to lowest, until the deficit is fully covered. We will denote this strategy as π_Q and define it via the algorithm that implements it.

Given a haircut budget $B = \theta_{\pi_Q} D_T(\mathcal{P}_n)$, the algorithm to choose h_i works as follows:

1. For each position $\mathbf{p} = (q, c, \ell, t, b) \in \mathcal{P}_n$, define the score⁹ $s_T(\mathbf{p}, p^{ref})$. For the PNL-leverage ranking, we have $s_T(\mathbf{p}, p^{ref}) = \ell \frac{\hat{p}_T}{p^{ref}}$.
2. Let $\sigma \in \text{Perm}(\mathcal{P}_n)$ be any permutation of the positions ranked by decreasing score $s_T(\mathbf{p}, p^{ref})$.
3. For $i \in [n]$, define $A_i = (B - \sum_{j=1}^i e_T(\mathbf{p}_{\sigma(j)})_+)_+$, where $e_T(\mathbf{p}_{\sigma(j)})$ is the equity of the j th largest position under the ranking σ .
4. Define $h \in [0, 1]^n$ as

$$h_{\pi_Q, \sigma(i)} = \begin{cases} 1 & \text{if } A_{i-1} - A_i = e_T(\mathbf{p}_{\sigma(i)})_+, A_i > 0 \\ \frac{A_{i-1}}{e_T(\mathbf{p}_{\sigma(i)})} & \text{if } A_i = 0, A_{i-1} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

Note this definition of h_{π_Q} ensures that $\sum_{i=1}^n h_{\pi_Q, i} e_T(\mathbf{p}_i) = B$, satisfying (17). On the other hand, since we only subtract positive equity positions (*e.g.* j such that $e_T(\mathbf{p}_{\sigma(j)})_+ > 0$), we satisfy (18) by construction.

There is one remaining question to address: how should we interpret the score $s_T(\mathbf{p}, p^{ref})$? Most exchanges justify this form for the score by arguing that it can represent the risk a single position holds, so that the exchange ranks positions to ADL by risk-level. Winning, positive equity positions that are autodeleveraged first will tend to be higher risk positions (*i.e.*, used more leverage for their winnings). The price terms in the numerator, $p_{\hat{T}}$, represents the last mark price of the position whereas the parameter p^{ref} represents the initial or opening price of the user.

Binance chooses the scoring parameter $p^{ref} = p^{bk}$ of the position, *i.e.*, $s_T^{\text{Binance}}(\mathbf{p}) = \ell \frac{p_{\hat{T}}(\mathbf{p})}{p^{bk}(\mathbf{p})}$. Hyperliquid, on the other hand, chooses the parameter p^{ref} to simply be the last mark price, $p^{ref} = \hat{p}_t$. In other words, Binance's choice effectively ranks a position based on the maximum possible PNL that a user could realize whereas Hyperliquid simply chooses the current PNL.

Pro-Rata Rules. This haircut strategy is even simpler than the Queue and has some benefits in terms of fairness, as we will show in §6. The Pro-Rata strategy, simply put, haircuts each winning user with positive equity by the same multiplicative fraction. The fraction is subject to budget constraints, as we describe below.

Given a valid feasible shortfall $\theta_{\pi} D_T(\mathcal{P}_n \leq W_T(\mathcal{P}_n))$, the pro-rata haircut policy is defined as:

$$h_{\pi_{PR}}(\mathbf{p}) = \frac{\theta_{\pi} D_T(\mathcal{P}_n)}{W_T(\mathcal{P}_n)} (e_T(\mathbf{p}))_+ \quad (20)$$

⁹We note that technically, many exchanges define the leverage used in the score differently (*i.e.*, instead of the score being linear in ℓ , some exchanges multiply by $\frac{\ell pq}{pq+c}$). This does not change our main results and mainly adds technical complications to the example.




Figure 2: ADL severity example comparing queue and pro-rata coloring. Purple shading equals the negative equity mass (deficit) while blue shading shows the haircut mass allocated to winning traders. The queue panel’s dashed blue block at rank 2 highlights residual equity when the queue method allows partial closures; exchanges that close winners fully (*e.g.* Hyperliquid) would shave this bar completely. Haircut mass matches deficit mass in each panel, illustrating severity $\theta = 0.50$.

This ranking simply says that the positive equity positions are socialized proportional to their size. Some exchanges slightly modify this formula to include leverage (*e.g.* Binance’s ADL Priority Index includes leverage and unrealized PNL [20]; Aevo documentation discusses leverage-weighted priority [3]), which penalizes higher leverage positions much like the PNL-leverage ranking:

$$h_{\pi_{LPR}}(\mathbf{p}) = \left(\frac{\ell e_T(\mathbf{p})_+}{\sum_{\mathbf{p} \in \mathcal{P}_n} \ell e_T(\mathbf{p})_+} \right) h_{\pi_{PR}}(\mathcal{P}_n) \quad (21)$$

Per-account constraints. Suppose that one wants to enforce some per-account constraints on the amount haircut. For instance, suppose that we guarantee that for user i , their haircut h_i always satisfies

$$h_i \leq \bar{h}_i \quad (22)$$

This corresponds to an exchange guaranteeing that a user will never lose more than $\bar{h}_i\%$ of their position in a single ADL round. Moreover, an exchange might also offer an absolute guarantee to users: if your equity is positive and sufficiently large, you will never be have your equity cut to below some threshold $\underline{e}_i$. This can be represented by the constraint

$$(1 - h_{\pi,i})e_T(\mathbf{p}_i) \geq \underline{e}_i \quad (23)$$

Numerous exchanges offer one or both of these guarantees on a per-ADL-round basis, including Hyperliquid, BitMEX, and Aevo [3, 24, 68].

Numerical example. Consider five accounts with equities $e_T = (10, 5, 1, -3, -12)$ and queue order $\sigma = (2, 1, 3, 5, 4)$ induced by the PNL-leverage scores. Let the queue policy choose severity $\theta_{\pi_Q} = \frac{1}{2}$, so the haircut budget is $B = \theta_{\pi_Q} D_T(\mathcal{P}_n) = 7.5$. The construction in (19) yields $A = (2.5, 0, 0, 0, 0)$ and haircuts $h = (0.25, 1, 0, 0, 0)$: the top winner is cut from 10 to 7.5, the second winner is fully wiped out, and the remaining accounts are unchanged, giving $e_T = (7.5, 0, 1, -3, -4.5)$. If, before the next ADL decision, only the losing accounts partially recover to $e_T = (7.5, 0, 1, -1.5, -2.25)$, then a second round at the same severity $\theta_{\pi_Q} = \frac{1}{2}$ autodeleverages all negative-equity accounts against the largest winner, yielding final equities $e_T = (3.75, 0, 1, 0, 0)$. Under the pro-rata rule (20) with the same severity $\theta_{\pi_{PR}} = \frac{1}{2}$, we have $D_T = 15$ and $W_T = 16$, so the uniform haircut factor

is $\alpha = \theta_{\pi_{PR}} D_T / W_T = 15/32$ and each winning account's equity is scaled by $(1 - \alpha) = 17/32$, giving $e_T = (170/32, 85/32, 17/32, -3, -12)$. In particular, the ranking of positive-equity accounts is preserved under pro-rata, since all winners are multiplied by the same constant factor, illustrating its rank-preservation property.

2.5 ADL Trilemma

The preceding subsections defined trader equity, exchange solvency, insurance funds, and autodeleveraging (ADL) policies. In this subsection we introduce a high-level design principle that organizes the rest of the paper: an *ADL trilemma*. Informally, a perpetuals venue cannot simultaneously (i) keep insolvency and ADL breach events rare, (ii) protect solvent winners from large socialized losses, and (iii) extract maximal long-run exchange revenue from trading and funding activity. Any ADL design must pick (at most) two of these three goals.

Exchange revenue. To define the trilemma, we have to first formally describe what constitutes exchange revenue. Recall from §2.3 that, given parameters $\alpha, \beta, \eta \in [0, 1]$, the insurance fund IF_t evolves by

$$\text{IF}_{t+1} = \text{IF}_t + \alpha \sum_{j \in T_t} \tau_t(\Delta q_j) + \eta p_t V_t + \beta \sum_{i=1}^n |\gamma_t| p_t |q_i| - \min\{\text{IF}_t, D_t\},$$

where T_t is the multiset of liquidations at t , V_t is traded volume, γ_t is the funding rate, and D_t is the period bad debt. For notational convenience, define the per-period gross fee flows

$$\text{Fee}_t^{\text{liq}} = \sum_{j \in T_t} \tau_t(\Delta q_j) \quad \text{Fee}_t^{\text{trade}} = p_t V_t \quad \text{Fee}_t^{\text{fund}} = \sum_{i=1}^n |\gamma_t| p_t |q_i|.$$

Most exchanges fully collect liquidation and trading fees and potentially collect fees on funding rates.¹⁰

Let $\zeta \leq (1 - \beta)$ be the expected fraction of funding rates that the exchange keeps as revenue. Then the exchange's *gross revenue* at time t is

$$\text{Revenue}_t^{\text{gross}} = \text{Fee}_t^{\text{liq}} + \text{Fee}_t^{\text{trade}} + \zeta \text{Fee}_t^{\text{fund}}.$$

By construction, the fractions α, η, β of these fee flows are diverted into the insurance fund. The remaining share accrues as *net exchange revenue*

$$\text{Revenue}_t := (1 - \alpha) \text{Fee}_t^{\text{liq}} + (1 - \eta) \text{Fee}_t^{\text{trade}} + (1 - \beta - \zeta) \text{Fee}_t^{\text{fund}}.$$

We note that in §8 that we refine this notion of revenue to a discounted *exchange long-term value* (LTV). The LTV accounts for future expected fee flows and traders leaving the exchange as a function of the realized ADL haircuts.

Three competing desiderata. There are three main goals that an ADL policy aims to enforce for traders and exchanges: revenue, solvency, and fairness. We formalize these notions in the sequel:

- *Solvency.* (§4, §5) The exchange aims to ensure that $R_t = 0$ for most times t and that the total shortfall $\sum_t R_t$ is small relative to the expected insurance fund size. This desideratum is beneficial to both the exchange and traders.

¹⁰In particular, we note that HLP-like systems where fees are charged on quoting offsetting positions can be viewed as a form of exchange revenue [69]. Moreover, a number of perpetuals AMM DEXs collect protocol revenue from funding flows [59, 91].

- *Fairness and moral hazard.* (§6, §7) Traders want to know that if they face socialization, they will not be asked to absorb a portion of R_t that is too large relative to their notional exposure on the exchange. This desideratum is mainly beneficial to traders.
- *Exchange revenue and participation.* (§8) Heavy ADL on high-value winners can trigger traders exiting the exchange permanently, shrinking future fee flows. This desideratum is mainly beneficial to the exchange.

It is clear that α, β, η and the ADL policy π jointly control these three dimensions. But a natural question is what the trade-offs are between these three desiderata at different values of α, β, η and parametrizations of π . Raising α, β, η builds the reserve faster and decreases breaches (*i.e.*, times with $R_t > 0$), but diverts exchange revenue into the fund. Increasing ADL severity (or concentrating haircuts) makes breaches rarer and accelerates solvency, but worsens fairness and drives away the highest-value winning traders. Keeping severities (*i.e.*, total solvency resolved by ADL) small preserves fairness and participation, but leaves the exchange exposed to repeated shortfalls. The goal of this paper is to formalize these three statements via a trilemma:

Proposition 2.1 (Trilemma, Informal). *Fix a sequence of perpetuals exchanges $\mathcal{P}_n$ and static ADL policies π_n with insurance parameters (α, β, η) . Under the heavy-tailed shortfall assumptions of §5, no policy family (π_n) can simultaneously satisfy the following uniformly in n :*

- (S) **Solvency:** *Residual risk is controlled, *i.e.*, $\sum_t R_t(\pi_n) = O_p(1)$ and $\mathbb{P}[R_t(\pi_n) > 0] = O(1)$.*
- (F) **Fairness:** *Moral hazard is bounded, *i.e.*, $\omega_T^\pi/D_T^\pi = \Theta(1)$ and $\omega_T^\pi/\Delta_T^\pi = \Theta(1)$.*
- (R) **Revenue:** *Exchange revenue is preserved relative to a benchmark, *i.e.*, $LTV(\pi_n) \asymp LTV_{\text{benchmark}}$.*

Enforcing (S) requires sacrificing (F) (via concentrated haircuts) or (R) (via excessive insurance diversion). Conversely, preserving (F) and (R) necessitates frequent solvency breaches.

A formal statement with precise definitions and a complete proof of the proposition appears in Appendix J.

Proof sketch and roadmap. The remainder of the paper establishes the Trilemma by analyzing each desideratum in turn:

- *Solvency Ratios (§5, Appendix C).* We analyze the ratio ω_T^π/D_T^π under heavy-tailed assumptions. Theorem C.3 proves that for any budget-balanced static ADL with severity θ_n , the ratio scales as $b_n/(\theta_n n)$. Consequently, enforcing rare breaches (large θ_n) drives the ratio to zero, violating (F).
- *Fairness Models (§6–§7, Appendices D–F.2).* We formalize fairness in two regimes. First, for a single round (§6), we show that capped pro-rata is the unique policy satisfying Sybil resistance and monotonicity. Second, under external shocks (§7), we prove that leverage-weighted rules (RAP) minimize future shortfall but sacrifice the most systemically important winners, highlighting the tension between short-term robustness and long-term participation.
- *Revenue and Price of Anarchy (§8, Appendix G.3).* We model the revenue-solvency tradeoff as a Stackelberg game between the exchange and traders. We show that static severities yield an unbounded Price of Anarchy ($J_n(\pi) \sim \theta_n n/b_n$) compared to dynamic policies. This confirms that achieving (R) and (S) simultaneously requires dynamic intervention, as static rules cannot balance the trilemma asymptotically.

3 Risk and Fairness Preliminaries

To analyze how different ADL policies balance solvency and winner survival, we require metrics that capture both the *magnitude* of losses (risk) and their *distribution* across traders (fairness). Classical risk measures like Value-at-Risk (VaR) and Expected Shortfall (ES) quantify aggregate solvency but ignore how the burden is shared. To address this, we introduce fairness-aware metrics and distributional orderings that allow us to rank policies based on how they concentrate or spread losses among winners. We then provide with an brief presentation of extreme-value scaling that we use in the sequel, before concluding with a brief introduction to algorithmic fairness. The preliminaries in this section are meant to be incomplete and we refer the reader to the literature cited for more details.

3.1 Risk Metrics

3.1.1 Traditional Risk Metrics

We briefly review standard metrics used in finance and regulation [18, 28]. For a loss $X \geq 0$ and confidence level $\alpha \in (0, 1)$:

- *Value-at-Risk (VaR)*: The quantile $\text{VaR}_\alpha(X) = \inf\{x : \mathbf{Prob}(X > x) \leq \alpha\}$.
- *Expected Shortfall (ES)*: The average loss in the worst α -fraction of cases, $\text{ES}_\alpha(X) = \mathbf{E}[X \mid X \geq \text{VaR}_\alpha(X)]$.

In ADL, we apply these to the residual shortfall R_t to measure solvency risk.

3.1.2 ADL-Specific Efficiency Metrics

To capture the trade-off between solvency and trader welfare, we define two ratios that normalize winner survival by the scale of the default:

- *Profitability-to-Total-Solvency Ratio (PTSR)*:

$$\text{PTSR}_T(\pi) = \mathbf{E} \left[\frac{\omega_T^\pi}{D_T^\pi} \right].$$

This measures the survival of the top winner relative to the total socialized loss. A low PTSR indicates that the most profitable trader is being disproportionately cannibalized to cover deficits.

- *Profitability-to-Maximum Solvency Ratio (PMR)*:

$$\text{PMR}_T(\pi) = \mathbf{E} \left[\frac{\omega_T^\pi}{\Delta_T^\pi} \right].$$

This compares the top winner’s equity to the largest single loser’s shortfall, capturing the concentration of risk on both sides of the trade.

These ratios mirror the VaR/ES distinction: PTSR captures aggregate efficiency, while PMR captures tail concentration. We formalize this connection and derive exact relationships in Appendix C.

3.2 Fairness and Distributional Comparisons

While scalar metrics like PTSR provide summary statistics, they cannot fully capture the fairness of a policy across the entire population of winners. Two policies might achieve similar solvency but distribute the pain very differently — one by wiping out a few large winners (Queue), another by shaving everyone slightly (Pro-Rata). To rank policies robustly, we use tools from majorization theory [81].

Schur-Convexity and Submajorization. We compare haircut vectors $h \in \mathbf{R}^n$ using *submajorization* ($\prec_w$). We say a policy π_A is *more fair* (or less concentrated) than π_B if its haircut vector is submajorized by π_B 's, i.e., $h(\pi_A) \prec_w h(\pi_B)$. A function $\phi : \mathbf{R}^n \rightarrow \mathbf{R}$ is said to be *Schur-convex* if $x \prec_w y$ implies $\phi(x) \leq \phi(y)$, i.e., it increases with concentration. This implies that for *any* convex, symmetric cost function ϕ (representing trader disutility), the aggregate distress is lower under π_A :

$$h(\pi_A) \prec_w h(\pi_B) \implies \sum \phi(h_i(\pi_A)) \leq \sum \phi(h_i(\pi_B)).$$

This gives us a powerful, parameter-free way to claim that Pro-Rata is “fairer” than Queue: it minimizes the collective pain for all convex risk attitudes. We provide a detailed treatment of these orderings and their application to ADL in Appendix F.2.

Comonotonicity. In our worst-case analysis, we often consider *comonotonic* couplings, where winners’ profits and losers’ deficits are perfectly correlated (move in lockstep). This represents the most dangerous regime for ADL, as large deficits coincide with large winner profits, testing the policy’s ability to extract liquidity without destroying the best traders. Our negative results in §5 exploit this structure to derive tight bounds on the Trilemma. See Appendix F.2 for further details on comonotonic risk bounds.

3.3 Extreme-value scaling

Our asymptotic analysis relies on *extreme-value scales*: deterministic sequences that characterize the typical magnitude of the largest winner equity and the largest loser shortfall in the limit of a large market. Recall the winner and loser index sets $\mathcal{W}_T$ and $\mathcal{L}_T$ with cardinalities k_n and m_n , respectively. Under mild mixing we assume the aggregate winner and loser masses concentrate at linear scales:

$$\left(\frac{W_T(\mathcal{P}_n)}{n}, \frac{D_T(\mathcal{P}_n)}{n} \right) \xrightarrow{p} (\mu_+, \mu_-). \quad (24)$$

Here $\mu_+, \mu_- > 0$ summarize the average winner and loser magnitudes. We define the respective maxima as

$$(\omega_T, \Delta_T) = \left(\max_{i \in \mathcal{W}_T} (e_T(\mathbf{p}_{i,T}))_+, \max_{i \in \mathcal{L}_T} (-e_T(\mathbf{p}_{i,T}))_+ \right). \quad (25)$$

A pair of deterministic, increasing sequences $\{b_k^+\}_{k \geq 1}$ and $\{b_m^-\}_{m \geq 1}$ constitute *extreme-value scales* if there exist constants $c_+, c_- \in (0, \infty)$ such that, as $n \rightarrow \infty$,

$$\left(\frac{\omega_T}{b_{k_n}^+}, \frac{\Delta_T}{b_{m_n}^-} \right) \xrightarrow{p} (c_+, c_-). \quad (26)$$

We abbreviate $b_n = b_{k_n}^+$ when the context (winner side) is unambiguous. A canonical choice, consistent with classical extreme-value theory, is the set of *upper-quantile normalizers*:

$$(b_k^+, b_m^-) = \left(F_+^{-1} \left(1 - \frac{1}{k} \right), F_-^{-1} \left(1 - \frac{1}{m} \right) \right). \quad (27)$$

where F_+ and F_- denote the distribution functions of winner-side positive equities and loser-side shortfalls, respectively.

Extreme-value severity scales. We refer to $\theta_n = \Theta(b_n/n)$ as the *extreme-value severity scale*. This choice aligns the haircut magnitude with the typical largest winner equity $b_n = b_{k_n}^+$ in a large market where $k_n = \Theta(n)$.

Examples. Two prototypical cases appear throughout our analysis:

- *Light tails (sub-Gaussian).* For distributions with sub-Gaussian scales σ_+, σ_- ,

$$b_k^+ \asymp \sigma_+ \sqrt{2 \log k}, \quad b_m^- \asymp \sigma_- \sqrt{2 \log m}.$$

- *Power-law tails (Pareto/Fréchet).* For distributions satisfying $\mathbf{Prob}\{(e_T)_+ > x\} \sim C_+ x^{-\alpha_+}$ and $\mathbf{Prob}\{-(e_T)_+ > x\} \sim C_- x^{-\alpha_-}$ with $\alpha_{\pm} > 0$,

$$b_k^+ \asymp (C_+ k)^{1/\alpha_+}, \quad b_m^- \asymp (C_- m)^{1/\alpha_-}.$$

3.4 Fairness

Algorithmic Fairness in ADL. Algorithmic fairness literature distinguishes between *axiomatic* approaches, which posit structural invariants (often yielding impossibility results [11, 73]), and *optimization* approaches, which minimize disparate impact via convex programs [50]. ADL spans both: queue-based policies resemble discrete ranking mechanisms subject to Arrow-style impossibilities, while pro-rata policies admit continuous convex formulations. We adopt an axiomatic perspective to characterize “ideal” fairness, then relax it to optimize for robustness.

Axiomatic Properties. We formalize three operational desiderata for fair ADL policies:

- *Sybil Resistance:* Splitting an account into multiple smaller positions (or merging them) should not alter the aggregate haircut. This ensures outcomes depend on total exposure, not identity or fragmentation. Queue-based rules violate this, as splitting a large winner can push parts of its equity lower in the priority stack.
- *Monotonicity (Stable Ordering):* The policy should preserve the relative ordering of winners. If trader A has more equity than trader B pre-ADL, A should also have more equity post-ADL. This prevents arbitrary re-ranking of survivors.
- *Scale Invariance:* If all deficits and equities scale by $\lambda > 0$, the relative haircuts should remain unchanged. This ensures the mechanism responds to *leverage* and *risk distribution*, not nominal sizes.

In §6, we prove that *capped pro-rata* is the unique policy satisfying all three properties, whereas standard queueing fails Sybil resistance and monotonicity.

3.5 Moral hazard

Moral hazard is a generic incentive problem that arises when one party (the *agent*) can take actions that affect risk or payoffs, but those actions are not perfectly observable, verifiable, or contractible by the party bearing some of the downside (the *principal*); see, *e.g.* Holmström [67], Laffont and

Martimort [77]. Classic examples include workers choosing effort under wage contracts, managers taking portfolio risk on behalf of investors, or insurers providing coverage to policyholders whose behavior may change once insured. Because the agent’s action is hidden, the principal cannot simultaneously provide full insurance (so that the agent’s payoff is insulated from shocks) and preserve strong incentives for the agent to take socially desirable actions: risk sharing necessarily distorts effort incentives. The central goal of the moral-hazard literature is therefore to characterize *second-best* contracts (typically linear or otherwise simple sharing rules) that optimally trade off incentive provision against insurance under information constraints [31, 49].

Moral hazard also couples to robustness: contracts must work across environments and counterparties, pushing outcomes away from the first-best benchmark [49]. The same friction arises whenever losses are socialized (*e.g.* insurance pools or centrally cleared derivatives) as risk sharing dulls the ability for individuals to manage risk efficiently. ADL fits this template: winners and losers supply mutual insurance while the exchange chooses the loss-sharing rule. The negative results in Section 5 formalize the resulting second-best trade-off.

4 Severity Optimization

We decompose the ADL problem into two orthogonal components: *severity* (how much to socialize) and *haircuts* (who pays). Solvency is driven almost entirely by the scalar severity sequence (θ_t) , while fairness and revenue depend on the haircut distribution h_t . Since §6–§7 cover haircut design, here we focus on optimizing θ_t to balance solvency against the risk of driving away traders.

Severity Policies. A *causal severity policy* maps available information $\mathcal{F}_t$ (deficits, funding rates, risk limits) to a severity fraction $\theta_t \in [0, 1]$. We contrast three approaches:

1. *Static Interpolation:* θ_t is fixed or linearly moves between deficit-matching $\theta_t^{\text{match}} = \min\{1, \sum_i e_{t,i}/D_t\}$ (so $H_t = B_t^*$ when feasible) and a cap $\bar{\theta}$ (*e.g.* 1).
2. *Exponential Backoff:* $\theta_t = \theta_0 \alpha^{k_t}$ decays as the number of recent shocks k_t increases. This prevents cascading failures from wiping out the entire book during prolonged stress.
3. *Online Control (Mirror Descent):* We treat severity selection as an online convex optimization problem. The controller adjusts θ_t to minimize a regret bound combining solvency costs (R_t) and revenue loss (excessive haircuts). This approach adapts to non-stationary market conditions and respects caps Θ_t derived from the Price of Anarchy analysis in §8.

Separation Principle. Under the budget-balance constraint (17), any one-round ADL policy π with severity θ_π and haircuts h_π satisfies

$$H(\pi) = \sum_i h_{\pi,i} e_T(\mathbf{p}_i)_+ = \theta_\pi D_T(\mathcal{P}_n),$$

so the post-ADL deficit is

$$R_T(\pi) = (D_T(\mathcal{P}_n) - H(\pi))_+ = (1 - \theta_\pi)_+ D_T(\mathcal{P}_n),$$

which depends only on θ_π and $D_T(\mathcal{P}_n)$, not on how $H(\pi)$ is distributed across winners. Thus, any solvency functional that depends on π only via $\{R_t(\pi)\}_t$ (*e.g.* breach frequencies, VaR/ES of $\sum_t R_t$) can be optimized over the scalar sequence (θ_t) , holding the haircut rule fixed. Conversely, fairness and revenue functionals depend on the winner-side post-ADL equities $\{(1 - h_{\pi,i})e_T(\mathbf{p}_i)_+\}_i$ and are invariant to the choice of θ_t once the total budget H_t is fixed. We exploit this separability in our numerics (§9) by tuning θ_t (severity control) and h_t independently.

5 Negative Results

We begin with two structural limits for ADL. They formalize a moral-hazard trade-off between solvency and winner payoffs that cannot be removed by better policy design. The tension is instantaneous and zero-sum at the shock: covering the aggregate deficit D_T eventually requires shaving the profits of the best winner. We first argue that this tension can be viewed as a form of moral hazard, in the sense used in principal-agent contracts.

An exchange that earns fees from trading volume has an incentive to let very risky, highly leveraged users accumulate large positions. When these users default, their negative equity shows up as a large deficit D_T , which is then socialized onto the winners via ADL haircuts. In our model this is the basic moral-hazard question: the exchange enjoys fee revenue from risky losers, while the cost of their tail losses is borne by solvent winners. The results below quantify how this wedge scales as the market thickens, showing that the severity needed for solvency grows with the number of positions and with the leverage imbalance between losers and winners.

5.1 Impossibility of Avoiding Moral Hazard as $n \rightarrow \infty$

We first show that the tension between solvency and winner survival scales unfavorably with market size, creating an unavoidable moral hazard wedge. Viewing traders as principals and the exchange as an agent, preserving the agent's utility (solvency) asymptotically destroys the principals' utility (winner profits) unless severity vanishes.

Example. Consider a book of size n where loser positions are i.i.d. with mean loss $\mu > 0$, so aggregate deficit scales as $D_T \approx \mu n$. Suppose the top winner's equity scales sub-linearly, e.g. $b_n \sim c \log n$ (light tails). If the exchange commits to a fixed severity $\bar{\theta} > 0$ to maintain solvency, the total haircut budget is $H_n \approx \bar{\theta} \mu n$. As $n \rightarrow \infty$, this linear haircut cost overwhelms the logarithmic winner equity:

$$e_n^{\text{post}} \approx (c \log n - \bar{\theta} \mu n)_+ \rightarrow 0.$$

Thus, to preserve any winner profit, severity must vanish at the rate $\theta_n = O(b_n/n)$. Constant severity policies therefore represent a form of growing moral hazard: they transfer an increasingly large share of tail risk from the exchange to the most profitable traders. We formalize this scaling in Appendix C.6.

Main Result. In our model, socialization is instantaneously zero-sum: every dollar of deficit D_T covered by ADL must come from the winners' equity. If ADL usage does not vanish as the market grows larger ($n \rightarrow \infty$), then one cannot simultaneously keep deficits small (solvency, low D_T) and preserve winner payoffs (high ω_T). The proposition below quantifies this tradeoff in terms of extreme-value scales and can be read as a moral-hazard impossibility result in the principal-agent sense.

Proposition 5.1 (Informal). *Suppose that the extreme value scale of the winning trader's equity, $\omega_T^\pi(\mathcal{P}_n)$, is b_n . Then,*

$$\text{PTSR}_T(\mathcal{P}_n, \pi) = \mathbf{E}_\pi \left[\frac{\omega_T^\pi(\mathcal{P}_n)}{D_T^\pi(\mathcal{P}_n)} \right] \asymp \frac{b_n}{\theta_n n}$$

In particular, the example above shows the order b_n/n is tight: unless the severity vanishes at the extreme-value scale, $\theta_n = \Theta(b_n/n)$, one cannot preserve the best trader's profits as $n \rightarrow \infty$.

Example. The bound b_n/n dictates the maximal safe severity for different distributions:

- *Gaussian Winners (Light Tails):* For $e_i \sim \mathcal{N}(0, 1)$, the maximum scales as $b_n \sim \sqrt{2 \log n}$. Thus, severity must vanish rapidly, $\theta_n \lesssim \sqrt{2 \log n}/n$, to avoid autodeleveraging the top winner. This implies that the shortfall that can be covered goes to zero as $n \rightarrow \infty$.
- *Pareto Winners (Heavy Tails):* For $e_i \sim \text{Pareto}(\alpha)$ with $\alpha > 1$, the maximum scales as $b_n \sim n^{1/\alpha}$. Here, the constraint is looser but still vanishing as $n \rightarrow \infty$: $\theta_n \lesssim n^{-(1-1/\alpha)}$.

Even in the heavy-tailed case (infinite variance, $\alpha < 2$), budget balance implies that a constant fraction severity θ wipes out the top winner almost surely as $n \rightarrow \infty$.

5.2 Excessive leverage guarantees large maximal trader loss

Another natural form of moral hazard, one that quantifies the role of the exchange as an agent, is the ratio of leverage held by the winning and losing sides. If losers generally have much more leverage than winners, then the exchange is giving winning traders a higher likelihood of paying socialized losses. We quantify this deficiency in risk management by the exchange by providing a more quantitative extreme value result: we show that the ratio of winner to loser leverages is a multiplicative term in how fast the winning trader's position is likely to be autodeleveraged.

Example. Consider a regime where loser leverage mass ℓ^- dominates winner leverage mass ℓ^+ . Let b_n^+ (resp. b_n^-) denote the extreme value (EV) scales of winner and loser equities, respectively. Aggregate shortfall scales with loser liability, $D_T \approx \ell^- b_n^-$, while the top winner's equity scales as $\omega_T \approx \ell^+ b_n^+$. If severity θ_n is fixed, the haircut budget $H_n = \theta_n D_T \approx \theta_n \ell^- b_n^-$ is large. Comparing budget to winner equity:

$$\frac{H_n}{\omega_T} \approx \theta_n \frac{\ell^- b_n^-}{\ell^+ b_n^+}.$$

If $\ell^- \gg \ell^+$, even a small severity θ_n creates a haircut larger than the top winner's entire position. Thus, safe severity is throttled not just by market size but by the *leverage imbalance* ℓ^+/ℓ^- . See Appendix C.6 for the rigorous construction.

Main Result. The next result makes the leverage–imbalance effect precise. When loser positions carry much more leverage than winners, the ratio of winner to loser leverage masses ℓ_n^+/ℓ_n^- enters multiplicatively in the EV-scale bound for the top winner's survival. In moral-hazard terms, an exchange that tolerates very high loser leverage effectively exposes its best winners to autodeleveraging: even with small severities θ_n , a large imbalance ℓ_n^-/ℓ_n^+ can still force substantial haircuts on the most profitable traders.

Proposition 5.2 (Informal). *Let b_n^+, b_n^- be the extreme value scales of the winning and losing traders, respectively. Furthermore, let the winner and loser side leverage masses be ℓ_n^+, ℓ_n^- , respectively. Then we have the general EVT scaling:*

$$\mathbf{E}_\pi \left[\frac{\omega_T^\pi(\mathcal{P}_n)}{D_T^\pi(\mathcal{P}_n)} \right] \asymp \frac{\ell_n^+}{\ell_n^-} \cdot \frac{b_n^+}{b_n^-} \cdot \frac{1}{\theta_n}.$$

The full proof is in Appendix C.

Example. The scaling is now modulated by the *leverage imbalance* $\Lambda_n = \ell_n^+ / \ell_n^-$.

- *Symmetric Leverage* ($\Lambda_n \asymp 1$): If winner and loser leverages are comparable, the bound reverts to the previous case (b_n^+ / b_n^-).
- *High Loser Leverage* ($\Lambda_n \ll 1$): If losers are $10\times$ more leveraged, the safe severity drops by $10\times$. In the extreme, if $\ell_n^- \sim n$ (e.g. a few massive whales) while $\ell_n^+ \sim 1$ (retail winners), the required severity θ_n becomes vanishingly small even faster than the EV scale suggests.
- *High Winner Leverage* ($\Lambda_n \gg 1$): Conversely, if winners are highly leveraged (so their equity is sensitive to small price moves, but large in notional), the ratio Λ_n buffers the impact, allowing for larger severities without autodeleveraging the top winner.

Moral Hazard Interpretation. These results formalize the “second-best” nature of ADL: perfect insurance without incentive distortion is impossible. Tail shocks necessarily haircut the top winners (Proposition 5.1), and high loser leverage magnifies that hit (Proposition 5.2). This is the same trade-off seen in clearinghouses: Variation Gains Haircutting trims winner gains to keep the CCP solvent [44, 60]. ADL is the perpetual futures analogue of Variation Gain Haircutting (i.e., an ex-post haircut on unrealized profits), so the same moral-hazard tension applies [48, 93]. PTSR/PMR play the same role as VGH fairness/solvency gauges in CCP recovery [44, 48, 60].

5.3 Queue-based Methods are the worst ADL policies for the top winning trade

One result that we see empirically in §9 is that the Queue mechanism is by far the worst ADL policy for the top winning trade. We briefly formalize this with an informal proposition that we prove in Appendix C.3.

Proposition 5.3 (Queue Maximizes Damage to Top Winner). *Among all valid ADL policies with fixed severity θ , the Queue policy (haircut largest winners first) uniquely minimizes the survival of the top winner. Let ω_T^{Queue} and ω_T^{PR} be the top winner’s post-ADL equity under Queue and Pro-Rata, respectively. Whenever the haircut budget $H = \theta D_T$ satisfies $H \leq e_{(1)}$, we have the strict gap*

$$\omega_T^{\text{PR}} - \omega_T^{\text{Queue}} = H \left(1 - \frac{e_{(1)}}{W_T} \right) > 0.$$

Consequently, Pro-Rata strictly dominates Queue in terms of fairness metrics:

$$\text{PTSR}_T(\text{PR}) - \text{PTSR}_T(\text{Queue}) = \theta \left(1 - \frac{e_{(1)}}{W_T} \right).$$

Thus, Queue maximizes the concentration of haircuts on the most profitable trader, leading to the worst possible PTSR/PMR scores.

The full proof is in Appendix C.3. Appendix C.5 further strengthens this statement by proving that Queue also minimizes the VaR and ES of the top survivor at every confidence level, so even classical risk metrics rank Queue as the worst policy for the top winner.

6 Fairness

A natural question is whether a particular ADL policy is optimal or fair. We demonstrate that a modified pro-rata rule respecting per-account constraints, the *capped pro-rata rule*, is the most fair policy in two senses: it minimizes total convex disutility and is the unique rule satisfying key axiomatic properties.

Definition of the Capped Pro-Rata Rule. The capped pro-rata rule is a water-filling algorithm that enforces per-account constraints before applying a pro-rata haircut. Combining haircut constraints (22) and equity constraints (23), we define an effective cap $\beta_i = \min\{\bar{h}_i, 1 - \underline{e}_i/e_T(\mathbf{p}_i)\}$. This β_i is the most we can haircut account i while honoring both its per-account haircut limit and its minimum post-ADL equity. The maximum haircut capacity is

$$C(\beta) = \sum_{i=1}^n e_T(\mathbf{p}_i) \beta_i. \quad (28)$$

If the deficit exceeds this capacity, no feasible haircut exists. Otherwise, we aim to equalize haircuts subject to these caps. The capped pro-rata rule π_{CP} is defined as:

$$h_{\pi_{CP},i} = \min\{\eta, \beta_i\} \quad (29)$$

where $\eta \in [0, 1]$ is chosen to satisfy the budget constraint

$$\sum_{i=1}^n h_{\pi_{CP},i} e_T(\mathbf{p}_i)_+ = \theta_\pi D_T(\mathcal{P}_n). \quad (30)$$

This can be computed efficiently by sorting the caps β_i and finding the threshold η (see eq. (30) and the algorithm in Appendix E).

Convex Optimality. Let $\phi : [0, 1] \rightarrow \mathbf{R}$ be an increasing convex function representing the disutility of a haircut. We formulate the problem of minimizing total disutility subject to per-account constraints:

$$\begin{aligned} \min_{h \in [0,1]^n} \quad & \sum_{i=1}^n \phi(h_i) \\ \text{subject to} \quad & \forall i \in [n] \quad h_i \leq \beta_i \quad (\text{effective cap}) \end{aligned} \quad (31)$$

In Appendix D, we show:

Proposition 6.1. *The unique solution to (31) is the capped pro-rata rule (29).*

This implies that for any reasonable model of financial disutility, capped pro-rata maximizes social welfare.

Axiomatic Fairness. We define three natural fairness properties:

- *Sybil resistance:* Outcomes are identity-agnostic and invariant to splitting or merging accounts.
- *Scale invariance:* Haircuts are unchanged if both shortfall and equities are scaled by a common factor.
- *Monotonicity (Stability):* The policy preserves the relative ordering of post-haircut equity. If $e_i \geq e_j$, then $(1 - h_i)e_i \geq (1 - h_j)e_j$.

In Appendix D, we formalize the following result:

Proposition 6.2 (Informal). *Under mild smoothness conditions, the unique sybil-resistant, scale-invariant, and monotone ADL policy satisfying per-account constraints is capped pro-rata.*

This establishes capped pro-rata as the unique fair ordering rule. In contrast, queue-based rules (ranking by leverage or PNL) are generally not stable or sybil-resistant. For example, a trader can split one large winning account into many small ones and push each down the queue, reducing any single haircut while keeping the same aggregate exposure. RAP introduces a risk tilt, deviating from strict equity stability to reduce risk, but remains “as stable as possible” within its risk-weighted framework. In the numerical examples we present in the next section, we will see explicit examples where capped-pro rata is better for both exchange solvency and trader profitability than the Queue and Pro-Rata.

7 Risk-aware Policies (RAP)

Fairness-focused rules like capped pro-rata treat all winning dollars equally, ignoring the heterogeneous risks imposed by different winners. Higher effective leverage implies thinner liquidation windows and greater sensitivity to execution costs [7, 57]. We introduce *Risk-Aware Pro-Rata* (RAP) rules, a family of rules that preserves the fairness structure of capped pro-rata while tilting haircuts toward higher-risk positions. This family includes Levered Pro-Rata (LPR) when the risk weighting is linear in leverage. We derive RAP from a robustness criterion that minimizes a one-step excess-deficit proxy.

One-step Next Deficit. To formalize robustness, we consider the *next-step deficit* D_{T+1}^{next} resulting from a price shock Z_T occurring immediately after the ADL procedure. Let π be an ADL policy yielding post-haircut notional $\tilde{n}_{T,i}$ and equity $e_{T,i}$. We model the shock Z_T using a Markov kernel sensitive to the state. The next-step deficit is the negative equity after the shock:

$$D_{T+1}^{\text{next}} = \sum_{i \in \mathcal{W}_T} (-e_{T,i} - \tilde{n}_{T,i} Z_{T,i})_+ = \sum_{i \in \mathcal{W}_T} \tilde{n}_{T,i} \left(\frac{1}{\lambda_{T,i}} + Z_{T,i} \right)_-.$$

Our objective is to minimize the conditional expected deficit $\delta_T = \mathbf{E}[D_{T+1}^{\text{next}} \mid \mathcal{F}_T]$. Defining $\psi_i(u) = \mathbf{E}[(u + Z_{T,i})_-]$, we have:

$$\delta_T = \sum_{i \in \mathcal{W}_T} \tilde{n}_{T,i} \psi_i \left(\frac{1}{\lambda_{T,i}} \right). \quad (32)$$

Appendix F.1 illustrates how pro-rata fails to control this deficit when shocks scale with leverage.

Risk Models. We assume the exchange employs a convex risk model $g(\lambda)$ characterizing the risk contribution of leverage λ . Common choices include:

1. *Linear*: $g(\lambda) = \lambda$, corresponding to LPR. Matches VaR/CVaR scaling under linear shocks.
2. *Power*: $g(\lambda) = \lambda^c$ ($c > 1$), penalizing high leverage super-linearly.
3. *CVaR*: $g(\lambda) = (\lambda - \tau)_+$, a surrogate for Conditional Value-at-Risk.

Appendix F.1 provides numerical examples of these weightings.

7.1 Risk-Aware Pro-Rata (RAP) Rule

The RAP rule allocates haircuts using weights $w_i = \lambda_{T,i} g(\lambda_{T,i})$. With per-account caps β_i , the haircuts are:

$$h_{\text{RAP}(g),i} = \min\{\beta_i, \tau w_i\}, \quad (33)$$

where τ is calibrated to meet the target budget. This concentrates haircuts on high-risk positions. See Appendix E for precise implementation of the algorithm.

7.2 Convex Dominance of RAP

We can construct an optimal risk model g^* that minimizes δ_T . Intuitively, an optimal risk model should maximize the marginal reduction of δ_T per unit of equity $e_{T,i}$. Differentiating δ_T with respect to the haircut $h_{T,i}$ reveals:

$$-\frac{\partial \delta_T(h)}{\partial h_{T,i}} = n_{T,i} \psi_i\left(\frac{1}{\lambda_{T,i}}\right) = e_{T,i} \lambda_{T,i} \psi_i\left(\frac{1}{\lambda_{T,i}}\right) = e_{T,i} \rho_i(\lambda_{T,i}), \quad (34)$$

where we define the *perspective transform* $\rho_i(\lambda) = \lambda \psi_i(1/\lambda)$. The greedy rule that maximizes the marginal reduction per equity-dollar is exactly the rule that prioritizes winners with the largest $\rho_i(\lambda_{T,i})$.

Perspective Transform and Economic Intuition. The function ρ_i is the perspective transform of the convex function ψ_i [27, §2.3.3], widely used to induce 1-homogeneity. Economically, $u = 1/\lambda$ is the bankruptcy buffer, which is the percentage price shock required to wipe out the position's equity (since loss equals equity when shock magnitude $Z = e/\tilde{n} = 1/\lambda$). Using convex duality, we can write $\rho_i(\lambda) = \sup_{y \geq 0} \{y - \lambda \psi_i^*(y)\}$, where y is a ‘yield’ (equity change per unit buffer change) and $\psi_i^*(y)$ is the cost to insure a position with yield y . Thus, maximizing ρ_i is equivalent to maximizing the net solvency gain per unit of equity. For example, if a trader with leverage $\lambda = 10$ faces an insurance cost of 0.5% for a yield of 8%, the net solvency gain is $0.08 - 10(0.005) = 0.03$ per unit equity.

Constructing the optimal risk model. This suggests the optimal weights $w_i^* = \rho_i(\lambda_{T,i})$ and risk model $g^*(\lambda) = \rho(\lambda)/\lambda$.

Proposition 7.1 (Informal). *For any weighted pro-rata rule comonotone with w^* , the residual risk vector is weakly submajorized by that of $\text{RAP}(g^*)$. This implies that $\text{RAP}(g^*)$ minimizes the total risk for any convex risk measure, assuming the shock process follows the kernel K .*

The proof is in Appendix F.2. This result implies that if we have prior knowledge about the kernel K (e.g. from historical data or backtesting), then constructing g^* via the perspective transform yields the optimally robust and fair haircut model.

7.3 Limitations of One-Step Optimality

While RAP minimizes the one-step next deficit, it may not be optimal under multiple correlated shocks or broader exchange objectives.

Correlated Shocks. Consider sequential shocks Z_{t+1}, Z_{t+2} with positive correlation. High-leverage positions might act as a hedge against future shocks. In such cases, aggressively liquidating high-leverage winners (as in RAP) can increase the two-step deficit compared to pro-rata. Appendix F.1 details an AR(1) shock example where the sum of deficits is lower for pro-rata.

Exchange Incentive Compatibility. Exchanges must also consider *exchange long-term value (LTV)* — the future fees and liquidity provided by traders. RAP targets high-leverage, often high-volume, traders. Fully liquidating these accounts to minimize immediate risk can reduce the exchange’s long-term utility. Appendix F.3 provides an example where preserving a high-value trader (via pro-rata) yields higher total utility than the risk-minimizing RAP.

8 Multiple-Round ADL as a Stackelberg Control Problem

As the previous section illustrates, risk-aware pro-rata mechanisms work well as robust, static strategies but not in a dynamic setting. Both future exchange revenue and correlated shocks require including a model of the dynamics of ADL over multiple rounds. We use the intuition gained from RAP mechanisms for a single round to help extend the model to the multiple round setting. A common place where there is a difference between one-shot equilibria and multi-round equilibria is in the difference between Nash and Stackelberg games. We will illustrate how the multi-round ADL problem is naturally a Stackelberg game whereas the ADL policies that we have studied so far are more like Nash optimal strategies. To do this, we will first briefly describe Stackelberg games and then provide examples demonstrating that it is the right formalism for this extension.

8.1 ADL is a Stackelberg Game

Definition. A Stackelberg (leader–follower) game is sequential: the leader first *commits* to a policy, the follower observes it, and then best-responds; the leader chooses its commitment anticipating that response [17, 56, 113].¹¹ Commitment here means the policy is observable and credible to followers. The move order is therefore leader first, follower second, and the leader optimizes against the follower’s best-response set. Over many rounds, the leader can *learn* a near-optimal commitment by updating $\{\theta_t, h_t\}$ using feedback from realized follower responses and shocks, recovering standard online-learning guarantees when losses are convex (or have convex surrogates) [42, 65, 100, 114]. The Stackelberg view also supports robust design: the exchange can pick policies that hedge against worst-case follower responses and price/oracle shocks, steering the system toward states with smaller tail externalities [5, 110].

ADL and Stackelberg games. ADL naturally fits the setup of a Stackelberg game: the exchange (leader) *publishes* its ADL policy each round by choosing a severity budget $\theta_t \in [0, 1]$ and a haircut rule. As we saw in §2.4, the queueing rules used by Binance and Hyperliquid imply both a static severity and haircut rule and are published on their websites (which is their form of commitment). Market participants (*e.g.* traders, LPs) and the stochastic environment (order flow, depth, liquidation behavior, oracle and price shocks) then respond as followers given the announced policy. After market participants (followers) respond, the residual exposure (*i.e.*, the remaining negative equity) is measured and the leader decides on a next action (which could include to stop the game and not perform an ADL policy). The commitment of the exchange allows followers to construct a *local* best response that can adapt as the number of rounds of the game increases.

Single vs. Multi-Round Models. One-shot (static) ADL treats each deficit in isolation, optimizing a fixed severity θ and haircut rule h for that round. The multi-round Stackelberg model links rounds: the leader commits to a policy path $\{\theta_t, h_t\}$, and followers adapt, creating state dependence

¹¹In finance and control, this appears as “dynamic optimal control with commitment” or a principal–agent model with rational expectations [19, 77, 94].

via churn, liquidity, and leverage. This coupling makes commitment valuable: a policy that is optimal in one round can be dominated once cross-shock feedback is considered.

Stackelberg vs. Nash. Nash equilibria are one-shot, *simultaneous-move* fixed points where each player best-responds to the other’s action but no one commits *ex ante* [56]. By contrast, Stackelberg equilibria embed the leader’s commitment advantage, often producing different (and more robust or revenue-enhancing) outcomes [17, 113]. Static ADL strategies such as venue queues (Binance/Hyperliquid) or single-round RAP can be analyzed under a Nash lens, as we assume there is no further response by a leader after a single execution of an ADL policy. In practice, ADL controllers are used repeatedly (such as on October 10, 2025; see §9) which more resembles a Stackelberg game. In Appendix G.3 we give an explicit example for ADL where a Stackelberg equilibrium differs from a Nash outcome and in which the Stackelberg equilibrium strictly improves exchange revenue and robustness.

8.2 Why do we need a dynamic model?

Stackelberg games are inherently dynamic due to the alternating move order between the leader and follower. A natural question is whether we really need to use a dynamic method or if a static method is ‘good enough.’ In this section, we will provide our first formal results showing that dynamic ADL rules are necessary if you want to optimize both time to recover solvency and exchange revenue after repeated shocks. This will motivate the construction of the dynamic ADL policy in the sequel.

Exchange Long-Term Value (LTV) and the Trilemma. Recall the ADL Trilemma from §2.5, which posits a conflict between Solvency (**S**), Fairness (**F**), and Revenue (**R**). To quantify the *Revenue* objective, we define the *Exchange Long-Term Value* (LTV). Unlike immediate fee capture, LTV accounts for the long-run impact of ADL on trader participation: heavy haircuts on winners cause them to exit or reduce activity (“churn”), eroding future fees.

For a policy π , let $h_{t,i} \in [0, 1]$ be the haircut fraction on winner i at time t . We model trader retention via a revenue retention curve $r_{t,i}(h_{t,i})$ which is decreasing in the haircut size (larger haircuts drive away more volume). The per-period expected revenue from trader i is then $R_{t,i} = e_{t,i} r_{t,i}(h_{t,i})$, where $e_{t,i}$ is their equity. The LTV is the discounted sum of these revenues:

$$\text{LTV}_T(\pi) = \sum_{t=0}^T \beta^t \sum_{i \in W_t} e_{t,i} r_{t,i}(h_{t,i}).$$

Maximizing LTV corresponds directly to the (*R*) vertex of the trilemma.

Competing Time Scales. To compare policies dynamically, we track two opposing clocks starting from a default event at $\tau_{\text{def}}: \tau_{\text{solv}}, \tau_{\text{rev}}$. Intuitively, τ_{solv} is how quickly deficits are cleared whereas τ_{rev} is how quickly fee revenue rebounds.

- *Solvency Recovery Time* (τ_{solv}): The time required for the insurance fund (or deficit) to fully recover to a safe level δ .
- *Revenue Recovery Time* (τ_{rev}): The time required for the exchange’s expected LTV to return to pre-shock levels $(1 - \epsilon)\text{LTV}_{\text{pre}}$.

Formally,

$$\tau_{\text{solv}} = \inf\{t \geq \tau_{\text{def}} : \mathbf{IF}_t \geq \delta\}, \quad \tau_{\text{rev}} = \inf\{t \geq \tau_{\text{def}} : \mathbf{E}[\text{LTV}_t] \geq (1 - \epsilon)\text{LTV}_{\text{pre}}\}.$$

As we will show, static policies trade these quantities off one another off. We also show that the concentration of how haircuts are distributed amongst winning traders controls this trade-off.

Example. We can rank policies by how concentrated their haircuts are using the Schur order (§3):

- *Queue (concentrated):* Haircuts hit the largest winners first, yielding a highly unequal (Schur-convex) vector. Deficits clear quickly (small τ_{solv}), but the biggest fee payers churn, so LTV recovers slowly (large τ_{rev}).
- *Pro-Rata (diffuse):* Haircuts are spread evenly (Schur-concave). Solvency takes longer because shallow accounts contribute little (larger τ_{solv}), but everyone stays engaged, so revenue rebounds faster (small τ_{rev}).

Any intermediate policy traces this frontier: more concentration shortens τ_{solv} but lengthens τ_{rev} . This intuition is formalized in the following proposition, which states that any policy that concentrates haircuts (like Queue) will recover solvency faster but revenue slower than a diffuse policy (like Pro-Rata).

Opposing Schur orderings. Let $z_t(\pi)$ be the vector of post-haircut residual equities. If policy A is more concentrated than policy B (i.e., $z_t(A) \succ_{\text{Schur}} z_t(B)$), then A recovers solvency faster but revenue slower.

Proposition 8.1 (Informal). *Under separable convex stage losses and monotone shock dynamics, if the residual exposure of policy A majorizes policy B (e.g. Queue vs. Pro-Rata), then:*

$$\tau_{\text{solv}}(A) \leq \tau_{\text{solv}}(B) \quad \text{but} \quad \text{LTV}_t(A) \leq \text{LTV}_t(B) \quad \text{for all } t.$$

This implies no static policy can simultaneously minimize both recovery times.

8.3 Incentive Compatibility and Constraints

We frame the problem as a Stackelberg game: the exchange (leader) commits to a policy, and traders (followers) best-respond. We impose follower incentive compatibility (FIC) and leader-side revenue and solvency constraints (LIC) to ensure the policy is feasible for traders and sustainable for the venue.

Follower Incentive Compatibility (FIC). Winners must remain solvent and within haircut caps:

$$0 \leq h_{t,i} \leq \beta_{t,i}, \quad (1 - h_{t,i})e_{t,i}^+ \geq \underline{e}_{t,i},$$

where $\beta_{t,i}$ enforces per-account haircut limits and $\underline{e}_{t,i}$ enforces minimum post-ADL equity (cf. (22)–(23)). These bounds guarantee feasibility (no negative winners) and rule out haircut paths that would violate venue-level haircut caps (as we saw in §2.4).

Leader Incentive Compatibility (LIC). Revenue must stay above a floor and solvency above a safety threshold:

$$\mathbf{E}[\text{LTV}_t] \geq (1 - \epsilon) \text{LTV}_{\text{pre}}, \quad \text{IF}_t \geq \delta \quad \text{or} \quad \text{CVaR}_\alpha(-\text{IF}_t) \leq \kappa.$$

These ensure the policy is not revenue-destructive (per the churn model in §9) and keeps the exchange solvent both on-path and in the tails (§2.3). These constraints define the feasible set for any policy path $\{\theta_t, h_t\}$; we defer the full optimization program to Appendix G.3.

Stackelberg vs. Nash Equilibria. We model the interaction between the exchange and traders as a Stackelberg game where the exchange (leader) commits to a policy π_t , and traders (followers) subsequently respond (*e.g.* by adjusting liquidity or closing positions). This commitment power is crucial. As we show in Appendix G.3 (Proposition G.6), a simultaneous-move (Nash) game admits a “bad” equilibrium where traders withdraw liquidity anticipating high severity, which in turn forces the exchange to set high severity to cover the deficit, often violating the LIC floor on revenue. In the Stackelberg model, the exchange commits to a policy that satisfies FIC/LIC and rules out this inefficient equilibrium, coordinating the market toward a high-liquidity, lower-severity outcome that keeps both trader feasibility and venue sustainability intact.

8.4 Learning the Optimal ADL Policy

Given these constraints, we can treat severity and haircuts as an online control problem: choose $\{\theta_t, h_t\}$ after each price shock, conditioned on past shocks and responses, to minimize loss while respecting FIC/LIC. Because losses admit convex surrogates, we can apply classical online schemes (*e.g.* mirror descent) and enforce constraints via primal–dual updates so FIC/LIC remain satisfied at each round. We use mirror descent because it naturally handles convex losses with simple geometry (RAP/Pro-Rata) and pairs cleanly with primal–dual updates, matching the Stackelberg setting of committing to a policy while respecting per-round constraints [61, 62, 65, 100, 114].

Mirror Descent with IC constraints (MDIC). We adopt a primal–dual mirror–descent scheme tailored to the RAP geometry and non–smooth constraints:

1. Choose a mirror map ψ on the policy domain Π (*e.g.* a separable entropy or a Bregman geometry aligned with capped pro–rata).
2. Initialize $\pi_1 \in \Pi$ and dual multipliers $\lambda_1 \geq 0$ for constraints $c(\cdot; \cdot) \leq 0$.
3. For rounds $t = 1, 2, \dots$:
 - (a) Observe state s_t ; publish π_t (commitment). The follower/environment best–responds; incur loss $\ell_t(\pi_t)$ and observe a subgradient $g_t \in \partial \ell_t(\pi_t)$ and constraint feedback $c_t = c(\pi_t; s_t)$.
 - (b) Primal update (mirror step with penalties):

$$\pi_{t+1} = \operatorname{argmin}_{\pi \in \Pi} \left\{ \langle g_t + \nabla_{\pi} \langle \lambda_t, c(\pi; s_t) \rangle, \pi \rangle + \frac{1}{\eta_t} D_{\psi}(\pi \parallel \pi_t) \right\}.$$

- (c) Dual update (projected subgradient ascent on violations):

$$\lambda_{t+1} = [\lambda_t + \gamma_t c_t]_+.$$

4. Tune (η_t, γ_t) as $O(t^{-1/2})$ to balance regret and constraint violations.

This controller reduces to standard OCO when constraints are absent, while handling non–smooth IC/solvency constraints via the dual iterates.

Why MDIC. The IC set is convex but non-smooth (*e.g.* caps, max over scenarios). Mirror descent matches this geometry (KL on the simplex, log-barrier on θ) and keeps updates stable while preserving $O(\sqrt{T})$ rates.

Algorithmic Guarantees and Why Regret Matters. If f_t is Lipschitz, the follower is ε -best-response, and the IC set has bounded diameter, MDIC with $\eta_t \propto 1/\sqrt{T}$ yields

$$\text{Regret}_T = O\left(\sqrt{T D_\phi(\theta^* \|\theta_1)} + \sqrt{T D_\Phi(h^* \|h_1)}\right) + \varepsilon T,$$

with per-round IC feasibility (or sublinear average violation in the long-term CVaR variant). We care about regret because it benchmarks any adaptive policy against the best dynamic policy in hindsight. Showing that MDIC achieves $O(\sqrt{T})$ regret means MDIC learns fast, while $\Omega(T)$ regret shows static rules (Queue, RAP, Pro-Rata) can be worst-case bad. Proofs are in Appendix G.2.

Static Policies Have Linear Regret. In an alternating two-regime example (liquidity vs. stress) any fixed θ incurs $\Omega(T)$ regret, whereas MDIC adapts and achieves $O(\sqrt{T})$ (proven in Appendix G.2). This captures why static Queue, RAP, Pro-Rata are fragile in adversarial environments.

8.5 Follower Robustness

So far, we have focused on the leader (*e.g.* the exchange), as they have to commit to a policy. However, it is possible that if the follower (*i.e.*, a trader who reduces the deficit) knows enough about the leader’s strategy to adversely select against other traders. We address the strategic behavior of traders (“followers”) in response to ADL policies, focusing on adverse selection and timing games. Our description of follower behavior will be minimal with most details in Appendix G.4.

Adverse Selection and the Pro-Rata Death Spiral. As we saw earlier, if an ADL policy socializes losses uniformly without regards to risk (like Pro-Rata), low-risk traders subsidize high-risk traders. If the cost of subsidy exceeds the utility of trading for low-risk traders, they exit. This increases both the exchange’s average risk and probability of future deficits and can be viewed as a form of adverse selection. This is a direct consequence of the follower’s strategic response: low-risk traders, observing a policy that penalizes them disproportionately, rationally choose to exit (or not participate), leaving the exchange with a riskier pool of traders.

Let μ_i be the baseline utility of trader i and let $H_\pi(\lambda_i \mathcal{P}_n)$ be the random haircut variable for a trader with leverage λ_i . The “utility is greater than subsidy” condition can be described as $\mu_i - \mathbf{E}[H_\pi(\lambda_i, \mathcal{P}_n)] \geq u_0$, where u_0 is a baseline minimum utility a low-risk trader accepts. Under Pro-Rata, $\mathbf{E}[H_{\pi_{PR}}]$ depends on the average leverage $\bar{\lambda} = \frac{1}{n} \sum_i \lambda_i$, which penalizes low-risk traders with $\lambda_i < \bar{\lambda}$. The worst case is if there is a ‘death spiral’ of exits, where low-risk traders leaving sequentially increase the probability of other traders leaving, as the average $\bar{\lambda}$ increases on each exit. This is adverse selection that the exchange faces as the pool of traders becomes increasingly risky from a leverage standpoint. We formalize this in Appendix G.4.1 and provide explicit examples of a death spiral.

The Waiting Game. Traders may act as “Backstop Liquidity Providers” (BLP) by voluntarily filling liquidation orders to reduce the deficit D_t . For instance, traders who deposit to Hyperliquid’s HLP pool around the time of an ADL event act as BLPs. This presents an optimal stopping problem: a BLP can absorb the deficit immediately (incurring market risk and execution costs) or wait, hoping that the exchange reduces ADL severity θ_t (and thus the potential haircut) in subsequent rounds. If the exchange employs a policy like an exponential backoff that lowers severity over time, it incentivizes traders to wait, potentially exacerbating the deficit. To prevent such strategic delays, the exchange needs to provide a financial incentive to ensure that BLPs add liquidity in a timely manner.

This incentive often comes in the form of yield premium Γ_t . This premium can be viewed as an incentive (e.g., a rebate or a yield) that the exchange must pay to compensate BLPs for the risk of absorbing the deficit early. Intuitively, Γ_t quantifies the cost of immediate intervention, capturing the execution slippage and inventory risk premium required for a BLP to step in. By adjusting these incentives, an exchange can encourage "no-wait" conditions to speed up solvency recovery. Intuitively, these "No-Wait" conditions correspond to solvency-like constraints: $\theta_t D_t \geq \Gamma_t$. We formalize this with the following proposition:

Proposition 8.2 (No-Wait Condition). *Let Γ_t be the liquidity premium required for a Backstop Liquidity Provider (BLP) to absorb deficit D_t immediately rather than waiting for Δt . If the severity policy θ_t is a decreasing function of time (e.g. exponential backoff), a BLP will withhold liquidity unless the immediate haircut cost is cheaper than the discounted future cost:*

$$\theta_t D_t \leq \mathbf{E}_t[\theta_{t+\Delta t} D_{t+\Delta t}] - \Gamma_t.$$

Violating this condition induces a "Waiting Game" where solvency recovery is delayed purely by the follower's strategy.

Appendix G.4.2 details the construction of Γ_t and proves this proposition.

8.6 Dynamic Phase Transition in ADL

In Section 5, we established that static ADL rules face fundamental scaling limits. However, our results in that section were qualitative in nature, making it hard to use those results to design better mechanisms. One question that one might have is if there is a way to classify, based on properties of an asset (e.g. volatility, liquidity, etc.) whether one can use a simple static ADL strategy or if one needs a full dynamic strategy. Here we will construct conditions for quantifying an answer to this question.

To quantify the efficiency gap between static (Nash) and dynamic (Stackelberg) control, we evaluate the system using welfare functions $W(\pi)$ corresponding to the trilemma (Proposition J.7) objectives. Welfare functionals quantify how fair, profitable, or solvent the entire system is (as opposed to individual traders). Specifically, we show there is a *phase transition* that depends on the liquidity and volatility of the asset pair that determines when one needs a dynamic methods (like MDIC) or if one can use a simpler algorithm (like RAP). The phase transition is defined by an *order parameter*, which is a dimensionless quantity that captures the relative strength of the ADL severity to the market's liquidity and volatility. This phase transition occurs in the *efficiency gap* between static and dynamic methods — i.e., how much better at preserving solvency and/or revenue that a dynamic method is over static method.

Price of Anarchy for Nash vs. Stackelberg. The Price of Anarchy (PoA) is the standard metric for quantifying efficiency loss from strategic behavior across equilibrium concepts [97, 98]. In Stackelberg settings, a leader's commitment reshapes follower best responses and can systematically lower inefficiency relative to simultaneous-move (Nash) play. This is well documented in dynamic game models [17] and in smoothness-based PoA analyses that extend uniformly across equilibrium notions, including sequential and composed mechanisms [98, 109]. In our ADL context, the venue is a natural leader that can commit to a severity path, while traders react myopically to current and anticipated haircuts. This makes a PoA comparison between one-shot (Nash-like) and multi-round (Stackelberg) control both natural and informative.

We define the PoA for a welfare function W as the ratio of the optimal dynamic Stackelberg welfare to the welfare of a static, one-shot policy π_n^{stat} :

$$\text{PoA}_n(W) = \frac{W(\pi_n^{\text{dyn}})}{W(\pi_n^{\text{stat}})} \quad (35)$$

We consider both a fairness welfare $W_{\text{Fair}}(\pi) = \text{PTSR}_T(\pi)$ and a revenue welfare $W_{\text{LTV}}(\pi) = \text{LTV}_T(\pi)$. The severity load $\kappa_n = \theta_n n / b_n$ serves as the order parameter for the phase transition in fairness, while the heavy-tail index α (see §3.1.2) governs the phase transition in revenue. Using extreme-value scaling and mean-field equilibrium assumptions, we obtain PoA bounds that are robust to the precise equilibrium concept. Appendix H formalizes the scaling assumptions and presents the theorem. Informally:

Proposition 8.3 (Informal). *Let $\kappa_n = \theta_n / b_n$ be the normalized severity relative to the winner equity scale b_n .*

- **Bounded PoA (Static Sufficient):** *If the severity load stays bounded ($\sup_n \kappa_n < \infty$), leverage masses are comparable ($\ell_n^+ \asymp \ell_n^-$), and deficits are light-tailed, then static ADL is constant-factor optimal: $\sup_n \text{PoA}_n^{\text{Fair}} < \infty$ and $\sup_n \text{PoA}_n^{\text{LTV}} < \infty$.*
- **Unbounded PoA (Dynamic Necessary):** *If the load diverges ($\kappa_n \rightarrow \infty$) or leverage is concentrated ($\ell_n^+ / \ell_n^- \rightarrow 0$), Fairness PoA blows up. If deficits are heavy-tailed ($\alpha < 2$) and solvency is enforced, Revenue PoA also diverges. Static policies must sacrifice at least one objective, while a dynamic controller can maintain both.*

This implies a screening rule: if estimates satisfy $\hat{\kappa}_n \leq K$ and $\hat{\Lambda}_n \in [c, C]$, static ADL suffices; otherwise, dynamic control is required to navigate the Trilemma.

Practically, this result can be used to provide simple heuristics for screening assets using historical data to determine whether a simple static rule suffices or if a dynamic policy is required. Formal statements and proofs are given by Theorem H.3 and Proposition H.5 in Appendix H, and by Example H.4.

9 Empirical Analysis: The October 10 Event

We evaluate the proposed ADL mechanisms using a high-fidelity replay¹² of the October 10, 2025 liquidation cascade on Hyperliquid [96]. We use this replay to stress-test the ADL Trilemma (§2.5), quantify fairness via PTSR/PMR (§2.5 and Figure 6), and benchmark RAP and other static baselines (§7).

9.1 Experimental Setup

Dataset. We utilize the canonical reconstruction of the event [41], covering a 12-minute window (21:16–21:27 UTC) with \$2,103,111,431 in liquidations across roughly 160 assets. The replay comprises $T = 1097$ ADL shocks (201 with $D_t > 0$). Each shock records the realized deficit D_t (total negative equity of the losers) and the winner-equity vector $e_t^{\mathcal{W}} = (e_{t,1}, \dots, e_{t,|\mathcal{W}_t|})$. The winner equity vector gives us the *maximum haircut mass*, $W(e_t^{\mathcal{W}}) = \sum_{i=1}^{|\mathcal{W}_t|} e_{t,i}$. This is the maximum profit that could be seized to cover a deficit.

¹²The author profusely thanks Mauricio Trujillo (@conejocapital) and Hydromancer for providing, cleaning, and verifying this data.

Across the 201 shocks with $D_t > 0$, the realized shortfall sums to $\sum_t D_t = \$23,245,151.37$. The feasible haircut mass $\sum_t \min\{D_t, W(e_t^V)\}$ is only $\$6,620,946.51$. Consequently at least $\$16.6\text{M}$ of the realized deficit is structurally uncovered regardless of policy, establishing a hard lower bound on cumulative residuals.

As an explicit example, consider the asset WLD on Hyperliquid. There was a realized deficit of $D_t = \$11,558,884.49$ despite having a maximum haircut $W(e_t^V) = \$2,507.11$. This means that every ADL strategy exits that round with $\$11,556,377.38$ of residual bad debt. We note that the dataset provides the exact sequence of exogenous price shocks and trader positions, allowing for a counterfactual analysis of ADL responses holding market conditions fixed.

Policies. We compare the static benchmark π_{queue} against the dynamic controllers defined in Appendix I.3:

- *Static Priority* ($\pi_{\text{queue}}, \pi_{\text{smart}}$): The standard greedy mechanism (Queue) and its deficit-capped variant (Smart Queue). π_{smart} enforces $H_t \leq D_t$ but retains the greedy priority ranking.
- *Dynamic Severity* ($\pi_{\text{exp}}, \pi_{\text{md}}$): The Exponential Backoff and Mirror Descent controllers that optimize the scalar severity θ_t while allocating haircuts pro-rata.
- *Vector Optimization* (π_{vec}): The full allocation policy minimizing the composite loss L_t via high-dimensional mirror descent.
- *Hybrid Blend* (π_{hyb}): A convex combination of the production Queue and optimized Levered Pro-Rata allocations (20%/80% mix). This approximates the recommended pairing of queue-style guardrails with pro-rata fairness without retraining the live mechanism.

Revenue proxy. We construct a *revenue proxy* that estimates the expected fee and revenue loss for the exchange given a set of ADL haircuts. At a high level, we model how much churn traders experience post haircut, akin to other customer retention models [1, 25, 99]. Given the churn, we model a fee proxy that estimate the loss in fees per unit of notional size that churns. Formally, define the *churn probability* of winner i as the probability that the account exits the venue after receiving a haircut in round t .

Empirically, this quantity controls how much future revenue the exchange forfeits. We model churn with the exponential hazard $p_{t,i} = 1 - \exp(-\beta h_{t,i}/w_{t,i})$, which keeps responses linear for small haircuts yet saturates toward one for severe cuts, matching the account-level attrition we measure in the Hyperliquid replay. This mirrors the survival-style retention curves widely used in both classic and modern churn models in economics [13, 14, 25, 78, 99]. Full details of the calculation and the fee proxy are in Appendix I.2.

9.2 Results

Overshoot and Efficiency. We measure *overshoot* on each shock as $(H_t - D_t)_+$, the haircuts applied beyond the realized deficit D_t . Recall that for our optimized algorithms, we enforce the feasibility constraint (18), so $H_t \leq D_t$. This means that most of the strategies will have very low if not zero overshoot.¹³

The production queue π_{queue} fully liquidates accounts and ignores feasibility, accumulating $\$653.6\text{M}$ of overshoot with a $\$47.1\text{M}$ single-shock spike. All other controllers keep overshoot

¹³The non-zero overshoot is due to the fact that the strategies are generally local gradient steps, which accumulate some noise and error in how well they enforce constraints.




Figure 3: ADL overshoot ($H_t - D_t$). Dashed: per-shock overshoot. Solid: cumulative overshoot. The static policy π_{queue} accumulates $\approx \$653\text{M}$ of excess haircutting. Dynamic policies ($\pi_{\text{md}}, \pi_{\text{vec}}$) keep overshoot negligible by construction, while π_{exp} maintains it within a bounded constant factor.

negligible: π_{smart} is exactly capped with $H_t = D_t$; π_{md} and π_{dyn2} have maxima of $\$89,537.21$ and $\$268,670.09$, respectively; π_{exp} peaks at $\$223,886.86$; the vector optimizer peaks at $\$974,644.37$. We also study the effect of the hybrid policy π_{hyb} which trades solvency for revenue smoothly. This strategy reduces queue overshoot by $\approx 80\%$ (to $\$134.4\text{M}$ total, $\$9.4\text{M}$ max) while retaining much of the revenue retention of the dynamic controllers.

The ADL Trilemma. We analyze the ADL trilemma (Proposition 2.1) empirically by reporting, for each policy, its solvency (residual shortfall), fairness (PTSR/PMR), and revenue (fee-retention) metrics. These three objectives summarize the trade-offs that each ADL policy makes.

Figure 7 anchors the “solvency-at-any-cost” corner of the trilemma: π_{queue} slashes residual deficits (*i.e.*, higher solvency) yet incurs runaway overshoot and fairness penalties (*i.e.*, lower fairness and revenue). Figure 8 shows that capping overshoot via π_{smart} still leaves concentrated losses (*i.e.*, low fairness) and elevated residual deficits once churn feedback kicks in (*i.e.*, lower long term revenue). Figures 9 and 10 illustrate how dynamic mirror-descent controllers match the queue’s residual deficits without its welfare cost, *i.e.*, they have better revenue preserving properties than Queue while preserving approximately the same amount of solvency.

On the other hand, Figure 11 demonstrates that (static) exponential backoff strategies can be competitive with dynamic policies if the deficit is sufficiently bounded. By jointly optimizing allocation and severity, Figure 12 shows π_{vec} delivering the fairest loss sharing and the steepest long term revenue recovery. Finally, the hybrid π_{hyb} in Figure 13 blends the production queue with Levered Pro-Rata, eliminating $\approx 80\%$ of Queue’s overshoot and cutting revenue loss by nearly an order of magnitude while keeping residual deficits low.

Regret Decomposition. Figure 7 makes clear that overshoot (orange) and fairness (red) dominate the queue baseline. Figure 8 cuts the overshoot bar almost entirely but replaces it with a larger residual (blue) mass as liquidity flees. Mirror-descent controllers in Figures 9 and 10 shrink every component simultaneously by aligning severity with observed deficits, with the warm-started variant tightening the residual tail. Figure 11 trades a small overshoot allowance for even lower fairness




Figure 4: Outstanding negative equity (cumulative residual deficits). Lower is better. Despite its aggression, π_{queue} fails to clear the deficit effectively due to feedback loops. π_{exp} achieves competitive solvency without the welfare cost.

cost, while Figure 12 compresses all three components and traces the Pareto frontier across solvency, fairness, and revenue.

9.3 Interpretation

Taken together, the replay shows the production Queue is behaving exactly as the unbounded PoA story predicts: it buys solvency with overshoot and churn, and pays with fairness and future fees. Policies that separate “how much” (θ_t) from “who” (h_t) — *e.g.* Exponential Backoff (π_{exp}) or Mirror Descent (π_{md}) — land in a strictly better region of the trilemma: they keep solvency comparable while lifting PTSR/PMR and retaining revenue. The convex hybrid π_{hyb} makes this trade-off explicit: a simple 20%/80% mix of Queue and Levered Pro-Rata inherits most of the fairness and revenue gains of the vector policy without giving up the familiar solvency profile.

Based on these results, we offer three concrete prescriptions for Hyperliquid and similar venues:

1. *Decouple severity from allocation.* Use a scalar feedback rule (*e.g.* π_{exp} or π_{md}) to set θ_t from the realized deficit D_t , then allocate separately.
2. *Adopt levered pro-rata.* Replace the greedy queue with Levered Pro-Rata (Eq. 21) or the optimized vector h_t ; this keeps PTSR/PMR near 1 and aligns haircuts with capacity.
3. *Blend for transparency.* A light hybrid (20% Queue, 80% LPR) preserves the familiar queue feel while cutting overshoot by $\approx 80\%$ and recovering most of the revenue and fairness gains.

10 Conclusion and Future Work

In this paper, we provide the first formalism for ADL mechanisms that allows us to compare the performance of different policies. We first that a negative result: it is impossible for a perpetuals to simultaneously optimize for exchange solvency, long term revenue, and fairness to traders. We then try to get around the negative result by analyzing strategies that optimize each of these components




Figure 5: Largest winner haircut (USD). π_{queue} places a massive, spikey burden on single participants. Dynamic policies ($\pi_{\text{md}}, \pi_{\text{exp}}$) distribute the load, respecting fairness constraints.

individually. Our results demonstrate that the heuristic strategy employed by Hyperliquid, Binance, and others that stems from 2015 is actually suboptimal in all dimensions. We empirically validate our results on Hyperliquid data from October 10, 2025 and demonstrate the superiority of dynamic and even other static algorithms to the live policy.

Our results suggest a few directions for future inquiry. Implementing these policies in production has a number of challenges that we don’t consider in the paper. For instance, as exchanges grow and have many accounts with very smaller notation sizes, it makes more sense to choose a subset of users who are “eligible for ADL”. This lowers the computational burden of executing ADL during times of duress (when execution latency is high) while also ensuring that smaller users are not subsidizing larger users. Practical limits on ADL applicability likely change the results we have here and need to be studied further.

ADL policies also likely change in their efficacy if they were able to be executed in a fully privately manner. In this paper, we ignored modeling execution costs for exchanges upon liquidation and position closure. The modeling of such execution costs is nuanced and likely leads to incorrect conclusions about ADL if not carefully handled¹⁴. When private execution is possible, the impact of these costs is muted as follower strategies are not able to react to ADL shocks.

While centralized exchanges effectively offer ADL privacy (*i.e.*, other users do not know who else was autodeleveraged besides themselves), they still leak information to the market via public changes to the order book. With superior privacy guarantess (likely provided by fully homomorphic encryption, zero knowledge proofs, and secure multi-party computation), one could imagine ADL policies being able to avoid the adverse selection effects we describe in the follower strategies section. As ADL mechanisms rely increasingly on backstop liquidity vaults (such as Hyperliquid’s HLP and Lighter’s LLP), privacy of ADL execution becomes increasingly important to reduce costs for the exchange and traders.

¹⁴See, *e.g.* Storm [106, 107] for examples of overly naive analyses that conflate execution costs with ADL mechanism design.




Figure 6: PTSR (top) and PMR (bottom) on log scale. π_{queue} exhibits ratios $> 10^6$, indicating extreme moral hazard. Adaptive policies remain within the theoretical corridor (≈ 1 – 10).

Finally, we note that the results in this paper focused on a single margin model, where the posted collateral was the numéraire. Many exchanges, including Hyperliquid, have cross-margin ADL support. This changes the modeling of execution costs and forces us to explicitly model market impact on cash and/or collateral balances.

11 Acknowledgments

The author would like to thank Udai Parvathaneni, Nathan Sheng, JD Maturen, Kamil Yusubov, and Luke Sterle from Gauntlet for helpful discussions around how to quantify realistic ADL scenarios (such as October 10, 2025). The author would also like to thank Yi Sun, Anirudh Pai, Kshitij Kulkarni, Theo Diamandis, Matheus V. X. Ferreira, Guillermo Angeris, Victor Xu, and Vinayak Kurup for helpful discussions. Finally, the author really appreciates the help with procuring and cleaning Hyperliquid data provided by SonarX, Hydromancer, and Mauricio Trujillo (@conejocapital).




Figure 7: Regret decomposition for the production Queue baseline π_{queue} . Overshoot (orange) and fairness (red) penalties dominate, illustrating static inefficiency.




Figure 8: Regret decomposition for the Smart Queue π_{smart} . Capping overshoot helps, but churn-driven residual deficits and fairness costs still dominate regret.




Figure 9: Regret decomposition for Mirror Descent π_{md} . Adaptive severity minimizes overshoot while balancing fairness and residual losses.




Figure 10: Regret decomposition for the Dynamic2 controller π_{dyn2} . Warm-started mirror descent keeps overshoot tame while moderating fairness penalties.




Figure 11: Regret decomposition for Exponential Backoff π_{exp} . Feedback-tuned severity contains overshoot without sacrificing solvency or fairness.




Figure 12: Regret decomposition for Vector Mirror π_{vec} . Joint severity and allocation updates compress every component, tracing the Pareto frontier.




Figure 13: Regret decomposition for the Hybrid blend π_{hyb} (20% Queue, 80% Levered Pro-Rata). Mixing the production policy with optimized allocations slashes overshoot and revenue loss while keeping residual deficits comparable to Queue.

Appendix

A Notation and conventions

Throughout the appendix we use consistent notation:

- Shocks are indexed by $t = 1, 2, \dots, T$; winners W_t with equities $e_{t,i} > 0$ and losers L_t .
- Deficit $D_t \geq 0$; severity $\theta_t \geq 0$; haircut vector $h_t \in [0, 1]^{|W_t|}$ with survivors $s_{t,i} = (1 - h_{t,i})e_{t,i}$.
- Per-round budget $B_t = \min\{\theta_t D_t, \sum_i w_{t,i}\}$ for severity policies and $B_t = w_t^\top h_t$ for vector policies; $B_t^* = \min\{D_t, \sum_i w_{t,i}\}$.
- Per-account caps are $\beta_{t,i} \in [0, 1]$ so that $0 \leq h_{t,i} \leq \beta_{t,i}$.
- Risk weights: $\rho(\lambda) = \lambda \psi(1/\lambda)$ and $g(\lambda) = \rho(\lambda)/\lambda$; when ordering by risk we use ρ .
- Orders: $x \prec_w y$ denotes weak submajorization (on decreasing rearrangements).
- Asymptotics: $X_n \asymp_p Y_n$ means there exist constants $c, C > 0$ such that $c|Y_n| \leq |X_n| \leq C|Y_n|$ with high probability.

B Liquidations, Autodeleveraging, and Insurance Funds

B.1 Bankruptcy Price Example

We illustrate the bankruptcy price calculation with an example. Fix $\ell^{\max} = 10$ (so $m_I = 0.10$) and $p_0 = 1$. Consider the five running positions from $\mathcal{P}_5$:

$$\begin{aligned} \mathbf{p}_A &= (q, c, b) = (1, 2, +1), & \mathbf{p}_B &= (1, 2/3, +1), & \mathbf{p}_C &= (4, 8/3, -1), \\ \mathbf{p}_D &= (1, 2/19, +1), & \mathbf{p}_E &= (1, 10/99, -1). \end{aligned}$$

Applying Eq. (10) with $\Gamma = 0$ and $p_t = p_0 = 1$ gives

$$\begin{aligned} p^{bk}(\mathbf{p}_A) &= \max\{0, 1 - 2\} = 0, \\ p^{bk}(\mathbf{p}_B) &= 1 - \frac{2}{3} = \frac{1}{3}, \\ p^{bk}(\mathbf{p}_C) &= 1 - \frac{8/3}{-4} = 1 + \frac{2}{3} = \frac{5}{3} \text{ (1.6667)}, \\ p^{bk}(\mathbf{p}_D) &= 1 - \frac{2}{19} \approx 0.8947, \\ p^{bk}(\mathbf{p}_E) &= 1 + \frac{10}{99} \approx 1.1010. \end{aligned}$$

Thus A is robust to a full drop to zero; B (long 1.5x) has a low bankruptcy price; C (short 1.5x) bankrupts only if the mark rises above $\frac{5}{3}$; D/E (long 9.5x/9.9x) have high bankruptcy prices close to 1, making negative equity likely if liquidations lag.

B.2 Liquidation Price Example

For $\mu = 0.10$ the liquidation prices evaluate to factors $\frac{1}{1-\mu}$ for longs and $\frac{1}{1+\mu}$ for shorts. Using the same five positions with $p_{t_i} = 1$ and $\Gamma = 0$,

$$\begin{aligned}\hat{p}^{liq}(\mathbf{p}_A, 0.10) &= \frac{1}{0.9} \cdot 0 = 0, \\ \hat{p}^{liq}(\mathbf{p}_B, 0.10) &= \frac{1}{0.9} \cdot \frac{1}{3} \approx 0.3704, \\ \hat{p}^{liq}(\mathbf{p}_C, 0.10) &= \frac{1}{1.1} \cdot \frac{5}{3} \approx 1.5152, \\ \hat{p}^{liq}(\mathbf{p}_D, 0.10) &= \frac{1}{0.9} \cdot \frac{17}{19} \approx 0.9942, \\ \hat{p}^{liq}(\mathbf{p}_E, 0.10) &= \frac{1}{1.1} \cdot \frac{109}{99} \approx 1.0010.\end{aligned}$$

Long A would liquidate only if the oracle fell to 0 (*i.e.*, never occurs in practice). The short C becomes liquidatable only when the oracle exceeds its threshold; the high-leverage long D and short E become liquidatable close to 1 if collateral is not topped up.

B.3 Execution Price Example

We summarize execution with a directional linear impact rule consistent with our notation: selling to close a long uses $P^{sell}(x) = p_t - \alpha x$ and buying to close a short uses $P^{buy}(x) = p_t + \alpha x$ with $\alpha > 0$; the volume-weighted execution for a slice is $p^{exec} = p_t - \frac{\alpha}{2}\Delta q$ if $b_i = +1$ (sell) and $p^{exec} = p_t + \frac{\alpha}{2}\Delta q$ if $b_i = -1$ (buy). Directional linear impact with a single α : selling (closing a long) uses $P^{sell}(x) = p_t - \alpha x$, buying (closing a short) uses $P^{buy}(x) = p_t + \alpha x$. The slice VWAP over $[0, \Delta q]$ is $p^{exec} = p_t \mp \frac{\alpha}{2}\Delta q$ (minus for sells, plus for buys). Fix $\alpha = 1.0$ and choose Δq per case:

- $\mathbf{p}_A$ (Long): $p_t = 1.30$, $\Delta q = 0.5$ gives $p^{exec} = 1.30 - 0.5 \cdot 0.5 = 1.05$. Here $p^{bk}(\mathbf{p}_A) = \max\{p_t - 2, 0\} = 0$, so $p^{exec} > p^{bk}$.
- $\mathbf{p}_B$ (Long): $p_t = 0.95$, $\Delta q = 0.2$ gives $p^{exec} = 0.95 - 0.5 \cdot 0.2 = 0.85$. With $p^{bk}(\mathbf{p}_B) = \max\{0.95 - \frac{2}{3}, 0\} \approx 0.2833$, we have $p^{exec} > p^{bk}$.
- $\mathbf{p}_C$ (Short): $p_t = 1.60$, $\Delta q = 2.0$ gives $p^{exec} = 1.60 + 0.5 \cdot 2.0 = 2.60$. Since $p^{bk}(\mathbf{p}_C) = p_t + \frac{8/3}{4} = p_t + \frac{2}{3} = 2.2667$, $p^{exec} > p^{bk}$ (adverse for a short).
- $\mathbf{p}_D$ (Long; targeted): $p_t = 0.98$, $\Delta q = 0.4$ gives $p^{exec} = 0.98 - 0.5 \cdot 0.4 = 0.78$. With $p^{bk}(\mathbf{p}_D) = p_t - \frac{2}{19} \approx 0.8747$, we achieve $p^{exec} < p^{bk}$.
- $\mathbf{p}_E$ (Short): $p_t = 1.05$, $\Delta q = 0.4$ gives $p^{exec} = 1.05 + 0.5 \cdot 0.4 = 1.25$. With $p^{bk}(\mathbf{p}_E) = p_t + \frac{10}{99} \approx 1.151$, we have $p^{exec} > p^{bk}$.

B.4 Liquidation Costs Example

To ground the fee model, let $\tau(\Delta q) = \tau^{fix} + \phi^{mark} p_t \Delta q + \phi^{exec} p^{exec} \Delta q$ as in practice. Consider $p_t = 1.30$, a slice $\Delta q = 0.50$, and a realized $p^{exec} = 1.32$. Two parameterizations:

- Binance: $(\tau^{fix}, \phi^{mark}, \phi^{exec}) = (0, 40 \text{ bps}, 0)$ [22]. Then $\tau = 0.0040 \cdot 1.30 \cdot 0.50 = 0.0026$.
- Hyperliquid: $(\tau^{fix}, \phi^{mark}, \phi^{exec}) = (0, 20 \text{ bps}, 10 \text{ bps})$ [68]. Then $\tau = 0.0020 \cdot 1.30 \cdot 0.50 + 0.0010 \cdot 1.32 \cdot 0.50 \approx 0.00130 + 0.00066 = 0.00196$.

Rates and formulas vary by venue and contract; the above are illustrative parameterizations consistent with public documentation that liquidation fees are charged and, on centralized venues like Binance, credited to the insurance fund.

B.5 Liquidation Strategy Example

Consider short $\mathbf{p}_E$ when the mark jumps to $p_t = 5.5$ (ignore funding for this step). Equity before liquidation is $e \approx c_E - q_E(p_t - p_0) = \frac{10}{99} - 1 \cdot 4.5 \approx -4.399$. Let $\mu = 0.10$ and a linear fee $\tau(\Delta q) = \phi p_t \Delta q$ with $\phi = 30\text{bps}$. Suppose execution is $p^{exec} = 5.55$. Using (11) and $b = -1$, the minimal slice that restores maintenance solves

$$\Delta q = \frac{\mu p_t q - e}{b(p^{exec} - p_t) - \phi p_t + \mu p_t} = \frac{0.1 \cdot 5.5 \cdot 1 - (-4.399)}{-0.05 - 0.003 \cdot 5.5 + 0.1 \cdot 5.5} \approx \frac{4.949}{0.4835} \approx 10.24.$$

Since $\Delta q > q_E$, a greedy policy would fully close E (cap at $\Delta q = q_E = 1$).

B.6 Bad Debt Example

Consider the high-leverage long $\mathbf{p}_D$ and a slice of size $\Delta q = 0.4$ at $p_t = 0.98$. Suppose the realized execution is $p_D^{exec} = 0.78$ while the bankruptcy level is $p_D^{bk} \approx 0.8747$. Since $p_D^{exec} < p_D^{bk}$, the realized shortfall from this slice is

$$(p_D^{bk} - p_D^{exec}) \Delta q \approx (0.8747 - 0.78) \cdot 0.4 \approx 0.0379,$$

which contributes this amount to the period bad debt D_t (cf. Eq. (13)). Coverage follows the solvency waterfall: the insurance fund pays $\min\{\text{IF}_t, D_t\}$ and any residual shortfall is socialized via ADL (see §2.3, §2.4). The short case is symmetric: buying to close above bankruptcy ($p^{exec} > p^{bk}$) realizes a positive contribution to D_t .

B.7 Anatomy of a Liquidation

Given the bankruptcy, liquidation, and execution prices, we can now describe the high-level algorithm that liquidations follow. We note that many live liquidation systems will have much more complex liquidation algorithms. These complexities deal with the coordination costs of coordinating many parties (*e.g.* oracle provider, liquidators, spot order book liquidity) and precise models that exchanges use for their liquidation strategy. However, we effectively lump all of these complexities into the definition of the liquidation strategy. The following liquidation loop is run on every oracle update received by a perpetuals exchange:

- For $\mathbf{p}_{i,t} \in \mathcal{P}_n$
 - If the maintenance margin condition (7) is violated for $\mathbf{p}_{i,t}$
 1. Remove the position $\mathcal{P}_n \leftarrow \mathcal{P}_n - \{\mathbf{p}_{i,t}\}$
 2. Estimate quantity to liquidate $\Delta q_i \leftarrow L(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T})$
 3. Liquidator executes Δq_i -sized liquidation and returns their execution price $p^{exec}(\Delta q_i)$
 4. Update position: $\mathbf{p}'_{i,t} = (q_i - \Delta q_i, c_i + p^{exec} \Delta q_i - \tau_t(\Delta q_i), t_i, b_i)$
 5. Re-add the position position: $\mathcal{P}_n \leftarrow \mathcal{P}_n \cup \{\mathbf{p}'_{i,t}\}$
 6. Update equity using (12)
 - If $\mathbf{p}_{i,t}$ has bad debt, $\tilde{e}(\mathbf{p}_{i,t}, p_{1:T}, \hat{p}_{1:T}, \Delta q_i) < 0$, then
 - Attempt to use the insurance fund, if it exists, to cover the bad debt (§2.3)
 - If the insurance fund is insufficiently sized, utilize an ADL mechanism (§2.4)

Example. We illustrate a five-step path using the running set $\mathcal{P}_5$ from above. Take $T = 5$, $p_0 = 1$ and

$$p_{0:5} = (1.00, 0.96, 0.94, 0.97, 1.05, 1.12), \quad \hat{p}_t = p_t \quad (t = 0, \dots, 5), \quad \mu = m_I = 0.10.$$

Executions follow the directional linear impact rule introduced in the execution example: for a slice of size Δq at time t , the volume-weighted execution is

$$\begin{aligned} p^{exec} &= p_t - \frac{\alpha}{2} \Delta q \quad (\text{sell to close a long, } b = +1), \\ p^{exec} &= p_t + \frac{\alpha}{2} \Delta q \quad (\text{buy to close a short, } b = -1). \end{aligned}$$

with $\alpha > 0$. We take $\alpha = 1.0$ and choose Δq via the loop's liquidation size $\Delta q_i = L(\mathbf{p}_{i:t}, p_{1:T}, \hat{p}_{1:T})$.

D liquidates at $t = 2$ (no bad debt). At $t = 2$ we have $p_2 = 0.94$ and the maintenance condition (7) is violated for $\mathbf{p}_{D,2}$, so the loop attempts a partial liquidation. Take $\Delta q_D = L(\mathbf{p}_{D,2}, \cdot) = 0.20$ for illustration. By (10),

$$p^{bk}(\mathbf{p}_{D,2}) = p_2 - \frac{2}{19} \approx 0.8347, \quad p_D^{exec} = p_2 - \frac{\alpha}{2} \Delta q_D = 0.94 - 0.10 = 0.84.$$

Since $p_D^{exec} > p^{bk}(\mathbf{p}_{D,2})$, this slice executes without bad debt; the position is updated to $\mathbf{p}'_{D,2} = (q_D - \Delta q_D, c_D + p_D^{exec} \Delta q_D - \tau_2(\Delta q_D), t_D, b_D)$ and equity is updated per (12) before reinserting $\mathbf{p}'_{D,2}$ into $\mathcal{P}_n$.

E becomes bad debt at $t = 4$ (short; liquidation fails). At $t = 4$ we have $p_4 = 1.05$ and (7) is violated for $\mathbf{p}_{E,4}$ with $b_E = -1$. The loop selects a liquidation size; take a full close $\Delta q_E = L(\mathbf{p}_{E,4}, \cdot) = 1$. By (10),

$$p^{bk}(\mathbf{p}_{E,4}) = p_4 + \frac{10}{99} \approx 1.1510, \quad p_E^{exec} = p_4 + \frac{\alpha}{2} \Delta q_E = 1.05 + 0.50 = 1.55.$$

For a short, $p^{exec} > p^{bk}$ realizes bad debt. The loop records the shortfall

$$D_4 = (p_E^{exec} - p^{bk}(\mathbf{p}_{E,4})) \Delta q_E \approx 0.399, \quad \tilde{e}(\mathbf{p}_{E,4}, p_{1:5}, \hat{p}_{1:5}, \Delta q_E) = -D_4 < 0,$$

and then attempts coverage via the insurance fund (up to $\min\{IF_4, D_4\}$); any residual shortfall R_4 defined by Eq. (16) is socialized by ADL (see §2.4).

B.8 Optimal Capital Structure Derivation

In this section, we compute the optimal static insurance fund size IF^* that trades off the opportunity cost of capital and expected uncovered losses beyond IF .

Setup. Let D_T denote the round deficit with pdf f_D and tail $\bar{F}_D(x) = \Pr[D_T > x]$. Let $r > 0$ be the per-unit capital cost and $\kappa > 0$ the per-unit social loss weight of uncovered deficits. The objective is

$$\min_{IF \geq 0} \mathcal{J}(IF) = r IF + \kappa \mathbb{E}[(D_T - IF)_+]$$

which equals $r IF + \kappa \int_{IF}^{\infty} (x - IF) f_D(x) dx$ when D_T is continuous.

Optimality condition. Differentiating yields,

$$\mathcal{J}'(IF) = r - \kappa \bar{F}_D(IF).$$

Hence any interior minimizer satisfies $\bar{F}_D(IF^*) = r/\kappa$, i.e.,

$$IF^* = \bar{F}_D^{-1}\left(\frac{r}{\kappa}\right) = \text{VaR}_{1-r/\kappa}(D_T).$$

Assumptions. In order for this argument to hold, we assume that $\mathcal{J}$ is convex and differentiable. Moreover, if $r \geq \kappa$ then we define $IF^* = 0$, whereas if $r \rightarrow 0$, then $IF^* \rightarrow \sup D_T$.

C Moral Hazard and Extreme Value Analysis

In this appendix, we formalize the moral hazard properties of ADL mechanisms. We analyze the optimal control of the Profitability-to-Total-Solvency Ratio (PTSR) and the Profitability-to-Maximum Solvency Ratio (PMR) defined in §3.1.2, and derive their asymptotic behavior under distributional assumptions.

C.1 Setup and Assumptions

We work in the *large-market limit* ($n \rightarrow \infty$) under the heavy-tailed assumptions characteristic of crypto markets. Recall that $D_T^\pi = \theta_\pi D_T$ is the total socialized loss and $\Delta_T^\pi = \theta_\pi \Delta_T$ is the maximum socialized shortfall under policy π . The survivor of the top winner is denoted ω_T^π .

Assumption A (regular variation). The right tail of the winner equity distribution $\bar{F}_+(x)$ and the right tail of the loser shortfall distribution $\bar{F}_-(x)$ are regularly varying with indices $\alpha_+ > 0$ and $\alpha_- > 0$, respectively. That is, $\bar{F}_\pm(x) = L_\pm(x)x^{-\alpha_\pm}$ where $L_\pm$ are slowly varying functions.

Assumption B (LLN and EVT). We assume the standard Law of Large Numbers (LLN) and Extreme Value Theory (EVT) scaling limits apply:

- *Aggregates:* $W_T/n \xrightarrow{P} \mu_+$ and $D_T/n \xrightarrow{P} \mu_-$, for constants $\mu_\pm \in (0, \infty)$.
- *Extremes:* The maximum winner ω_T scales as $b_{k_n}^+ = F_+^{\leftarrow}(1 - 1/n)$, and the maximum loser shortfall Δ_T scales as $b_{m_n}^- = F_-^{\leftarrow}(1 - 1/n)$.

C.2 Optimal Control of Moral Hazard

C.3 Queue maximizes top-winner damage

We first establish that the *Queue* (or Top-First) rule minimizes the moral hazard metrics defined in the main text for any fixed budget.

Proposition C.1 (Queue Minimizes Top Survivor). *Fix a budget $H = D_T^\pi$. Let ω_T be the equity of the largest winner. For any feasible haircut vector h satisfying $\sum h_i e_i = H$, the top-winner survivor ω_T^π satisfies*

$$\omega_T^\pi \geq \max\{\omega_T - H, 0\}.$$

Equality is attained by the Queue rule, which sets the haircut on the top winner to $h_{(1)} = \min(H/\omega_T, 1)$ and others to 0 (until $h_{(1)}$ saturates).

Proof. Let $h_{(1)}$ be the haircut applied to the top winner. Since $h_i e_i \geq 0$ for all i , we have $h_{(1)} \omega_T \leq \sum_i h_i e_i = H$. The survivor is $\omega_T^\pi = \omega_T - h_{(1)} \omega_T \geq \omega_T - H$. Since equity cannot be negative, $\omega_T^\pi \geq \max\{\omega_T - H, 0\}$. The Queue rule greedily allocates the budget to the largest position, achieving $h_{(1)} \omega_T = \min(H, \omega_T)$, thus attaining the lower bound. $\square$

Corollary C.2 (Minimality of PTSR/PMR). *Since D_T^π and Δ_T^π are fixed for a given policy severity, the Queue rule minimizes both PTSR_T and PMR_T among all budget-balanced policies.*

Gap versus pro-rata. The Queue rule minimizes moral hazard but concentrates the entire loss on the top winner (extreme inequality). In contrast, the Pro-Rata rule spreads the loss proportionally across all winners, prioritizing *smoothness* (treating similar positions similarly) over minimizing the top survivor's burden. For $H \leq \omega_T$, the survivor gap is

$$\omega_T^{\text{PR}} - \omega_T^{\text{Queue}} = H \left(1 - \frac{\omega_T}{W_T} \right).$$

This gap scales linearly with the budget H , quantifying the “cost of fairness”: by choosing the smoother Pro-Rata allocation, the system allows the top winner to retain more profit than is strictly necessary to cover the deficit.

C.4 Asymptotic Scaling Results

We now characterize the asymptotic behavior of PTSR and PMR under “gentle” policies (like Pro-Rata) where the top winner is not specifically targeted.

Theorem C.3 (PTSR scaling). *Under Assumptions A and B, for any policy π with severity θ_n where $\omega_T^\pi \sim \omega_T$ (e.g., Pro-Rata with $H \ll W_T$), the PTSR scales as*

$$\text{PTSR}_T(\pi) \asymp_p \frac{b_{k_n}^+}{\theta_n n}.$$

Proof. By definition, $\text{PTSR}_T(\pi) = \mathbf{E}[\omega_T^\pi / D_T^\pi]$. Under the hypothesis, the numerator scales as $\omega_T \sim b_{k_n}^+$. The denominator is $D_T^\pi = \theta_n D_T$. By the LLN, $D_T \sim \mu_- n$, so $D_T^\pi \sim \theta_n \mu_- n$. Thus, the ratio scales as $b_{k_n}^+ / (\theta_n n)$. Using bounded convergence for the expectation yields the result. $\square$

Implication. The behavior of PTSR depends critically on the tail class of winner equities:

- *Pareto (Heavy) Tails:* $b_{k_n}^+ \asymp n^{1/\alpha_+}$. Here $\text{PTSR}_T \asymp n^{1/\alpha_+ - 1} / \theta_n$. Moral hazard vanishes ($\text{PTSR} \rightarrow 0$) if and only if winners have finite mean ($\alpha_+ > 1$). If $\alpha_+ < 1$, the top survivor grows faster than the aggregate deficit, making the moral hazard wedge permanent.
- *Exponential/Gaussian (Light) Tails:* $b_{k_n}^+ \asymp (\log n)^\gamma$. Here $\text{PTSR}_T \asymp (\log n)^\gamma / (n \theta_n)$. Since polylog growth is slower than linear, moral hazard vanishes rapidly for any non-vanishing severity θ_n , as the aggregate deficit overwhelms the largest individual winner.

Theorem C.4 (PMR Scaling). *Assume winner equities have mass ℓ_n^+ and loser deficits have mass ℓ_n^- (representing total leverage), and that the underlying normalized distributions satisfy Assumption A. The PMR scales as:*

$$\text{PMR}_T(\pi) \asymp_p \frac{1}{\theta_n} \cdot \frac{\ell_n^+}{\ell_n^-} \cdot \frac{b_{k_n}^+}{b_{m_n}^-} \asymp \frac{1}{\theta_n} \cdot \frac{\ell_n^+}{\ell_n^-} \cdot n^{\frac{1}{\alpha_+} - \frac{1}{\alpha_-}}.$$

Proof. We have $\text{PMR}_T(\pi) = \mathbf{E}[\omega_T^\pi / \Delta_T^\pi]$. The top winner scales with total winner leverage mass: $\omega_T \sim \ell_n^+ b_{k_n}^+$. The maximum loser shortfall scales with total loser leverage mass: $\Delta_T \sim \ell_n^- b_{m_n}^-$. The budget balance condition implies $\Delta_T^\pi = \theta_n \Delta_T$. Thus, the ratio scales as

$$\frac{\ell_n^+ b_{k_n}^+}{\theta_n \ell_n^- b_{m_n}^-} = \frac{1}{\theta_n} \frac{\ell_n^+}{\ell_n^-} \frac{b_{k_n}^+}{b_{m_n}^-}.$$

Substituting the regular variation scalings $b_{k_n}^+ \sim n^{1/\alpha_+}$ and $b_{m_n}^- \sim n^{1/\alpha_-}$ yields the result. $\square$

Theorem C.4 decomposes moral hazard into three components: (1) *Policy Severity* ($1/\theta_n$): Lower severity amplifies PMR. (2) *Leverage Imbalance* (ℓ_n^+/ℓ_n^-): If the winning side holds more leverage mass, PMR increases. (3) *Tail Risk* ($n^{1/\alpha_+ - 1/\alpha_-}$): Heavier winner tails relative to losers drive PMR divergence. This decomposition highlights that even with fair tails ($\alpha_+ = \alpha_-$), a systemic leverage imbalance ($\ell_n^+ \gg \ell_n^-$) can sustain a high PMR. Specifically, if the exchange allows winners to be significantly more leveraged than losers (a "risk-on" imbalance), the top winner's survival will systematically outstrip the worst-case socialized loss, creating a persistent moral hazard where maximal profits are privatized while maximal losses are capped.

C.5 Relationship to Classical Risk Measures

These two metrics have natural interpretations in terms of financial risk measures. The deficit D_T corresponds to the aggregate *Expected Shortfall* (ES) of the losing tail, while Δ_T corresponds to the *Value-at-Risk* (VaR) at the extreme quantile ($1/n$). Specifically, PTSR compares the *Maximum Profit* to the *Aggregate Socialized Loss* (ES-like), measuring efficiency in bulk. PMR compares the *Maximum Profit* to the *Maximum Socialized Loss* (VaR-like), measuring efficiency in the extreme tail. A high PMR implies that the system permits "unicorn" wins that vastly exceed the worst-case individual losses, potentially incentivizing excessive risk-taking if traders perceive a capped downside but unbounded upside.

We further strengthen the connection to classical risk measures by showing that Queue not only minimizes the top survivor in expectation, but also minimizes it in the sense of VaR and ES at *every* tail level.

Proposition C.5 (Queue minimizes VaR/ES of the top survivor). *Fix any budget $h \geq 0$ and $\alpha \in (0, 1)$. For any feasible haircut vector h with $\sum_i h_i e_i = h$,*

$$\omega_T^\pi \geq (\omega_T - h)_+ \quad a.s.$$

Consequently,

$$\text{VaR}_\alpha(\omega_T^\pi) \geq \text{VaR}_\alpha((\omega_T - h)_+), \quad \text{ES}_\alpha(\omega_T^\pi) \geq \text{ES}_\alpha((\omega_T - h)_+).$$

The Queue rule attains equality. Moreover, the following identities hold:

$$\begin{aligned} \text{VaR}_\alpha((\omega_T - h)_+) &= \max\{\text{VaR}_\alpha(\omega_T) - h, 0\}, \\ \text{ES}_\alpha((\omega_T - h)_+) &= \frac{1}{1 - \alpha} \int_\alpha^1 \max\{\text{VaR}_u(\omega_T) - h, 0\} du. \end{aligned}$$

Proof. The pointwise lower bound $\omega_T^\pi \geq (\omega_T - h)_+$ follows from the budget constraint and nonnegativity of haircuts, as in Proposition C.1. Monotonicity of risk measures implies that if $X \geq Y$ almost surely, then $\text{VaR}_\alpha(X) \geq \text{VaR}_\alpha(Y)$ and $\text{ES}_\alpha(X) \geq \text{ES}_\alpha(Y)$. For the identities, observe that $x \mapsto (x - h)_+$ is nondecreasing; hence quantiles shift: $\text{VaR}_\alpha((X - h)_+) = \max\{\text{VaR}_\alpha(X) - h, 0\}$. The ES identity follows from the Kusuoka representation [75] $\text{ES}_\alpha(Z) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_u(Z) du$ applied to $Z = (X - h)_+$. $\square$

Implication for severity design. For a random budget $H = \theta_n D_T$, apply Theorem C.5 conditionally on H to conclude that Queue minimizes the conditional VaR/ES of the top survivor at every tail level. When $\alpha_+ > 1$ (finite mean winners), the tail-equivalence property of regularly varying distributions yields

$$\frac{\text{ES}_u(\omega_T)}{\text{VaR}_u(\omega_T)} \rightarrow \frac{\alpha_+}{\alpha_+ - 1} \quad \text{as } u \uparrow 1,$$

so VaR- and ES-based moral hazard conclusions coincide asymptotically with those of PTSR and PMR.

C.6 Randomized constructions for moral-hazard examples

Extreme-value moral hazard (principal-agent). Fix $\rho \in (0, 1)$ and $k_n = \lfloor \rho n \rfloor$. Draw winner equities $Y_i^{(n)}$ i.i.d. Pareto(α_+) and loser equities $X_i^{(n)}$ with mean μ . Then $M_n^+ = \max Y_i^{(n)} \asymp n^{1/\alpha_+}$ while non-max winners sum to $o(M_n^+)$. Losers sum to $D_T \approx \mu n$. For fixed severity $\theta_n \equiv \bar{\theta}$, the haircut $H_n \approx \theta \mu n$ exceeds the capacity of non-max winners, forcing the top winner to cover the bulk. Post-ADL equity is $(M_n^+ - \bar{\theta} \mu n)_+ \rightarrow 0$ since $n^{1/\alpha_+} \ll n$ for $\alpha_+ < 1$.

Leverage-imbalance construction. Fix leverage masses $\ell_n^- \gg \ell_n^+$. Draw loser equities $X_i^{(n)}$ i.i.d. Pareto(α_-), so $D_T \approx M_n^- \asymp n^{1/\alpha_-}$. Assign winner leverage $c \ell_n^+$ to a random index I_n and distribute the rest evenly. Then $\omega_n \asymp (\ell_n^+ / \ell_n^-) n$. This satisfies Proposition 5.2 assumptions, yielding the claimed threshold.

D Theoretical Properties of Capped Pro-Rata

We formalize the theoretical properties of the capped pro-rata rule. We note that the capped pro-rata algorithm in Algorithm 1 is a standard water-filling algorithm [27]. The most similar known prior work to this appendix is the study of how such algorithms provide sybil resistance in concave games in decentralized systems [71]. One can view our result as a generalization of this result.

Properties of ADL rules. Fix time T , state $\mathcal{P}_n$, winners $\mathcal{W}_T$, and equities $e_i = (e_T(\mathbf{p}_i))_+$. Let $s_{\pi,i} = (1 - h_{\pi,i})e_i$ be the surviving equity of position i under ADL policy π . We say that a feasible ADL rule π (§2.4) satisfies:

1. *Sybil resistance:* Outcomes are invariant to splitting/merging accounts. For any split $e_i = \sum_{a=1}^r z_a$, the sum of survivors equals the original survivor: $\sum_a s_{\pi,a}^{(i \rightarrow z)} = s_{\pi,i}$.
2. *Scale invariance:* $s_{\pi}(ce; \theta_{\pi}) = c s_{\pi}(e; \theta_{\pi})$ for $c > 0$.
3. *Monotonicity:* If $e_1 \geq \dots \geq e_k$, then $s_{\pi,1} \geq \dots \geq s_{\pi,k}$.
4. *Interior regularity:* The map $(e, H) \mapsto s_{\pi}(e; H)$ is C^1 on the interior, i.e., for $e_i > 0$ for all i and $0 < H < \sum_i e_i$.

Collectively, we refer to these as the *fairness properties* for an ADL rule.

Proposition D.1 (Uniqueness of the Pro-Rata Rule). *If a feasible ADL policy π satisfies the fairness properties, then $s_{\pi}(e; H) = s^{\text{PR}}(e; H)$ for all feasible inputs.*

Proof. Fix a feasible budget H (e.g. maximum value of $\theta_n D_T$) and write $\beta_i \in (0, 1]$ for the haircut cap on winner i (as in the capped pro-rata rule (29)), and sort so that $e_1 \geq \dots \geq e_k$. View the haircuts as a function of the realized budget $b \in [0, H]$ and write $h_i(b)$ for the haircut on winner i when total budget b has been allocated. For each b , define the *active set*

$$A(b) = \{i : h_i(b) < \beta_i\} \quad (36)$$

of winners whose caps are not yet binding. Since there are only finitely many caps β_i , there exists a partition $0 = b_0 < b_1 < \dots < b_L = H$ such that on each open interval $I_\ell := (b_{\ell-1}, b_\ell)$ the active set is constant. Fix one such interval $I = (b_-, b_+)$ and write $A = A(b)$ for any $b \in I$. On this interval, by scale invariance on A (per-unit budget increases all active haircuts at the same rate) and feasibility,

$$\frac{dh_i}{db} = \frac{1}{\sum_{j \in A} e_j} \quad (i \in A), \quad \frac{dh_i}{db} = 0 \quad (i \notin A).$$

Using the interior regularity property, we can integrate these terms on I . Integrating from b_- to any $b \in I$ gives the *unconstrained* evolution

$$\tilde{h}_i(b) = \begin{cases} h_i(b_-) + \frac{b - b_-}{\sum_{j \in A} e_j} & (i \in A), \\ h_i(b_-) & (i \notin A). \end{cases}$$

All active coordinates in A move in lockstep, so on I there exists a scalar function $\eta_I(b)$ such that

$$h_i(b) = \min\{\eta_I(b), \beta_i\} \quad (i \in A), \quad h_i(b) = h_i(b_-) \quad (i \notin A).$$

The budget identity can then be written on I as

$$b = \sum_{i \notin A} e_i h_i(b_-) + \sum_{i \in A} e_i \min\{\eta_I(b), \beta_i\},$$

which, for fixed $b \in I$, is a continuous strictly increasing function of $\eta_I(b)$ as long as $A \neq \emptyset$. Thus for each $b \in I$ there is a unique $\eta_I(b)$ solving the budget identity. In particular, on the *interior* interval where no caps bind ($A = \{1, \dots, k\}$ and $b \in (0, W_T)$), we have $h_i(b) = \eta_I(b)$ for all i so the budget identity reduces to

$$b = \sum_{i=1}^k e_i h_i(b) = \eta_I(b) \sum_{i=1}^k e_i = \eta_I(b) W_T,$$

which implies $\eta_I(b) = b/W_T$. Thus on this interval $s_{\pi,i} = (1 - \eta_I(b))e_i = (1 - b/W_T)e_i$, i.e., proportional to equity.

Sybil resistance implies allocations depend only on total equity: splitting $e_i = \sum_a z_a$ leaves $\sum_a z_a h_a(b)$ and hence the survivor $\sum_a (1 - h_a(b))z_a$ unchanged for each b , so the proportional form on interior intervals is preserved under arbitrary splits. Monotonicity further restricts how indices can exit the active set as b increases. At an endpoint b_ℓ where $\eta_I(b_\ell)$ first hits some cap β_m , all indices $j \geq m$ with $\beta_j \leq \beta_m$ must saturate together; otherwise we would have $e_m \geq e_j$ but $(1 - h_m(b_\ell))e_m < (1 - h_j(b_\ell))e_j$, violating monotonicity. Thus, as we pass from I_ℓ to $I_{\ell+1}$, a (possibly empty) tail $\{j > m\}$ leaves A , contributing a fixed amount $\sum_{j>m} e_j \beta_j$ to the budget, and the same water-filling argument applies on the remaining head with reduced budget.

Concatenating the solution across all intervals I_ℓ yields the reverse-waterfilling form $h_i = \min\{\eta^*, \beta_i\}$ with η^* chosen so that $\sum_i e_i h_i = H$, which is exactly capped pro-rata. Uniqueness follows either from the strict convexity of the Euclidean projection onto $\{h \in [0, 1]^k : \sum_i e_i h_i = H\}$ or from the monotone one-dimensional search that defines η^* . $\square$

Convex optimality We now formalize the convex-welfare interpretation of capped pro-rata from §6. Fix time T , winners $\mathcal{W}_T$ with equities $e_i = e_T(\mathbf{p}_i)_+$, and effective caps $\beta_i = \min\{\bar{h}_i, 1 - e_i/e_T(\mathbf{p}_i)\}$ defined by the haircut and equity constraints (22)–(23). Let $B_T = \theta_\pi D_T(\mathcal{P}_n)$ denote the haircut budget from (30), and let $\phi : [0, 1] \rightarrow \mathbf{R}$ be a strictly convex increasing function representing per-unit haircut disutility as in (31). We consider choosing haircuts $h = (h_i)_{i \in \mathcal{W}_T}$ to minimize the equity-weighted total disutility $\sum_{i \in \mathcal{W}_T} e_i \phi(h_i)$ subject to the per-account bounds $0 \leq h_i \leq \beta_i$ and the aggregate budget constraint $\sum_{i \in \mathcal{W}_T} e_i h_i = B_T$.

Proposition D.2 (Convex optimality). *For any strictly convex increasing ϕ , the unique solution to*

$$\min_h \sum_{i \in \mathcal{W}_T} e_i \phi(h_i) \quad \text{s.t.} \quad \sum_{i \in \mathcal{W}_T} e_i h_i = B_T, \quad 0 \leq h_i \leq \beta_i$$

is the capped pro-rata rule (29), i.e.,

$$h_{\pi_{CP}, i}^* = \min\{\eta^*, \beta_i\},$$

where η^ satisfies $\sum_{i \in \mathcal{W}_T} e_i h_{\pi_{CP}, i}^* = B_T$.*

Proof. The optimization problem is convex with a strictly convex objective and linear constraints, so any point satisfying the Karush–Kuhn–Tucker (KKT) conditions is the unique global minimizer [27, Ch. 5]. For B_T in the interior of the feasible region ($0 < B_T < C(\beta)$, where $C(\beta)$ is defined in (28)), Slater’s condition holds, so KKT conditions are necessary and sufficient. The Lagrangian is

$$\mathcal{L}(h, \lambda, \mu, \nu) = \sum_{i \in \mathcal{W}_T} e_i \phi(h_i) + \lambda \left(\sum_{i \in \mathcal{W}_T} e_i h_i - B_T \right) + \sum_{i \in \mathcal{W}_T} \mu_i (h_i - \beta_i) - \sum_{i \in \mathcal{W}_T} \nu_i h_i,$$

with multipliers $\lambda \in \mathbf{R}$ and $\mu_i, \nu_i \geq 0$. Stationarity with respect to h_i gives

$$e_i \phi'(h_i) + \lambda e_i + \mu_i - \nu_i = 0 \quad (i \in \mathcal{W}_T),$$

together with complementary slackness $\mu_i (h_i - \beta_i) = 0$ and $\nu_i h_i = 0$. For any index i with $0 < h_i < \beta_i$, we must have $\mu_i = \nu_i = 0$, so $\phi'(h_i) = -\lambda$. Because ϕ' is strictly increasing, this implies $h_i = c$ for some common scalar c independent of i . If $h_i = \beta_i$ then $\mu_i \geq 0$ and $\nu_i = 0$, and if $h_i = 0$ then $\nu_i \geq 0$ and $\mu_i = 0$, so in all cases the KKT conditions imply the water-filling form

$$h_i = \min\{c, \beta_i\} \quad (i \in \mathcal{W}_T).$$

The budget constraint $\sum_{i \in \mathcal{W}_T} e_i h_i = B_T$ then reduces to finding c such that

$$\sum_{i \in \mathcal{W}_T} e_i \min\{c, \beta_i\} = B_T.$$

The left-hand side is continuous and strictly increasing in c on $[0, 1]$ as long as some $\beta_i > 0$, so there is a unique $c = \eta^*$ solving this equation. Thus every KKT point has the capped pro-rata form $h_i = \min\{\eta^*, \beta_i\}$ with η^* determined by the budget, which is exactly the rule in (29). Strict convexity of the objective implies this KKT point is the unique global minimizer, proving the claim (see also the standard water-filling derivations in [27, Ch. 5]). $\square$

E Algorithms for Pro-Rata Haircut Rules

This appendix collects the explicit water-filling procedures for both the capped pro-rata rule from Section 6 and the Risk-Aware Pro-Rata (RAP) rule from Section 7. Each algorithm takes the winner equities $\{e_i\}$, effective caps $\{\beta_i\}$, and the target budget $B_T = \theta_\pi D_T(\mathcal{P}_n)$, returning the optimal haircut vector h or declaring infeasibility if the aggregate capacity $\sum_i e_i \beta_i$ is insufficient.

E.1 Capped Pro-Rata Water-Filling

The procedure below enforces the constraints in [Equations \(29\)](#) and [\(30\)](#) by leveling caps until the target budget is met.

Algorithm 1 Capped Pro-Rata Haircut Allocation (Water-Filling)

Require: Winner equities $e = \{e_1, \dots, e_n\}$, Effective caps $\beta = \{\beta_1, \dots, \beta_n\}$, Target budget B_T

Ensure: Haircut vector $h = \{h_1, \dots, h_n\}$ or **Infeasible**

```

1:  $C \leftarrow \sum_{i=1}^n e_i \beta_i$  ▷ Compute total maximum capacity
2: if  $B_T > C$  then
3:   return Infeasible
4: end if
5: Sort indices  $p$  such that  $\beta_{p_1} \leq \beta_{p_2} \leq \dots \leq \beta_{p_n}$ 
6:  $\beta_{p_0} \leftarrow 0$ 
7:  $V \leftarrow 0$  ▷ Cumulative value covered
8:  $R \leftarrow \sum_{i=1}^n e_i$  ▷ Remaining uncapped equity mass
9: for  $k = 1$  to  $n$  do
10:   $\Delta\beta \leftarrow \beta_{p_k} - \beta_{p_{k-1}}$ 
11:   $V_{\text{step}} \leftarrow \Delta\beta \cdot R$ 
12:  if  $V + V_{\text{step}} \geq B_T$  then
13:     $\eta \leftarrow \beta_{p_{k-1}} + (B_T - V)/R$  ▷ Found the water level  $\eta$ 
14:    break
15:  end if
16:   $V \leftarrow V + V_{\text{step}}$ 
17:   $R \leftarrow R - e_{p_k}$  ▷ User  $p_k$  becomes fully capped
18: end for
19: if  $V < B_T$  then ▷ Handling numerical edge cases
20:   $\eta \leftarrow \beta_{p_n}$ 
21: end if
22: for  $i = 1$  to  $n$  do
23:   $h_i \leftarrow \min\{\eta, \beta_i\}$ 
24: end for
25: return  $h$ 

```

E.2 Risk-Aware Pro-Rata Water-Filling

The RAP algorithm augments the capped procedure by prioritizing accounts according to their “cap-to-weight” ratios β_i/w_i .

F Risk-Aware Pro-Rata (RAP)

F.1 Examples of Risk-Aware Pro-Rata (RAP) and Next Deficit

In this section, we provide detailed numerical examples illustrating the properties of the RAP rule and the impact of post-haircut shocks.

RAP weighting example. We illustrate the three choices of risk models on an example at $T = 2$ where the winners are $\mathcal{W}_2 = \{A, C, E\}$ with effective leverages $\lambda_{A,2}^+ \approx 1.031$, $\lambda_{C,2}^+ \approx 0.925$,

Algorithm 2 Risk-Aware Pro-Rata Haircut Allocation (Weighted Water-Filling)

Require: Winner equities $e = \{e_1, \dots, e_n\}$, Effective caps $\beta = \{\beta_1, \dots, \beta_n\}$, Risk weights $w = \{w_1, \dots, w_n\}$, Target budget B_T

Ensure: Haircut vector $h = \{h_1, \dots, h_n\}$ or **Infeasible**

```

1:  $C \leftarrow \sum_{i=1}^n e_i \beta_i$  ▷ Total capacity
2: if  $B_T > C$  then
3:   return Infeasible
4: end if
5: Compute ratios  $r_i \leftarrow \beta_i / w_i$  for all  $i$  (treat  $w_i = 0$  as  $r_i = \infty$ )
6: Sort indices  $p$  such that  $r_{p_1} \leq r_{p_2} \leq \dots \leq r_{p_n}$ 
7:  $r_{p_0} \leftarrow 0$ 
8:  $V \leftarrow 0$  ▷ Cumulative value covered
9:  $W_{\text{rem}} \leftarrow \sum_{i=1}^n e_i w_i$  ▷ Remaining weighted equity mass
10: for  $k = 1$  to  $n$  do
11:    $\Delta\tau \leftarrow r_{p_k} - r_{p_{k-1}}$ 
12:    $V_{\text{step}} \leftarrow \Delta\tau \cdot W_{\text{rem}}$ 
13:   if  $V + V_{\text{step}} \geq B_T$  then
14:      $\tau \leftarrow r_{p_{k-1}} + (B_T - V) / W_{\text{rem}}$  ▷ Found the scaling factor  $\tau$ 
15:     break
16:   end if
17:    $V \leftarrow V + V_{\text{step}}$ 
18:    $W_{\text{rem}} \leftarrow W_{\text{rem}} - e_{p_k} w_{p_k}$  ▷ User  $p_k$  becomes fully capped
19: end for
20: if  $V < B_T$  then ▷ Numerical edge case
21:    $\tau \leftarrow r_{p_n}$ 
22: end if
23: for  $i = 1$  to  $n$  do
24:    $h_i \leftarrow \min\{\beta_i, \tau w_i\}$ 
25: end for
26: return  $h$ 

```

and $\lambda_{E,2}^+ \approx 1.548$. Recall that under pro-rata, the normalized shares are $s_i^{\text{PR}} \propto e_{T,i}$, yielding $(s_A^{\text{PR}}, s_C^{\text{PR}}, s_E^{\text{PR}}) \approx (0.163, 0.728, 0.109)$. For RAP with $w_i = \lambda_i g(\lambda_i)$, the shares allocate proportional to $e_{T,i} w_i$. The resulting shares (order A, C, E) are:

- **Linear** $g(\lambda) = \lambda$: $\approx (0.164, 0.589, 0.246)$.
- **Power** $g(\lambda) = \lambda^2$: $\approx (0.155, 0.498, 0.348)$.
- **CVaR** $g(\lambda) = (\lambda - 0.9)_+$: $\approx (0.149, 0.114, 0.737)$.

RAP shifts haircut mass toward high-leverage winners; the tilt is mild for linear g , stronger for λ^2 , and concentrates almost entirely on the over-threshold tail for the CVaR model.

Next deficit and leverage sensitivity. Consider the setup where $T = 2$ with $D_T \approx 0.705$ and $W_T \approx 7.72$. Under the normal pro-rata rule, the haircut rate is $h_T^{\text{PR}} \approx 0.0913$. We consider a simple Markovian shock whose direction is uniformly random and whose magnitude grows with the winner leverage mass:

$$Z_{T,i} = \xi_T \zeta_T, \quad \xi_T \in \{-1, +1\} \text{ equiprobable}, \quad \zeta_T = \alpha \frac{\ell_T^+}{k_T}.$$

With $\alpha = 1.2$, the expected next deficit is:

$$\mathbf{E}[D_{T+1}^{\text{next}} \mid \mathcal{F}_T] \approx 1.46 > D_T.$$

This illustrates a failure mode for pro-rata when the shock kernel scales strongly with leverage: pro-rata shrinks all winners uniformly and leaves effective leverages $\lambda_{T,i}$ unchanged, so the shock magnitude ζ_T is unaffected while residual exposure remains large on high-leverage winners.

Correlated shocks example. Consider two winners with equal equity e and leverage levels $\lambda_{t,1} > \lambda_{t,2} \geq 1$, and budget $b_t = 2\epsilon e$. Assume price shocks are AR(1): $Z_{t+2} = \rho Z_{t+1} + \varepsilon_{t+2}$, $\rho \in (0, 1)$. RAP with $w_{t,i} \propto \lambda_{t,i} \psi(1/\lambda_{t,i})$ puts more haircut mass on account 1 (the higher leverage account). Assume account 1's exposure to Z_{t+1} offsets the loss term (a "hedge") in the next step deficit: $D_{t+1}^{\text{next}} = (\alpha - \beta s_{t,1}) Z_{t+1}$. Shrinking $s_{t,1}$ (winner 1's residual equity) weakens the hedge into $t+2$. Since $s_{t,1}^{\text{RAP}} < s_{t,1}^{\text{PR}}$, the two-step sum of deficits S_t can satisfy $S_t(h_t^{\text{PR}}) < S_t(h_t^{\text{RAP}})$, even if RAP minimizes the one-step deficit.

Exchange incentive compatibility example. Consider two positions $\mathbf{p}_A$ (high value, $\theta_A = 100$) and $\mathbf{p}_B$ (low value, $\theta_B = 5$) with equal initial leverage $\ell = 1$. Suppose the exchange must reduce total leverage by 1 unit.

- **RAP (Targeted):** Fully liquidates $\mathbf{p}_A$. Continuation value: 5β .
- **Pro-Rata:** Reduces both by 50%. Continuation value: 52.5β .

Pro-rata yields significantly higher utility by preserving the high-value trader, demonstrating that RAP need not be incentive compatible for the exchange.

F.2 RAP Optimality and Convex Dominance

In this appendix, we prove two statements: 1) RAP optimizes the one-step next deficit (32) and 2) RAP has a smaller residual than any other comonotone haircut rule.

RAP optimizes δ_T . In this section, we briefly show that RAP optimizes the one-step deficit proxy δ_T . We do this by showing that the weights determined by the perspective transform $\rho(\lambda)$, which define g^* , optimize δ_T .

Proposition F.1. *Fix a round t with budget $b_t = \theta_t |D_t|$ and per-account caps $0 \leq H_{t,i} \leq 1$. For*

$$\delta_t(h) = \sum_{i \in W_t} (1 - h_{t,i}) \lambda_{t,i} e_{t,i} \psi\left(\frac{1}{\lambda_{t,i}}\right)$$

the capped reverse-waterfilling with weights $w_{t,i} = \rho(\lambda_{t,i}) = \lambda_{t,i} \psi(1/\lambda_{t,i})$ minimizes $\delta_t(h)$ among all h with $\sum_i e_{t,i} h_{t,i} = b_t$ and $0 \leq h_{t,i} \leq H_{t,i}$.

Proof. Using $\rho(\lambda) = \lambda \psi(1/\lambda)$,

$$\delta_t(h) = \sum_i (1 - h_{t,i}) \lambda_{t,i} e_{t,i} \psi(1/\lambda_{t,i}) \equiv C_h - \sum_i \rho(\lambda_{t,i}) e_{t,i} h_{t,i}$$

where C_h is a constant independent of h (can depend on $\lambda_{t,i}$). Maximizing $\sum_i \rho(\lambda_{t,i}) e_{t,i} h_{t,i}$ under $\sum_i s_{t,i} = b_t$ and $0 \leq s_{t,i} \leq e_{t,i} H_{t,i}$ is a fractional knapsack problem solved by sorting the values $\rho(\lambda_{t,i})$. The optimizer for this problem is reverse-waterfilling [27]:

$$h_{t,i} = \min\{H_{t,i}, \tau_t^* w_{t,i}\}, \quad w_{t,i} = \rho(\lambda_{t,i}),$$

with τ_t^* set by $\sum_i e_{t,i} h_{t,i} = b_t$. This choice of w minimizes δ_t among all weighted reverse-waterfilling rules. $\square$

RAP realizes Schur-convex dominance. We first note that theoretical results from the measure theoretic literature imply that RAP should provide Schur-convex dominance. RAP weights $w_i = \lambda_{T,i} g(\lambda_{T,i})$ can be interpreted as allocating budget proportional to a coherent, law-invariant risk measure. Specifically, let $\rho(\lambda) = \lambda \psi(1/\lambda)$ be a risk density with ψ convex and nonincreasing. This corresponds to a spectral risk measure $\rho(X) = \int_0^1 \text{ES}_\alpha(X) d\mu(\alpha)$ via the Kusuoka representation [1, 75]. Choosing $g(\lambda) = \rho(\lambda)/\lambda$ aligns the RAP allocation with this spectral density. Known results from [75] then imply RAP with a “more convex” risk density (in the Schur sense) will weakly Schur-dominate any other weighted pro-rata rule with a less concentrated density. Instead of utilizing such strong measure theoretic tools, we instead prove this directly below in an elementary manner.

Theorem F.2 (Constructive Schur-convex dominance). *Fix T , budget b_T , and caps (β_i) . Let ρ be nondecreasing and $g^*(\lambda) = \rho(\lambda)/\lambda$. Consider any weighted pro-rata rule $h^{(w)}$ with weights w . If the haircut share vector of $h^{(w)}$ is no more concentrated on high- ρ indices than that of $\text{RAP}(g^*)$ on any fixed active set, then the residual vector of $\text{RAP}(g^*)$ weakly submajorizes that of $h^{(w)}$:*

$$z_T(\text{RAP}(g^*)) \preceq_w z_T(h^{(w)}).$$

Thus, for any convex increasing ϕ , $\sum \phi(z_{T,i}(\text{RAP})) \leq \sum \phi(z_{T,i}(h^{(w)}))$.

Proof. We split this proof into three steps. The first step is analogous to the proof of Proposition D.1, where we analyze the change to the weights on a piecewise constant set of intervals. The second step is to analyze how the residuals change with budget using the piecewise-constant representation. Given the change in residual, the final step uses the majorization inequality to show that the residual vector of RAP weakly submajorizes that of any other weighted pro-rata rule.

Step 1: Active set and parameterization. We first define the active set $A(b)$ as in (36). Any weighted reverse-waterfilling with weights w admits the water-level form

$$h_{T,i}^{(w)}(b) = \min\{\beta_i, \tau^{(w)}(b) w_i\},$$

where $\tau^{(w)}(b)$ is chosen so that $\sum_i e_{T,i} h_{T,i}^{(w)}(b) = b$. On an interval $[b_0, b_1]$ with fixed $A(b)$ we have

$$\frac{db}{d\tau^{(w)}} = \sum_{j \in A(b)} e_{T,j} w_j, \quad \frac{dh_{T,i}^{(w)}}{db} = \begin{cases} \frac{w_i}{\sum_{j \in A(b)} e_{T,j} w_j}, & i \in A(b), \\ 0, & i \notin A(b). \end{cases}$$

Step 2: Residual dynamics. Next we look at how the residuals (e.g. post haircut equity) change with budget. Write the reweighted residuals as

$$z_{T,i}(b) = \rho(\lambda_{T,i}) e_{T,i} (1 - h_{T,i}(b)).$$

Then on $[b_0, b_1]$,

$$\frac{dz_{T,i}^{(w)}}{db} = \begin{cases} -\frac{\rho(\lambda_{T,i}) e_{T,i} w_i}{\sum_{j \in A(b)} e_{T,j} w_j}, & i \in A(b), \\ 0, & i \notin A(b). \end{cases}$$

For $\text{RAP}(g^*)$ we take $w_i = \rho(\lambda_{T,i})$, giving

$$\frac{dz_{T,i}^{(\text{RAP})}}{db} = \begin{cases} -\frac{\rho(\lambda_{T,i})^2 e_{T,i}}{\sum_{j \in A(b)} e_{T,j} \rho(\lambda_{T,j})}, & i \in A(b), \\ 0, & i \notin A(b). \end{cases}$$

Step 3: Majorization at each budget. On a fixed $A(b)$, sort indices by decreasing $\rho(\lambda_{T,i})$. By the hypothesis that the haircut share vector of $h^{(w)}$ on $A(b)$,

$$\sigma_i^{(w)}(b) := \frac{e_{T,i} w_i}{\sum_{j \in A(b)} e_{T,j} w_j},$$

is no more concentrated on high- ρ indices than the RAP share $\sigma_i^{(\text{RAP})}(b) = \frac{e_{T,i} \rho(\lambda_{T,i})}{\sum_{j \in A(b)} e_{T,j} \rho(\lambda_{T,j})}$, the rearrangement/majorization inequality implies that for every k ,

$$\sum_{i \leq k} \frac{dz_{T,(i)}^{(\text{RAP})}}{db} \leq \sum_{i \leq k} \frac{dz_{T,(i)}^{(w)}}{db}$$

where (i) denotes the order by decreasing ρ . Hence the instantaneous decrease vector under RAP weakly submajorizes that under $h^{(w)}$ on $[b_0, b_1]$. Integrating over b preserves $\prec_w$ on the interval, and concatenating the finitely many intervals where $A(b)$ changes preserves the order overall: $z_T(\text{RAP}(g^*)) \preceq_w z_T(h^{(w)})$. Schur-convexity then yields the separable convex loss comparison. $\square$

F.3 Example of the Solvency vs. Long-Term Revenue Trade-Off

This example illustrates a fundamental tension: the risk-minimizing policy (RAP) may be suboptimal for the exchange's long-term value (LTV) because it disproportionately liquidates high-leverage users who generate the most fees. Under certain conditions, a "fairer" policy like Pro-Rata (PR), which preserves these high-value users, yields higher total utility for the exchange.

Setup. Consider an exchange with two profitable users $i \in \{H, L\}$. User H is high-leverage ($\lambda_H > \lambda_L$) and high-revenue; user L is safer but generates less fee volume. The exchange must raise a budget b via haircuts $h = (h_H, h_L)$ to cover a deficit. Its objective combines immediate safety (minimizing insolvency risk) and future revenue (LTV):

$$U^{\text{exch}}(h) = \underbrace{-\text{Loss}(h)}_{\text{Immediate Safety}} + \underbrace{\beta \sum_{i \in \{H, L\}} \theta_i (1 - h_i) \lambda_i}_{\text{Future Revenue (LTV)}}$$

where $\text{Loss}(h) = L_0 - \alpha_H h_H - \alpha_L h_L$ is the expected insurance fund draw, and $\theta_i \lambda_i$ is the expected future fee revenue per unit of equity from user i . We assume $\alpha_H/e_H > \alpha_L/e_L$, meaning user H provides the cheapest risk reduction per dollar of haircut.

Policy comparison. We compare two policies:

- *RAP (Risk-Minimizing):* Prioritizes risk reduction above all. Since H offers the best “bang for the buck” in safety ($\alpha_H/e_H > \alpha_L/e_L$),¹⁵ RAP fully targets H first: $h^{\text{RAP}} = (b/e_H, 0)$ (assuming $b < e_H$).
- *Pro-Rata (Revenue-Preserving):* Spreads the pain evenly, setting $h_i^{\text{PR}} = \frac{b}{e_H + e_L}$ for both users. This is less efficient for immediate safety but preserves more of user H ’s position.

When an exchange prefers Pro-Rata to maximize long-term revenue. The exchange prefers PR over RAP when the LTV gain from saving user H outweighs the increased immediate risk. The utility difference is:

$$\Delta U = U^{\text{exch}}(h^{\text{PR}}) - U^{\text{exch}}(h^{\text{RAP}}) = \underbrace{\Delta h_H (\beta \theta_H \lambda_H - \alpha_H)}_{\text{Gain from saving } H} - \underbrace{h_L^{\text{PR}} (\alpha_L - \beta \theta_L \lambda_L)}_{\text{Cost of cutting } L}.$$

If user H is sufficiently profitable (θ_H is large), then $\Delta U > 0$. Specifically, PR dominates RAP if the relative revenue of the high-leverage user exceeds a threshold:

$$\frac{\theta_H}{\theta_L} \geq \Theta^* = \frac{h_L^{\text{PR}}}{\Delta h_H} \cdot \frac{\lambda_L}{\lambda_H} + \frac{\alpha_H - \frac{h_L^{\text{PR}}}{\Delta h_H} \alpha_L}{\beta \lambda_H \theta_L}.$$

While RAP is “optimal” for preventing immediate insolvency, it can be myopic. If high-leverage traders are the exchange’s cash cows, the exchange has a rational incentive to use Pro-Rata to keep them active, even at the cost of slightly higher short-term risk.

G Stackelberg Control

In this appendix, we formalize the results of §8, where we model ADL as a Stackelberg game.

¹⁵The marginal reduction in loss per unit of budget is $\frac{\partial \text{Loss}}{\partial h_i} \frac{dh_i}{d(\text{budget})} = \frac{\alpha_i}{e_i}$. Since user H has higher leverage, they have a higher risk coefficient α_H , making α_H/e_H the steepest descent direction for the loss function.

G.1 Opposing Schur orderings for time to solvency and LTV

We formalize the fundamental trade-off between aggressive debt reduction (safety) and trader fee retention (value), formalizing the example in Appendix F.3. Let $n_t = |\mathcal{W}_t|$. Given any feasible strategy π , write $z_t(\pi) \in \mathbf{R}_+^{n_t}$ for the vector of residual debts at time t (sorted in decreasing order). Let $h_{t,i}^\pi \in [0, 1]$ be the haircut fraction for agent i with equity $e_{t,i}$, so that the haircut mass is $m_{t,i}(\pi) = h_{t,i}^\pi e_{t,i}$. We write $m_t(\pi) \in \mathbf{R}_+^{n_t}$ for the corresponding vector of haircut masses. Let $Z_t \in \mathbf{R}_+^{n_t}$ denote the equity shock at time t , following the notation of Section 7. These evolve componentwise as

$$z_{t+1,i}(\pi) = z_{t,i}(\pi) + Z_{t+1,i} - m_{t,i}(\pi),$$

so summing from $\tau = 0$ to $t - 1$ yields the conservation-of-mass identity

$$z_t(\pi) = z_0 + \sum_{\tau=1}^t Z_\tau - \sum_{\tau=0}^{t-1} m_\tau(\pi). \quad (37)$$

This ensures that the equity at time t is either initial equity, was gained or lost in a price shock, or haircut.

Proposition G.1 (Solvency-Revenue Trade-off). *Let A and B be two strategies facing the same shock sequence $(Z_t)_t$. Assume:*

- (i) *Safety Dominance: For all $t < \tau_{\text{solv}}(A)$, strategy A maintains weakly smaller residuals than B in the weak submajorization order: $z_t(A) \preceq_w z_t(B)$.*
- (ii) *Retention Value: Let $M_t(\pi) := \sum_{\tau=0}^{t-1} m_\tau(\pi)$ be the cumulative haircut vector and suppose the lifetime value takes the form*

$$\text{LTV}(\pi) = \sum_t \beta^t G_t(M_t(\pi)),$$

where each stage value $G_t : \mathbf{R}_+^{n_t} \rightarrow \mathbf{R}$ is Schur-concave and coordinate-wise nonincreasing in M_t (more cumulative liquidations in the weak submajorization order reduce exchange LTV).

Then:

- (a) $\tau_{\text{solv}}(A) \leq \tau_{\text{solv}}(B)$ (Strategy A is safer).
- (b) $\text{LTV}(A) \leq \text{LTV}(B)$ (Strategy B generates more value).

Proof. We prove this in two steps.

Step 1: Solvency time. For any $t < \tau_{\text{solv}}(A)$ we have $z_t(A) \preceq_w z_t(B)$ by (i). If B is solvent at some such time t so that $z_t(B) = 0$, then weak submajorization on $\mathbf{R}_+^{n_t}$ forces $z_t(A) = 0$ as well, since the zero vector is minimal in this order. Hence B cannot become solvent strictly before A , and $\tau_{\text{solv}}(A) \leq \tau_{\text{solv}}(B)$ almost surely.

Step 2: LTV. From the conservation-of-mass identity (37) and the fact that z_0 and $(Z_\tau)_\tau$ are common across strategies, the cumulative haircuts $M_t(\pi) := \sum_{\tau=0}^{t-1} m_\tau(\pi)$ satisfy

$$M_t(\pi) = z_0 + \sum_{\tau=1}^t Z_\tau - z_t(\pi),$$

so that

$$M_t(A) - M_t(B) = z_t(B) - z_t(A).$$

For each $t < \tau_{\text{solv}}(A)$, assumption (i) gives $z_t(A) \preceq_w z_t(B)$, and subtracting from the common vector $z_0 + \sum_{\tau=1}^t \xi_\tau$ reverses the weak submajorization order, yielding $M_t(A) \succeq_w M_t(B)$. By Schur-concavity and coordinate-wise monotonicity of each G_t in (ii), this implies $G_t(M_t(A)) \leq G_t(M_t(B))$ for all t , so

$$\text{LTV}(A) = \sum_t \beta^t G_t(M_t(A)) \leq \sum_t \beta^t G_t(M_t(B)) = \text{LTV}(B).$$

□

Examples. We first give a one-step illustration of Proposition G.1 and then show how its hypotheses can fail for the literal queue versus capped pro-rata policies.

Single-period trade-off. Consider a single round with two winners and common initial residuals $z_0 = (1, 1)$. Let strategy A use aggressive haircuts $M_1(A) = (1, 1)$, fully clearing both accounts so that $z_1(A) = (0, 0)$. Let strategy B use milder haircuts $M_1(B) = (1, 0)$, clearing only the first account and leaving $z_1(B) = (0, 1)$. Then $z_1(A) \preceq_w z_1(B)$ and $M_1(A) \succeq_w M_1(B)$, so the hypotheses of Proposition G.1 hold with $T = 1$. For a Schur-concave, coordinate-wise nonincreasing stage value such as

$$G_1(M) := -\|M\|_2^2,$$

we obtain $G_1(M_1(A)) = -2$ and $G_1(M_1(B)) = -1$, so $\text{LTV}(A) \leq \text{LTV}(B)$, exactly exhibiting the safety-versus-value trade-off.

Queue versus capped pro-rata. Now consider the familiar one-round example with two winners, equities $e = (1, 1)$, and budget $b = 1$. Under capped pro-rata we have haircuts $h^{\text{PR}} = (\frac{1}{2}, \frac{1}{2})$, masses $m^{\text{PR}} = (\frac{1}{2}, \frac{1}{2})$, and residuals $z^{\text{PR}} = (\frac{1}{2}, \frac{1}{2})$. Under a queue policy we instead have $h^{\text{Q}} = (1, 0)$, so $m^{\text{Q}} = (1, 0)$ and $z^{\text{Q}} = (0, 1)$. In weak submajorization order, $z^{\text{PR}} \prec_w z^{\text{Q}}$ and likewise $m^{\text{PR}} \prec_w m^{\text{Q}}$. Thus no choice of labels A, B can make pro-rata simultaneously satisfy the safety-dominance condition $z_t(A) \preceq_w z_t(B)$ and the “more cumulative haircuts” condition $M_t(A) \succeq_w M_t(B)$ in Proposition G.1. Moreover, for any Schur-concave, coordinate-wise nonincreasing G (for instance $G(M) = -\|M\|_2^2$), we have $G(m^{\text{Q}}) \leq G(m^{\text{PR}})$, so in this toy case the safer policy (pro-rata) also delivers *higher* user value. This shows that while Proposition G.1 captures a structural opposing-orders phenomenon, its hypotheses do not hold mechanically for every queue versus capped pro-rata comparison.

G.2 Regret Analysis

In this appendix, we unify the regret analysis for the severity (scalar) and MDIC (vector) controllers. We present a master theorem for Online Mirror Descent (OMD) with constraints and then specialize it to our specific settings.

Master mirror descent bound. Consider a sequence of convex loss functions $f_t : \mathcal{X} \rightarrow \mathbb{R}$ on a convex set $\mathcal{X}$. A learner chooses $x_t \in \mathcal{X}$ and updates using Online Mirror Descent (OMD) with a proximal function (Bregman divergence) $D(\cdot \| \cdot)$.

Theorem G.2 (OMD Regret). *Let subgradients be bounded by $\|g_t\|_* \leq G$ in the dual norm, and let the domain diameter be bounded by $D_{\max}^2 \geq \max_x D(x \| x_1)$. With step size $\eta_t = \frac{D_{\max}}{G\sqrt{t}}$, the regret satisfies:*

$$\sum_{t=1}^T f_t(x_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^T f_t(x) \leq 2D_{\max} G \sqrt{T}.$$

When convex constraints $c_t(x) \leq 0$ are imposed, applying OMD to the Lagrangian guarantees the same $O(\sqrt{T})$ regret and $O(T^{-1/2})$ average constraint violation [65].

We refer the interested reader to the standard reference [65] for a detailed proof.

Residual value function and subgradients. We now describe the residual value function, which is the main subject of mirror descent analysis. For each round t , let $\mathcal{W}_t$ be the set of winners with equities $e_{t,i} > 0$, haircut caps $\beta_{t,i} \in [0, 1]$, and weights $w_{t,i} \geq 0$. The residual value function parametrized by a haircut budget b is

$$\tilde{L}_t(b) = \min_{h \in [0,1]^{|W_t|}} \sum_{i \in W_t} (1 - h_i) \lambda_{t,i} e_{t,i} \quad \text{s.t.} \quad \sum_{i \in W_t} e_{t,i} h_i = b, \quad 0 \leq h_i \leq \beta_{t,i}.$$

Let $\tau_t(b)$ be the KKT multiplier of the budget constraint at b .

Lemma G.3 (Convexity and subgradient). *For each t , $b \mapsto \tilde{L}_t(b)$ is convex, nonincreasing, and piecewise-linear on $[0, \bar{b}_t]$; any KKT multiplier $-\tau_t(b) \in \partial \tilde{L}_t(b)$ is a valid subgradient. Moreover $|\tau_t(b)| \leq \max_{i \in W_t} w_{t,i}$.*

Proof. The program is a linear minimization with a right-hand-side parameter b . Sensitivity analysis implies that the optimal value is convex and piecewise-linear in b , and that the negative of the budget multiplier is a subgradient. Complementary slackness yields $h_i = \min\{\beta_{t,i}, \tau_t(b)w_{t,i}\}$ on active coordinates, which bounds $|\tau_t(b)|$ by the maximum weight. $\square$

Lemma G.3 verifies the curvature and bounded-subgradient assumptions required by the master OMD regret theorem (Theorem G.2) hold, so we can directly plug $\tilde{L}_t$ into that framework.

Severity optimization. In the main text, the severity controller selects a scalar severity $\theta_t \in [0, \Theta_t]$ that determines the haircut budget via $b_t = \theta_t D_t$. Optimizing over b_t or θ_t is equivalent; we keep b_t as the decision variable in the appendix for notational convenience. The loss is $f_t(b) = \tilde{L}_t(b)$, which is convex by Lemma G.3. We use the Euclidean divergence $D(x||y) = \frac{1}{2}(x - y)^2$, reducing OMD to projected gradient descent.

Corollary G.4 (Severity Regret). *Let $G = \max_{t,i} w_{t,i}$ be the bound on subgradients (marginal haircut savings). The severity controller achieves:*

$$\text{Regret}_T^{(\text{sev})} \leq 2\bar{b}G\sqrt{T}.$$

Proof. The domain diameter is $\bar{b}$. By Lemma G.3, subgradients correspond to dual variables bounded by the maximum weight G . The result follows directly from Theorem G.2. $\square$

Haircut optimization (MDIC). The MDIC controller optimizes $u_t = (\theta_t, v_t)$, where θ_t is severity and v_t are ranking weights. We use a block-separable divergence $D = D_\phi \oplus D_\Phi$, where D_Φ is the KL-divergence for the weights.

Corollary G.5 (MDIC Regret). *The MDIC controller achieves regret bounded by:*

$$O\left(\sqrt{TD_\phi(\theta^*||\theta_1)} + \sqrt{TD_\Phi(v^*||v_1)}\right).$$

Initializing v_1 as the uniform distribution over the active winners (so every coordinate is strictly positive) minimizes $D_\Phi(v^||v_1)$ to $O(\log |\mathcal{W}_t|)$, yielding better scaling than Euclidean approaches for high-dimensional weight vectors.*

Proof of the $\mathcal{O}(\log |\mathcal{W}_t|)$ claim. The divergence D_Φ induced by the entropy mirror map equals the Kullback–Leibler divergence:

$$D_\Phi(v^\star \| v_1) = \sum_{i \in \mathcal{W}_t} v_i^\star \log \frac{v_i^\star}{v_{1,i}}.$$

Uniform initialization over active winners sets $v_{1,i} = 1/|\mathcal{W}_t|$. Hence

$$D_\Phi(v^\star \| v_1) = \sum_i v_i^\star \log v_i^\star + \log |\mathcal{W}_t|.$$

The first term is the negative Shannon entropy of $v^\star$ and is therefore non-positive, implying $D_\Phi(v^\star \| v_1) \leq \log |\mathcal{W}_t|$. Thus the initialization costs at most $\mathcal{O}(\log |\mathcal{W}_t|)$. $\square$

Linear regret for static policies. The following toy instance shows that any static severity $\bar{\theta}$ can incur $\Omega(T)$ regret. Consider two regimes A (liquidity) and B (stress) alternating every round, with loss

$$f_t(\theta) = \begin{cases} \theta & t \in A \quad (\text{minimize severity to protect revenue}), \\ \bar{\theta} - \theta & t \in B \quad (\text{maximize severity to clear deficits}). \end{cases}$$

The per-round optimizer is $\theta_t^\star = 0$ in A and $\theta_t^\star = \bar{\theta}$ in B , so

$$\sum_{t=1}^T f_t(\theta) - f_t(\theta_t^\star) = \frac{T}{2}\theta + \frac{T}{2}(\bar{\theta} - \theta) = \frac{T}{2}\bar{\theta} = \Omega(T)$$

for any fixed $\theta \in [0, \bar{\theta}]$. By Theorem G.2, OMD (and MDIC) with $\eta_t \propto 1/\sqrt{T}$ attains $\mathcal{O}(\sqrt{T})$ regret on the same sequence.

G.3 Stackelberg vs. Nash in a Two-Round ADL Game

This appendix illustrates how the timing of moves affects equilibrium selection in ADL scenarios. We show that while simultaneous moves (Nash) can result in a coordination failure where no one unwinds, sequential moves (Stackelberg) allow a leader to induce the efficient high-unwind outcome.

Setup. Consider two agents $i \in \{1, 2\}$, each holding one unit of position. At $t = 1$, each agent chooses whether to *unwind* (voluntarily close) their position ($x_i = 1$) or maintain it ($x_i = 0$). Let $X = x_1 + x_2$ be the total volume of voluntary unwinds. We assume that forced ADL occurs at $t = 2$ if this volume is insufficient, i.e., if $X < T$ for some safety threshold $T \in (1, 2]$. Voluntary unwinding incurs a transaction cost $f > 0$. However, if ADL is triggered (because $X < T$), *every* agent suffers an additional penalty cost $c > f$, regardless of their individual choice.

Simultaneous play (Nash). In simultaneous play, there are two pure Nash equilibria and we effectively have a Coordination Game [56]:

- *Coordination failure* (0, 0): If the opponent plays 0, playing 1 results in volume $1 < T$. ADL still triggers, yielding a total cost $f + c > c$. Thus, the best response is 0, making (0, 0) stable.
- *Coordination success* (1, 1): If the opponent plays 1, playing 1 achieves volume $2 \geq T$ at cost f . Since $f < c$, this is preferred to playing 0 (which yields volume $1 < T$ and cost c), making (1, 1) stable.

This creates a coordination problem: agents may get stuck in the inefficient $(0, 0)$ equilibrium where ADL triggers.

Proposition G.6 (Stackelberg Dominance). *In sequential play where Agent 1 moves first, the unique subgame-perfect equilibrium is $(1, 1)$. Agent 2 observes Agent 1 and will match their action (per the logic above). Agent 1 anticipates this: playing 0 leads to $(0, 0)$ with cost c , while playing 1 induces $(1, 1)$ with cost f . Since $f < c$, Agent 1 chooses 1, eliminating the bad equilibrium.*

Numerical example. Let $f = 1$, $c = 5$, and $T = 1.5$. In simultaneous play, $(0, 0)$ is stable because deviating costs $1 + 5 = 6 > 5$. In Stackelberg play, the leader plays 1, knowing the follower will respond with 1 (cost 1) rather than triggering ADL (cost 5). From the exchange's perspective, if ADL losses exceed $2f$, the sequential outcome $(1, 1)$ is strictly preferred as it collects fees and preserves solvency.

G.4 Follower Strategic Responses

We study two types of follower strategic responses. First, we show that pro-rata haircuts can lead to low leverage users responding by leaving an exchange earlier than higher leverage users. This imposes a negative feedback loop as the exchange's remaining users are higher risk. Secondly, we study traders who aim to add liquidity to the exchange to cover a deficit. These traders are speculating on profits that can be made after an ADL event. We show that such users are incentivized to wait long than the exchange solvency time.

G.4.1 Adverse Selection Under Pro-Rata

We model the ADL interaction as a repeated Stackelberg game: in every round t the exchange (leader) moves first and the surviving winners (followers) best respond. Each stage looks like a one-round Stackelberg problem, but the outcomes feed into subsequent rounds through the evolving of winner equities $e_{T,i}$ and the winning set $\mathcal{W}_t$. At the start of round t the exchange publicly commits to a severity/haircut rule (*e.g.* queue, pro-rata, RAP) that maps the realized deficit D_t and winner book $\mathcal{W}_t$ to haircut shares. After observing θ_t and anticipating D_t , each winner chooses whether to keep its position active (accepting the induced haircut) or to exit/migrate, receiving outside option u_0 . Payoffs are $U_i^{(\pi)} = \mu_i - \mathbf{E}[H_{t,i}^{(\pi)}]$; type i exits whenever this falls below u_0 .

A *death spiral* occurs when pro-rata haircuts force the safest (low-leverage) winners to churn first, shrinking W_t , which in turn raises the future haircut share for the remaining winners, triggering additional exits and further eroding liquidity. We show that RAP breaks this feedback loop by tilting the follower game against high-leverage accounts.

Setup. Let i index a profitable trader with effective equity $e_{t,i}$, leverage $\lambda_{t,i}$, and expected per-round utility μ_i . Normalize severity and deficits by $\bar{\theta} = \mathbf{E}[\theta_t]$, $\bar{D} = \mathbf{E}[D_t]$, and write $\bar{W} = \mathbf{E}[W_t]$ for the expected equity mass of winners. Under pro-rata, the haircut share of i is $s_{t,i}^{\text{PR}} = e_{t,i}/W_t$, so the realized haircut mass is

$$H_{t,i}^{\text{PR}} = \theta_t D_t s_{t,i}^{\text{PR}}.$$

For RAP with a nondecreasing weight g , the share becomes

$$s_{t,i}^{\text{RAP}} = \frac{e_{t,i} \lambda_{t,i} g(\lambda_{t,i})}{\sum_{j \in \mathcal{W}_t} e_{t,j} \lambda_{t,j} g(\lambda_{t,j})}, \quad H_{t,i}^{\text{RAP}} = \theta_t D_t s_{t,i}^{\text{RAP}}.$$

Risk-intensity comparator. Define the equity-weighted market risk intensity at round t by

$$\mu_t^{(g)} := \frac{\sum_{j \in W_t} e_{t,j} \lambda_{t,j} g(\lambda_{t,j})}{\sum_{j \in W_t} e_{t,j}}.$$

Lemma G.7 (When RAP burdens a trader less than Pro-Rata). *For any winner i ,*

$$s_{t,i}^{\text{RAP}} \leq s_{t,i}^{\text{PR}} \iff \lambda_{t,i} g(\lambda_{t,i}) \leq \mu_t^{(g)}.$$

If g is strictly increasing and $\lambda_{t,i}$ is strictly below the equity-weighted market average in the g -scale, the inequality is strict.

Proof. Compute

$$\frac{s_{t,i}^{\text{RAP}}}{s_{t,i}^{\text{PR}}} = \frac{e_{t,i} \lambda_{t,i} g(\lambda_{t,i}) / \sum_j e_{t,j} \lambda_{t,j} g(\lambda_{t,j})}{e_{t,i} / \sum_j e_{t,j}} = \frac{\lambda_{t,i} g(\lambda_{t,i})}{\mu_t^{(g)}}.$$

The claim follows immediately. $\square$

Participation thresholds. Each trader type i has a reservation utility u_0 : the expected per-round payoff it can secure outside the ADL venue (*e.g.* by migrating flow to another exchange, posting liquidity in a different product, or simply investing idle cash in a risk-free instrument). We treat $u_0 > 0$ as exogenous and, unless stated otherwise, common across trader types. We note that heterogeneity in reservation utility can be captured by indexing it as $u_{0,i}$ without changing the argument.

A trader remains active only if the ADL-adjusted payoff exceeds this fallback value. We formalize this by defining the net utility under policy π as

$$U_i^{(\pi)} := \mu_i - \mathbf{E}[H_{t,i}^{(\pi)}],$$

and saying that i participates in a given regime iff $U_i^{(\pi)} \geq u_0$. Equivalently, the *participation threshold* is the maximum haircut burden $\mathbf{E}[H_{t,i}^{(\pi)}]$ that keeps i indifferent, namely $\mu_i - u_0$.

Corollary G.8 (Pro-Rata death spiral vs. RAP retention). *Fix a type i and suppose there is a set of rounds of positive probability on which $\lambda_{t,i} g(\lambda_{t,i}) \leq \mu_t^{(g)}$ (so $s_{t,i}^{\text{RAP}} \leq s_{t,i}^{\text{PR}}$ by Lemma G.7). If, over the same distribution of rounds,*

$$\mu_i - \mathbf{E}[\theta_t D_t s_{t,i}^{\text{PR}}] < u_0 \leq \mu_i - \mathbf{E}[\theta_t D_t s_{t,i}^{\text{RAP}}],$$

then the participation constraint fails under pro-rata but holds under RAP: type i exits in the pro-rata regime while remaining under RAP.

Examples and calibration. Consider two winners, L (low leverage) and H (high leverage), who trade for two rounds. Equities are $(e_L, e_H) = (60, 40)$, leverage levels are $(\lambda_L, \lambda_H) = (2, 6)$, expected per-round utilities are $(\mu_L, \mu_H) = (12, 40)$, and the outside option is $u_0 = 2$. Each round the deficit equals the total haircut budget ($\theta_t = 1$) with $D_1 = 40$ and $D_2 = 30$.

Pro-rata. Shares equal equity weights: $s_i^{\text{PR}} = e_i / (e_L + e_H)$, so $s^{\text{PR}} = (0.6, 0.4)$ even though both traders lose the *same haircut factor* $h^{\text{PR}} = H_i^{\text{PR}} / e_i = \theta_t D_t / W_t = 0.4$. Round 1 haircut masses are $H^{\text{PR}} = (24, 16)$ and utilities are $U^{\text{PR}} = (-12, 24)$. Trader L churns because $U_L^{\text{PR}} < u_0$, leaving only H for round 2. With $W_2 = 40$ the next deficit $D_2 = 30$ forces $H_{2,H}^{\text{PR}} = 30$ (i.e., haircut factor 0.75),

giving $U_{2,H}^{\text{PR}} = 10$; solvency is preserved but the equity base has already halved, so any larger D_2 would wipe out the last winner, illustrating the death spiral.

RAP with $g(\lambda) = \lambda$. Weights scale as $e_i \lambda_i^2$, so L 's share collapses to $s_L^{\text{RAP}} = 240/(240 + 1440) \approx 0.14$ and H 's share rises to 0.86. The round 1 haircut factors become $h_L^{\text{RAP}} \approx 0.095$ and $h_H^{\text{RAP}} \approx 0.86$ (masses $H^{\text{RAP}} = (5.7, 34.3)$), yielding $U^{\text{RAP}} = (6.3, 5.7) > u_0$, so both types remain for round 2. With both accounts active in round 2, shares remain tilted toward H and haircuts $(4.3, 25.7)$ keep both traders above u_0 , preventing churn and stabilizing W_t . This concrete two-player, two-round example mirrors the equilibria predicted by Appendix C.6 and the empirical replay in §9: pro-rata drives the safest capital away first, whereas RAP reallocates the burden toward high-leverage accounts and keeps the equity base intact.

G.4.2 Waiting Game and MDIC-NW

Game. When an ADL event creates a deficit D_t , the exchange moves first: it announces the contemporaneous severity θ_t (hence the haircut budget $\theta_t D_t$) together with the liquidity premium κ_t it is willing to pay per unit of external capital.¹⁶ A Backstop Liquidity Provider (BLP) then decides whether to intervene immediately by injecting any $q_t \in [0, D_t]$ units, or to wait $u \geq 1$ additional rounds before posting the same quantity. Waiting exposes the BLP to time discounting (with a factor β^u) and to the future premium schedule $(\theta_{t+u}, \kappa_{t+u})$. The BLP's payoff from intervening in round $t + u$ with size q is $\beta^u q(\theta_{t+u} - \kappa_{t+u})$, while the exchange's objective is to restore solvency before a deadline by ensuring that the entire deficit is filled (so delay harms solvency).

This Stackelberg game involves the exchange posting incentives as leader and BLPs optimally choose a stopping time and quantity as a follower. The setup captures the “waiting game” intuition: unless the contract guarantees non-negative per-unit surplus right now, rational liquidity providers will defer, pushing resolution beyond the solvency window. We note that this resembles other waiting games in Maximal Extractable Value (MEV) auctions that have been studied empirically on live systems [80, 83].

Per-unit surplus and waiting. Let a Backstop Liquidity Provider (BLP) be able to absorb up to D_t units at time t . Write the per-unit liquidity premium as $\kappa_t \geq 0$ (execution cost plus risk per unit, e.g., $\kappa_t = \Gamma_t/D_t$ when the premium is linear in size) and define the per-unit net surplus

$$\delta_t := \theta_t - \kappa_t.$$

If the BLP executes $q \in [0, D_t]$ units at time t , the immediate surplus is $q \delta_t$; deferring the same q to a later time $t + u$ yields discounted surplus $\beta^u q \delta_{t+u}$.

No-wait condition. We say the exchange enforces a per-unit “No-Wait” constraint when $\delta_t \geq 0$ (equivalently, $\theta_t D_t \geq \Gamma_t$ in the linear-premium case).

Lemma G.9 (Per-unit No-Wait implies immediate action). *Suppose $\beta \in (0, 1]$, the exchange enforces $\delta_t \geq 0$, and the net per-unit surplus is nonincreasing over time: $\delta_{t+u} \leq \delta_t$ for all $u \geq 0$ (e.g., when $\theta_{t+u} \leq \theta_t$ and $\kappa_{t+u} \geq \kappa_t$). Then for any $q \in [0, D_t]$, executing q immediately at t weakly dominates waiting:*

$$q \delta_t \geq \beta^u q \delta_{t+u} \quad \forall u \geq 0.$$

In particular, the BLP's optimal stopping time is $\tau^ = t$ and, if capacity allows, $q_t^* = D_t$.*

¹⁶In practice, this occurs via either an increase in the expected payment for users who stake to an insurance fund (e.g. [47]) or add assets to an HLP/LLP style vault [69, 79]

Proof. By assumption, $\delta_{t+u} \leq \delta_t$ and $\beta^u \leq 1$ for all $u \geq 0$, hence $\beta^u q \delta_{t+u} \leq q \delta_t$ for any fixed $q \in [0, D_t]$. Summing over an optimal decomposition of D_t into infinitesimal units yields the stated dominance and the immediate-execution optimality. $\square$

Corollary G.10 (No-Wait bounds solvency recovery). *Let τ_{def} denote the default time and recall the solvency recovery clock τ_{solv} from §8.2. If the exchange enforces $\delta_t \geq 0$ for all rounds between τ_{def} and the first round in which the deficit is zero, then $\tau_{\text{solv}} = \tau_{\text{def}} + 1$. Hence solvency is restored within a single round of the default event.*

Proof. Lemma G.9 implies that at τ_{def} the BLP injects $q_{\tau_{\text{def}}}^* = D_{\tau_{\text{def}}}$, so the entire deficit is covered immediately. Therefore the insurance fund (or deficit buffer) reaches the safety level δ after the same round, yielding $\tau_{\text{solv}} = \tau_{\text{def}} + 1$. $\square$

This result states that if the no-wait condition is enforced, BLPs will inject liquidity such that the solvency time is minimized. In theory, an exchange can use some of its future revenue or profits (*e.g.* via token issuance) to enforce the no-wait constraint.

Proof of Proposition 8.2

Proof. Let $U(\tau, q) = \beta^{\tau-t} q(\theta_\tau - \kappa_\tau)$ be the discounted utility of a BLP who provides liquidity q at time $\tau \geq t$. Here, θ_τ represents the payment per unit of liquidity (derived from the ADL severity) and κ_τ represents the cost per unit (liquidity premium). The BLP solves the optimal stopping problem (see, *e.g.* [92]):

$$\max_{\tau \geq t, q \in [0, D_t]} \mathbf{E}_t [U(\tau, q)].$$

This objective captures the trade-off between acting immediately to capture the current spread versus waiting for potentially higher future premiums, discounted by the time cost of delay. Assuming the BLP is risk-neutral and has capacity $q = D_t$, the condition for immediate stopping at $\tau = t$ is that the immediate payoff exceeds the expected discounted future value:

$$D_t(\theta_t - \kappa_t) \geq \beta \mathbf{E}_t [V_{t+1}(D_{t+1})],$$

where $V_{t+1}(D_{t+1}) = \max_{\tau \geq t+1, q \in [0, D_{t+1}]} \mathbf{E}_{t+1} [U(\tau, q)]$ is the value function from $t + 1$ onwards. Substituting the per-unit surplus $\delta_t = \theta_t - \kappa_t$ and rearranging gives the condition:

$$\theta_t D_t \geq \Gamma_t + \beta \mathbf{E}_t [V_{t+1}(D_{t+1})].$$

In the simplified case where the BLP compares acting now vs. acting at $t + \Delta t$, and assuming capacity constraints are non-binding, the condition reduces to comparing marginal costs. Specifically, if the cost of waiting is strictly positive (due to discounting $\beta < 1$ or decreasing surplus $\delta_{t+\Delta t} < \delta_t$), then the optimal strategy is to act immediately. If, however, θ_t decays rapidly such that $\theta_t D_t > \mathbf{E}[\theta_{t+\Delta t} D_{t+\Delta t}]$, but the premium Γ_t makes immediate action costly, the inequality flips. Rearranging the condition in the proposition statement:

$$\theta_t D_t + \Gamma_t \leq \mathbf{E}_t [\theta_{t+\Delta t} D_{t+\Delta t}]$$

implies that the future expected payout (even after discounting) is higher than the current payout net of costs, incentivizing delay. Thus, to enforce no-wait, the exchange must ensure the reverse inequality holds. $\square$

H Price of Anarchy Phase Transitions

We characterize the efficiency gap between static (Nash) and dynamic (Stackelberg) ADL policies via Price of Anarchy (PoA) phase transitions. We analyze two distinct welfare objectives: *Fairness* (minimizing haircuts to winners) and *Revenue* (maximizing exchange value). In both cases, we find a sharp transition from a bounded regime, where static policies are constant-factor optimal, to an unbounded regime, where dynamic control is strictly necessary. Throughout, we work with the terminal book $\mathcal{P}_n$ of size n as $n \rightarrow \infty$.

H.1 Fairness Phase Transition

We first analyze fairness using the Profitability-to-Total-Solvency Ratio (PTSR), which measures the survival rate of winner equity relative to the total deficit covered.

H.1.1 Welfare and Assumptions

Fairness welfare. We define the fairness welfare of a policy π as its expected PTSR:

$$W_{\text{Fair}}(\pi) := \mathbf{E} \left[\frac{\omega_T^\pi}{D_T^\pi} \right].$$

This metric captures the efficiency of haircuts: higher values imply that the policy covers deficits D_T^π while preserving maximal winner equity ω_T^π . Extreme-value scaling implies W_{Fair} scales inversely with the severity load.

Price of anarchy. We compare the welfare of a static policy π^{stat} (simultaneous move) against the optimal dynamic Stackelberg policy π^* . The Price of Anarchy is defined as the ratio:

$$\text{PoA}_{\text{Fair}} := \frac{W_{\text{Fair}}(\pi^*)}{W_{\text{Fair}}(\pi^{\text{stat}})}.$$

We require the following regularity assumptions. *LLN and EV Scaling* (Prop. [H.1](#)) establish the baseline scales for deficits ($O(n)$) and equity ($O(b_n)$). *Anti-concentration* (Prop. [H.2](#)) ensures that $b_n \ll n$, creating the scarcity of winner equity that drives the phase transition.

Proposition H.1 (LLN and EV scaling). *The number of winners k_n and losers m_n satisfy $k_n, m_n \asymp n$. Winner equity and loser deficits satisfy extreme-value limits with scales $b_n := b_{k_n}^+$ and $b_{m_n}^-$.*

Proposition H.2 (Anti-concentration). *(i) Equity: The top winner is not dominant: $b_n = o(n)$. (ii) Leverage: Max leverage scales with average leverage: $\max_i \lambda_{i,T}^\pm \leq C\ell_n^\pm/n$. (iii) Balance: Winner and loser leverage masses are comparable: $\ell_n^+ \asymp \ell_n^-$.*

H.1.2 Phase Transition

The efficiency of static ADL depends on the load $\kappa_n = \theta_n n / b_n$, which measures the severity intensity relative to the tail of the winner distribution.

Theorem H.3 (Fairness PoA Phase Transition). *Suppose Assumptions [H.1–H.2](#) hold. Let π^* be the optimal dynamic policy and π^{stat} be any static policy with load κ_n^{stat} .*

1. Bounded Regime (Low Load): If $\sup_n \kappa_n^{\text{stat}} < \infty$, then static ADL is constant-factor optimal:

$$\limsup_{n \rightarrow \infty} \text{PoA}_{\text{Fair}} \leq C < \infty.$$

2. Unbounded Regime (High Load): If $\kappa_n^{\text{stat}} \rightarrow \infty$ (e.g., due to heavy tails or fixed severity with $b_n = o(n)$), then the Price of Anarchy diverges:

$$\text{PoA}_{\text{Fair}} \asymp \kappa_n^{\text{stat}} \rightarrow \infty.$$

Proof. We prove the result by establishing the scaling limits of the dynamic benchmark, the static cost, and their asymptotic ratio.

First, for the dynamic benchmark, the optimal Stackelberg policy π^* minimizes haircuts ex-post to cover the realized deficit $D_T \asymp n$. By targeting the haircut to preserve the top winner's equity $\omega_T \asymp b_n$, the dynamic policy achieves a Profitability-to-Total-Solvency Ratio (PTSR) that concentrates around a constant, $W_{\text{Fair}}(\pi^*) \asymp 1$ (Theorem C.3). This is possible because the controller can observe the realization of heavy-tailed variates and adjust the severity θ precisely to the minimal necessary level.

Second, for the static policy, the severity θ_n is fixed ex-ante and applies a uniform pressure $\theta_n n$ on the book, which must be absorbed by individual winners with capacity scaling as b_n . Under heavy-tailed scaling, Proposition 5.1 shows that this mismatch leads to a welfare decay $W_{\text{Fair}}(\pi) \asymp b_n/(\theta_n n) = 1/\kappa_n$, as the policy blindly destroys small winners or fails to extract sufficient capital from the tail without excessive rates. Crucially, as $n \rightarrow \infty$, the gap between the aggregate load $O(n)$ and individual capacity $O(b_n)$ widens (since $b_n = o(n)$), forcing κ_n to grow if θ_n does not vanish rapidly enough.

Finally, combining these estimates yields the Price of Anarchy $\text{PoA}_{\text{Fair}} \asymp 1/(1/\kappa_n^{\text{stat}}) = \kappa_n^{\text{stat}}$. In the bounded regime ($\kappa_n^{\text{stat}} = O(1)$), the ratio is constant, but in the unbounded regime ($\kappa_n^{\text{stat}} \rightarrow \infty$), the efficiency of the static policy collapses relative to the dynamic optimum, proving the divergence. $\square$

Example H.4 (Light-tailed Failure). If winner equities are sub-Gaussian ($b_n \asymp \sqrt{\log n}$) but the exchange uses fixed severity $\theta > 0$, then $\kappa_n^{\text{stat}} \asymp n/\sqrt{\log n} \rightarrow \infty$. Static ADL unnecessarily destroys winner equity compared to a dynamic policy that scales $\theta \sim 1/n$, leading to infinite PoA.

H.2 Revenue Phase Transition

We now extend the analysis to the *Revenue* objective, formalizing the trade-off between solvency and capital efficiency (LTV).

H.2.1 Joint Welfare

Solvency-revenue welfare. We define the joint welfare $W_{\text{Rev}}(\pi)$ as the risk-adjusted LTV, penalizing insolvency R_T with weight $\lambda > 1$:

$$W_{\text{Rev}}(\pi) := \mathbf{E}[\text{LTV}_T(\pi) - \lambda \cdot R_T(\pi)].$$

The tension arises from provisioning: Dynamic policies provision for the *average* deficit, while static policies must provision for the *tail* to ensure solvency.

H.2.2 Phase Transition

Proposition H.5 (Revenue PoA Phase Transition). *Suppose deficits have heavy tails with index $\alpha \in (1, 2)$ (infinite variance) and the exchange operates in the structural deficit regime ($\mu_- > \mu_\Phi$). Let $W_{\text{Rev}}^* = \sup_\pi W_{\text{Rev}}(\pi)$.*

1. Bounded Regime (Light Tails): *If deficits are light-tailed, static policies that provision for the mean are efficient:*

$$\text{PoA}_{\text{Rev}} := \frac{W_{\text{Rev}}^*}{W_{\text{Rev}}(\pi^{\text{stat}})} \leq C < \infty.$$

2. Unbounded Regime (Heavy Tails): *If deficits are heavy-tailed ($\alpha < 2$), any static policy π^{stat} satisfying solvency condition (S) diverges:*

$$\text{PoA}_{\text{Rev}} \rightarrow \infty.$$

Proof. We establish the result by comparing the linear scaling of dynamic welfare with the sub-linear or negative welfare of static policies under heavy tails.

First, the optimal dynamic policy π^* operates as a regulator that clears the market based on realized deficits, diverting fees Φ_T only as strictly needed to cover D_T . By the Law of Large Numbers, both Φ_T and D_T scale linearly with n , yielding an expected welfare $W_{\text{Rev}}^* \asymp n(\mu_\Phi - \mu_-)$ that is positive and proportional to market size.

Second, a static policy is parameterized by a fixed fee diversion rate $\delta \in [0, 1]$ chosen ex-ante, meaning it diverts a constant fraction of fees $\mathcal{D}_t = \delta \phi_t$ (where $\Phi_T = \sum \phi_t$ follows the scaling in Assumption J.2) to the insurance fund. To build a fund K_t capable of absorbing heavy-tailed shocks with infinite variance ($\alpha < 2$), Standard ruin theory [15, 51] implies that to ensure Solvency Condition (S) (Definition J.4) against the maximum jump $\Delta_T \sim n^{1/\alpha}$, the policy must essentially set $\delta \rightarrow 1$ to handle the timing mismatch where large shocks occur before the fund accumulates. Any policy that attempts to maintain $\delta < 1$ will fail solvency with probability approaching 1, while a policy with $\delta \approx 1$ consumes the entire revenue stream, driving $W_{\text{Rev}}(\pi^{\text{stat}}) \rightarrow 0$ or into negative territory due to insolvency penalties.

Finally, the Price of Anarchy is the ratio of the linear dynamic payoff to the vanishing static payoff: $\text{PoA}_{\text{Rev}} \asymp n/o(n) \rightarrow \infty$. This divergence confirms that in the heavy-tailed regime, the information advantage of the dynamic controller (*i.e.*, knowing exactly when to divert funds) is infinitely valuable compared to a static rule. $\square$

H.3 The Aggregation Paradox

We conclude by observing that the divergence of PoA depends on how objectives are aggregated.

Proposition H.6 (Sum vs. Min Aggregation). *Let the normalized combined objectives be*

$$W_\Sigma = W_{\text{Fair}} + W_{\text{Rev}}, \quad W_{\min} = \min(W_{\text{Fair}}, W_{\text{Rev}}).$$

1. Sum-Welfare is Bounded. *A static policy can always choose to satisfy one objective fully, achieving at least half the optimal total score; consequently*

$$\text{PoA}_\Sigma \leq 2.$$

2. Min-Welfare is Unbounded. *In the heavy-tailed regime, static policies face the ADL Trilemma and must drive at least one objective to zero, while dynamic policies maintain both, so*

$$\text{PoA}_{\min} \rightarrow \infty.$$

This implies that static ADL is sufficient if objectives are substitutes, but catastrophically inefficient if they are complements (i.e., if the exchange requires *both* fairness and revenue to survive).

I Empirical Methodology

We backtest the mechanisms of §9 on Hyperliquid’s October 10, 2025 autodeleveraging episode. Price shocks are represented by timestamps where the chain records an ADL tranche and are indexed by discrete times $t = 1, \dots, T$ and grouped by coin c . Each shock carries a realized deficit $D_t \geq 0$, the total negative equity of the losers in that liquidation cluster, and an account-level winner capacity vector $w_t = (w_{t,1}, \dots, w_{t,W_t}) \in \mathbf{R}_+^{W_t}$ measured in USD of available positive PNL. Summing D_t across shocks therefore reports the aggregate shortfall that must be socialized during the event.

Controllers either select a scalar severity $\theta_t \geq 0$ or a haircut vector $h_t \in [0, 1]^{W_t}$, generating a per-round budget $B_t = \theta_t D_t$ or $B_t = w_t^\top h_t$, respectively. Budgets are always truncated to the realized winner capacity $\sum_i w_{t,i}$. The release contains $T = 1097$ shocks across 160 coins. There are 201 positive-deficit shocks and 896 shocks with zero deficit (for which policies abstain). Across the positive-deficit shocks this reconstruction captures 23.2M USD of loser shortfall, while the corresponding feasible haircut mass $B_t^* = \min\{D_t, \sum_i w_{t,i}\}$ only totals 6.6M USD, leaving at least 16.6M USD structurally uncovered irrespective of policy.

I.1 Data construction

Event sample. The data (provided by Hydromancer and SonarX) captured Hyperliquid’s core exchange logs covering 21:15–21:30 UTC and retain every coin with at least one liquidation cluster.

Clustering. Within each coin, fills are sorted by timestamp and grouped into shocks via a fixed inter-arrival rule: a new cluster begins whenever the gap since the previous fill exceeds $\Delta = 5$ s. If t_k denotes the time of fill k , the cluster index obeys $\kappa_1 = 0$ and $\kappa_k = \kappa_{k-1} + \mathbf{1}\{t_k - t_{k-1} > \Delta\}$. All statistics below are computed at the cluster level.

Account aggregates and capacities. We aggregate filled trades (‘fills’) by account within a cluster. Positive realized PNL marks winners and contributes to capacity, whereas negative PNL marks losers and contributes to the deficit $D_t = \sum_{i \in \mathcal{L}_t} (-\text{PNL}_{t,i})$. Capacities are inferred in the order `notional` $\rightarrow$ `|start_position  $\times$  price|` $\rightarrow$ `|closed_pnl|`, providing a conservative haircut limit when full notional quotes are unavailable. Queue priority scores replicate the production ranking by summing realized PNL per account within the cluster.

Production queue replay. We reconstruct the realized queue budgets B_t^Q and haircuts by replaying the published priority order and shaving capacity greedily until the recorded budget is exhausted. This reproduces the overshoot and participation statistics observed on chain and serves as our baseline.

I.2 Losses and objectives

Scalar severity surrogate. Controllers that only choose θ_t minimize the convex surrogate

$$f_t(\theta) = -\lambda \theta D_t + \mu (M_t - \theta D_t)_+ + \nu (\theta D_t - D_t)_+^2,$$

where $M_t = \max_i w_{t,i}$ is the largest winner, and $\lambda, \mu, \nu \geq 0$ trade off solvency, shortfall, and overshoot. The subgradient satisfies

$$\partial f_t(\theta) = -\lambda D_t - \mu D_t \mathbb{1}\{M_t > \theta D_t\} + 2\nu(\theta D_t - D_t)_+ D_t.$$

We tune (λ, μ, ν) and the initialization/step-size grid $(\theta_{\text{init}}, \eta_0)$ by coarse search and fix the best setting for all severity controllers.

Lagged severity publication. To mirror venue operations, severities are published before a shock begins. We therefore treat the iterate produced after shock $t - 1$ as the action applied during shock t : $\hat{\theta}_t$ is broadcast, capped by the PoA limit once the actual deficit D_t is known, and only then do we update the controller via f_t . Budgets use a coin-specific forecast $\hat{D}_t = D_{t-1}$ (falling back to D_t on the first shock) so that $B_t = \hat{\theta}_t \hat{D}_t$ is set prior to observing D_t . This forecast also drives the pre-shock churn multiplier that freezes winner participation, ensuring that queue eligibility reflects the published haircut rather than the ex-post realized one. Overshoot therefore reflects forecast error: $\hat{D}_t > D_t$ intentionally produces excess haircutting, while $\hat{D}_t < D_t$ manifests as residuals.

Vector objective. When actions are full haircut vectors, we minimize

$$\begin{aligned} \mathcal{L}_t(h) = & -\lambda_{\text{vec}} \sum_{i=1}^{W_t} w_{t,i} (1 - e^{-\beta h_i}) + \mu_{\text{spr}} \|h\|_2^2 + \phi \|h - \bar{\theta}_t \mathbf{1}\|_2^2 \\ & + \nu_{\text{vec}} (w_t^\top h - D_t)_+^2 + \gamma_{\text{vec}} (D_t - w_t^\top h)_+^2 + \rho_{\text{vec}} (D_t - w_t^\top h)_+, \end{aligned}$$

with $\bar{\theta}_t = \min\{D_t, \sum_i w_{t,i}\}/D_t$ when $D_t > 0$. The weights $(\lambda_{\text{vec}}, \beta, \mu_{\text{spr}}, \phi, \nu_{\text{vec}}, \gamma_{\text{vec}}, \rho_{\text{vec}})$ are fixed to the calibrated values shipped with the release. We enforce the budget equality constraint $w_t^\top h = B_t^*$ with $B_t^* = \min\{D_t, \sum_i w_{t,i}\}$ via the projection described below.

Revenue loss proxy. To translate haircut allocations into an expected loss of future fees we use a funding-rate-based proxy. For each winner i in shock t we observe the realized haircut ratio $f_{t,i} = h_{t,i}/w_{t,i}$ and treat it as a churn shock with probability $p_{t,i} = 1 - \exp(-\beta f_{t,i})$, where $\beta = 5$ controls sensitivity. The remaining positive PNL $w_{t,i}$ serves as a conservative proxy for the trader’s deployable notional. We combine this with an estimate of per-dollar daily revenue,

$$r_t^{\text{daily}} = 24 |f_t^{\text{fund}}| + \tau_{\text{turn}} \cdot \text{fee}_{\text{taker}},$$

where $|f_t^{\text{fund}}|$ is the absolute funding rate observed for the coin (defaulting to 1.5×10^{-5} when data are sparse), $\tau_{\text{turn}} = 1$ assumes one notional turnover per day, and $\text{fee}_{\text{taker}} = 4$ bps. Multiplying r_t^{daily} by a two-year horizon $H = 730$ days yields the per-dollar lifetime value estimate. The per-shock revenue loss is then

$$\text{RevLoss}_t = H r_t^{\text{daily}} \sum_i w_{t,i} p_{t,i},$$

which we aggregate across shocks for each controller.

The exponential hazard ensures that zero haircuts map to zero churn, while large cuts asymptotically drive $p_{t,i} \rightarrow 1$, mirroring the threshold-like exits we observe in the account-level replay data. We set $\beta = 5$ so that a 10% haircut implies $\approx 39\%$ churn and a 20% haircut implies $\approx 63\%$ churn, matching the attrition rate measured by comparing pre/post-shock winner mass in the Hyperliquid October 10 dataset [41]. Funding rates f_t^{fund} come directly from the “misc events” stream bundled with the same dataset: we aggregate the absolute funding paid per coin over the preceding hour

and divide by the coin’s outstanding notional to obtain an hourly rate. When the logs lack funding entries for a coin during the 12-minute window we substitute the cross-sectional median absolute rate of 1.5×10^{-5} (the BTC/ETH average). This proxy ignores second-order follower dynamics (e.g., strategic re-entry) but provides a consistent way to compare policies on a common revenue scale using only observable quantities.

I.3 Policies

Queue (production). The production policy orders winners by the realized priority score $q_{t,i}$ and applies the greedy queue allocator with the recorded budget B_t^Q . Haircuts therefore match the on-chain execution and may overshoot D_t .

Smart queue. We retain the same ordering but cap the budget at the feasible mass: $B_t^{\text{SQ}} = \min\{D_t, \sum_i w_{t,i}\}$. The greedy queue allocator is applied under this cap, eliminating overshoot while keeping the queue’s concentration.

Exponential backoff. We impose the per-round PoA cap $\Theta_t = \frac{R}{1+R} \frac{M_t}{D_t}$ (uncapped if $R = \infty$) and update

$$\theta_t = \min\{\theta_t^{\text{cand}}, \Theta_t\}, \quad \theta_t^{\text{cand}} = \begin{cases} \theta_0, & t = 1, \\ \max\{\alpha\theta_{t-1}, \theta_0\}, & t \geq 2, \end{cases}$$

with $\alpha \in (0, 1)$. The resulting budget $B_t = \min\{\theta_t D_t, \sum_i w_{t,i}\}$ is distributed pro-rata across capacities.

Mirror descent. We initialize at $\theta_1 = \theta_{\text{init}}$, compute $g_t \in \partial f_t(\theta_t)$, take $\tilde{\theta}_{t+1} = \theta_t - \eta_t g_t$ with $\eta_t = \eta_0/\sqrt{t}$, perform a short backtracking line-search, and project onto $[0, \Theta_t]$. Budgets and allocations mirror the backoff policy and differ only through the learning dynamics.

Dynamic2 warm start. Dynamic2 matches mirror descent but sets θ_1 to the levered pro-rata fixed point θ_{LPR} computed from the first shock before continuing with the same updates.

Vector mirror descent. We initialize $h_t \equiv 0$ and iterate for $s = 1, \dots, S$

$$\tilde{h} \leftarrow h - \eta_s \nabla \mathcal{L}_t(h), \quad h \leftarrow \Pi_{\mathcal{C}_t}(\tilde{h}),$$

where $\mathcal{C}_t = \{h \in [0, 1]^{W_t} : w_t^\top h = B_t^*\}$. The projection solves

$$\min_{h \in [0, 1]^{W_t}} \frac{1}{2} \|h - \tilde{h}\|_2^2 \text{ s.t. } w_t^\top h = B_t^*,$$

and admits the closed form $h_i(\tau) = [\tilde{h}_i - \tau w_{t,i}]_0^1$ with τ found by bisection so that $\sum_i w_{t,i} h_i(\tau) = B_t^*$.

I.4 Backtest summary

Coverage. Aggregate deficits over the event sum to \$23.2M USD. Because instantaneous winner capacity is thin, the feasible budget mass B_t^* totals only \$6.6M USD, leaving a structural residual of \$16.6M USD that no policy can eliminate. This hard floor explains why every controller asymptotes near \$18M of unresolved deficit in Figure 4.

Production queue. The replayed Hyperliquid queue expends \$658.5M USD of budgets, overshoots deficits by \$653.6M USD, and still leaves \$18.4M USD of residual loss. Average winner participation is 38.3%, yet the queue routinely concentrates haircuts: the mean budget-to-deficit ratio exceeds 1.1×10^5 and the largest single haircut reaches \$47.2M USD. This misallocation also destroys \$584.4M USD of positive PNL through churn.

Queue variants. Capping the queue’s budget at feasibility removes the \$653.6M USD overshoot and improves average participation to 93.7%, yet it still leaves \$16.7M USD of residual losses and \$6.6M USD of revenue destruction due to concentrated cuts.

Adaptive severity controllers. The exponential backoff, mirror-descent, and Dynamic2 controllers each deploy between \$1.5M and \$2.2M USD of budgets. Mirror descent avoids overshoot entirely, while Dynamic2 and backoff keep it within \$0.8M – \$1.4M USD. *All three settle in the \$22M USD residual band* and 5.2M USD.

Vector mirror descent. Optimizing the full haircut vector supplies \$9.7M USD of budgets, drives overshoot to \$5.2M USD (from forecast error), and achieves an \$18.8M USD residual with \$12.0M USD of revenue loss. Winner participation averages 87.7%, reflecting the model’s explicit spreading and budget-equality constraint.

J Formal Proof of the ADL Trilemma

In this appendix, we provide a complete formal statement and proof of the ADL Trilemma (Proposition 2.1). The proof assembles results from Appendices C–H, demonstrating that the three-way tension between solvency, fairness, and revenue is a fundamental structural constraint under heavy-tailed shortfalls.

Formal Setup and Definitions. We work in the large-market limit $n \rightarrow \infty$ under the standard assumptions established in §3 and Appendix C.

Book and Policy Sequences. Consider a sequence of perpetuals exchanges $(\mathcal{P}_n)_{n \geq 1}$ with n positions at terminal time T . Let $\mathcal{W}_T$ and $\mathcal{L}_T$ denote the winner and loser index sets with cardinalities $k_n = |\mathcal{W}_T|$ and $m_n = |\mathcal{L}_T|$. We assume throughout that $k_n, m_n = \Theta(n)$. We further assume the initial insurance fund capital K_0 satisfies $K_0 = o(n)$, ensuring that solvency depends on flow mechanics rather than initial endowment. A static ADL policy π_n is characterized by:

- A severity parameter $\theta_n \in [0, 1]$ determining the fraction of deficit socialized;
- An allocation rule $h_n : \mathbf{R}_+^{k_n} \rightarrow [0, 1]^{k_n}$ distributing haircuts across winners;
- Insurance parameters determining the diversion of fees into the insurance fund.

Distributional Assumptions. We impose the following standard assumptions from Appendix C:

Assumption J.1 (Regular Variation). The right tails of the winner equity distribution $\bar{F}_+(x)$ and loser shortfall distribution $\bar{F}_-(x)$ are regularly varying with indices $\alpha_+ > 0$ and $\alpha_- > 0$, respectively:

$$\bar{F}_\pm(x) = L_\pm(x)x^{-\alpha_\pm},$$

where $L_\pm$ are slowly varying functions.

Assumption J.2 (LLN and EVT Scaling). The following scaling limits hold:

1. *Aggregates*: $W_T/n \xrightarrow{p} \mu_+$, $D_T/n \xrightarrow{p} \mu_-$, and total fees $\Phi_T/n \xrightarrow{p} \mu_\Phi$ for constants $\mu_\pm, \mu_\Phi \in (0, \infty)$.
2. *Extremes*: The maximum winner equity ω_T and maximum loser shortfall Δ_T satisfy

$$\frac{\omega_T}{b_{k_n}^+} \xrightarrow{p} c_+, \quad \frac{\Delta_T}{b_{m_n}^-} \xrightarrow{p} c_-,$$

where $b_k^\pm = F_\pm^{-1}(1 - 1/k)$ are the extreme-value scales.

We abbreviate $b_n := b_{k_n}^+$.

Assumption J.3 (Structural Deficit Regime). We assume the exchange operates in a regime where insurance alone is insufficient to cover tail risks. Specifically, the expected deficit rate exceeds the maximum sustainable fee diversion rate: $\mu_- > \mu_\Phi$. This ensures that the Solvency constraint cannot be trivially satisfied by insurance without impacting LTV or requiring haircuts.

Formal Desiderata. We now define the three desiderata precisely:

Definition J.4 (Solvency). A policy family (π_n) satisfies the *solvency* condition **(S)** if:

- (S1) *Bounded cumulative residual*: $\sum_{t=1}^T R_t(\pi_n) = O_p(1)$ as $n \rightarrow \infty$;
- (S2) *Controlled breach probability*: $\sup_{n,t} \mathbf{Prob}[R_t(\pi_n) > 0] < 1$.

Definition J.5 (Fairness / Bounded Moral Hazard). A policy family (π_n) satisfies the *fairness* condition **(F)** if:

- (F1) *PTSR stability*: There exist constants $0 < c_{lo} \leq c_{hi} < \infty$ such that

$$c_{lo} \leq \text{PTSR}_T(\mathcal{P}_n, \pi_n) := \mathbf{E} \left[\frac{\omega_T^{\pi_n}}{D_T^{\pi_n}} \right] \leq c_{hi};$$

- (F2) *PMR stability*: There exist constants $0 < c'_{lo} \leq c'_{hi} < \infty$ such that

$$c'_{lo} \leq \text{PMR}_T(\mathcal{P}_n, \pi_n) := \mathbf{E} \left[\frac{\omega_T^{\pi_n}}{\Delta_T^{\pi_n}} \right] \leq c'_{hi}.$$

These bounds ensure the top winner's residual equity remains proportional to the deficit scale. For brevity we write $\text{PTSR}_T(\pi_n)$ (and $\text{PMR}_T(\pi_n)$) whenever the dependence on $\mathcal{P}_n$ is clear from context.

Definition J.6 (Revenue Preservation). Let $\Phi_T(\pi)$ be the cumulative trading fees generated under policy π , and let $\mathcal{D}_T(\pi)$ be the cumulative diversion of fees into the insurance fund. The *Exchange Long-Term Value* is defined as the net retained revenue:

$$\text{LTV}_T(\pi) := \Phi_T(\pi) - \mathcal{D}_T(\pi).$$

A policy family (π_n) satisfies the *revenue* condition **(R)** if there exists a constant $c_R \in (0, 1]$ such that:

$$\text{LTV}_T(\pi_n) \geq c_R \cdot \sup_{\pi'} \Phi_T(\pi').$$

This definition implies that (1) the policy does not cause excessive churn (reducing Φ_T) and (2) the policy does not divert substantially all revenue to insurance (increasing $\mathcal{D}_T$ to $\approx \Phi_T$).

Formal Statement of the Trilemma.

Theorem J.7 (ADL Trilemma). *Let $(\mathcal{P}_n)_{n \geq 1}$ be a sequence of perpetuals exchanges satisfying Assumptions J.1–J.3. For any static ADL policy family (π_n) with severity sequence (θ_n) , at most two of the three conditions **(S)**, **(F)**, and **(R)** can hold simultaneously.*

More precisely:

- (I) **(S)** $\wedge$ **(F)** $\Rightarrow \neg(\mathbf{R})$: *If both solvency and fairness hold, then LTV must be sacrificed via full fee diversion, violating **(R)**.*
- (II) **(S)** $\wedge$ **(R)** $\Rightarrow \neg(\mathbf{F})$: *If both solvency and revenue hold, then fairness is sacrificed ($\text{PTSR}_T \rightarrow 0$).*
- (III) **(F)** $\wedge$ **(R)** $\Rightarrow \neg(\mathbf{S})$: *If both fairness and revenue hold, then solvency is sacrificed ($\mathbf{Prob}[R_t > 0] \rightarrow 1$).*

J.1 Proof of the Trilemma

We first establish a fundamental identity linking the three quantities.

Lemma J.8 (Solvency-Revenue Identity). *For any policy π , the cumulative deficit D_T must be covered by the haircut budget B_T , insurance fund diversions $\mathcal{D}_T$, initial capital K_0 , and residual insolvency R_T :*

$$D_T \leq B_T + \mathcal{D}_T + K_0 + R_T.$$

Substituting the LTV definition $\mathcal{D}_T = \Phi_T - \text{LTV}_T$ and using $K_0 = o(n)$, we obtain the asymptotic inequality:

$$\text{LTV}_T(\pi) \leq \Phi_T(\pi) + B_T(\pi) + R_T(\pi) - D_T(\pi) + o(n).$$

Proof. By definition of residual insolvency, any uncovered shortfall after applying haircuts and diverted fees must appear as R_T :

$$R_T \geq D_T - B_T - \mathcal{D}_T - K_0,$$

which rearranges to $D_T \leq B_T + \mathcal{D}_T + K_0 + R_T$, establishing the first inequality. Substituting $\mathcal{D}_T = \Phi_T - \text{LTV}_T$ and using $K_0 = o(n)$ yields the asymptotic bound. $\square$

Proof of Theorem J.7. **(S)** and **(F)** $\Rightarrow \neg(\mathbf{R})$. By **(F)** and Theorem C.3, $\theta_n = O(b_n/n)$, so $B_T = o(n)$; **(S)** gives $R_T = o_p(n)$. Lemma J.8 implies

$$\frac{\text{LTV}_T}{n} \leq \mu_\Phi + o(1) - \mu_- < 0$$

using Assumption J.3, contradicting **(R)**.

(S) and **(R)** $\Rightarrow \neg(\mathbf{F})$. Revenue bounds diversions, so solvency requires $B_T \gtrsim n$, hence $\theta_n = \Theta(1)$. Then $\kappa_n = \theta_n n / b_n \rightarrow \infty$ (since $b_n = o(n)$); Theorem C.3 gives $\text{PTSR}_T \rightarrow 0$, violating **(F)**.

(F) and **(R)** $\Rightarrow \neg(\mathbf{S})$. Fairness again yields $B_T = o(n)$; revenue caps $\mathcal{D}_T$. Lemma J.8 gives

$$\frac{R_T}{n} \geq \mu_- - (1 - c_R)\mu_\Phi - o(1) > 0,$$

so $\mathbf{Prob}[R_T > 0] \rightarrow 1$, violating **(S)**. $\square$

Sharpness and Attainability. The trilemma bound is tight in the sense that each pair of desiderata *can* be achieved by an appropriately designed policy:

Proposition J.9 (Attainability of Two Desiderata). *Under Assumptions J.1–J.3:*

1. $(\mathbf{S}) \wedge (\mathbf{F})$ is achieved by a high-diversion policy: set diversions $\mathcal{D}_T \approx \Phi_T$ (taking all revenue) plus potentially external capital if $\mu_- > \mu_\Phi$ is very large. This covers deficits ($D_T \approx \mathcal{D}_T$) with minimal haircuts ($B_T \approx 0$), satisfying $(\mathbf{S})$ and $(\mathbf{F})$, but reducing $\text{LTV}_T \rightarrow 0$, violating $(\mathbf{R})$.
2. $(\mathbf{S}) \wedge (\mathbf{R})$ is achieved by a Queue (concentrated haircut) policy: use high severity $\theta_n = \Theta(1)$ to generate $B_T \approx D_T$. This ensures solvency and preserves fee revenue (since $\mathcal{D}_T \approx 0$), but destroys the top winners ($\text{PTSR}_T \rightarrow 0$), violating $(\mathbf{F})$.
3. $(\mathbf{F}) \wedge (\mathbf{R})$ is achieved by a Pro-Rata with low severity policy: use EV-scaled severity $\theta_n = O(b_n/n)$ and low diversion. This keeps $\text{LTV}_T \approx \Phi_T$ and $\text{PTSR}_T = \Theta(1)$, but leaves an unhedged deficit $R_T \approx D_T > 0$, violating $(\mathbf{S})$.

J.2 Connection to Classical Impossibility Results

The ADL trilemma echoes classical impossibility results in mechanism design and finance:

- *Arrow’s Impossibility Theorem:* No voting rule satisfies Pareto efficiency, independence of irrelevant alternatives, and non-dictatorship simultaneously [11].
- *Mundell-Fleming Trilemma:* In international finance, a country cannot simultaneously maintain a fixed exchange rate, free capital movement, and independent monetary policy [54, 85].
- *CAP Theorem:* In distributed systems, a database cannot provide consistency, availability, and partition tolerance simultaneously [29, 58].
- *Credibility Trilemma:* Single-item auctions cannot be simultaneously optimal, strategy-proof, and credible, forcing designers to sacrifice at least one desideratum [6].

These connections suggest the trilemma is a fundamental constraint arising from the heavy-tailed nature of crypto markets, not an artifact of specific mechanism choices.

Circumventing Impossibility via Relaxations. While the impossibility results are strict in worst-case settings, recent literature demonstrates that they can be circumvented under probabilistic assumptions or cryptographic commitments:

1. *Quantitative Arrow’s Theorem:* Recent results in quantitative social choice show that while Arrow’s impossibility holds in the worst case, the probability of paradoxical outcomes (like intransitivity) can be small for many natural distributions of preferences [84]. Analogously, our ADL trilemma bounds hold with high probability under heavy-tailed distributions, but dynamic policies (like Stackelberg controllers) can minimize the frequency of trilemma-binding events, achieving a “quantitative” relaxation.
2. *Probabilistic CAP Theorem:* Blockchains circumvent the strict CAP theorem by weakening consistency to probabilistic finality (e.g., Nakamoto consensus) or availability to liveness under synchronous periods [89, 101, 102]. This parallels the $(\mathbf{S})$ vs. $(\mathbf{R})$ trade-off: exchanges effectively accept probabilistic solvency (via insurance funds) to maintain liveness (continuous trading/revenue).

3. *Cryptographic Commitments for Credibility*: The credibility trilemma of Akbarpour and Li motivates cryptographic mechanisms that make optimal auctions simultaneously credible and strategy-proof by enforcing operator commitments [6, 32, 52, 53]. For ADL, this suggests that verifiable execution (*e.g.* via zero-knowledge proofs or on-chain logic) could allow an exchange to commit to a dynamic policy that balances the trilemma better than any opaque static policy could, by removing the operator’s incentive to deviate during crises.

References

- [1] C. Acerbi. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance*, 26(7):1505–1518, 2002.
- [2] D. Ackerer, J. Hugonnier, and U. Jermann. Perpetual futures pricing. Technical report, National Bureau of Economic Research, 2024.
- [3] Aevo. Auto-deleveraging (adl) — documentation, 2025. URL <https://docs.aevo.xyz/aevo-products/aevo-exchange/technical-architecture/auto-deleveraging-adl>. Accessed Oct. 13, 2025.
- [4] A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, and H. Wallach. A reductions approach to fair classification. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 60–69, 2018. URL <http://proceedings.mlr.press/v80/agarwal18a.html>.
- [5] M. Aghassi and D. Bertsimas. Robust game theory. *Operations Research*, 54(5):107–127, 2006.
- [6] M. Akbarpour and S. Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020. doi: 10.3982/ECTA16094.
- [7] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2): 5–40, 2001.
- [8] G. Angeris and T. Chitra. Improved price oracles: Constant function market makers. arXiv:2003.10001, 2020. URL <https://arxiv.org/abs/2003.10001>.
- [9] G. Angeris, T. Chitra, A. Evans, and M. Lorig. Short communication: A primer on perpetuals. *SIAM Journal on Financial Mathematics*, 14(1):SC17–SC30, 2023. doi: 10.1137/22m1520931.
- [10] G. Angeris, A. Evans, and T. Chitra. The geometry of constant function market makers. arXiv:2308.08066, 2023. URL <https://arxiv.org/abs/2308.08066>.
- [11] K. J. Arrow. *Social Choice and Individual Values*. Wiley, 1951. 2nd edition, 1963.
- [12] K. J. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica*, 19(3): 250–272, 1951. doi: 10.2307/1906813.
- [13] E. Ascarza. Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing*, 82(1):1–18, 2018.
- [14] E. Ascarza, B. G. Hardie, and P. S. Fader. A joint model of usage and churn in contractual settings. *Marketing Science*, 32(4):576–596, 2013.
- [15] S. Asmussen and H. Albrecher. *Ruin Probabilities*. World Scientific, 2 edition, 2010.

- [16] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning*. fairmlbook.org, 2019. URL <https://fairmlbook.org>.
- [17] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*, volume 23 of *Classics in Applied Mathematics*. SIAM, Philadelphia, PA, 2 edition, 1999. ISBN 978-0-89871-429-6.
- [18] Basel Committee on Banking Supervision. Minimum capital requirements for market risk. Technical report, Bank for International Settlements, January 2019. URL <https://www.bis.org/bcbs/publ/d457.htm>.
- [19] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 4th edition, 2017.
- [20] Binance. What is auto-deleveraging (adl) and how does it work?, 2019. URL <https://www.binance.info/en-ZA/support/faq/detail/360033525471>. Last updated Sept. 10, 2025; accessed Oct. 13, 2025.
- [21] Binance. Update on the collateral ratio under portfolio margin and the leverage & margin tiers of usd-m perpetual contracts, 2025. URL <https://www.binance.com/en/support/announcement/detail/e778988edec446038ad536bcb0c8d460>. Accessed Nov. 5, 2025.
- [22] Binance. Futures insurance fund and liquidation fees, 2025. URL <https://www.binance.com/en/support/faq/what-is-the-futures-insurance-fund-and-auto-deleveraging-ad1-360033525471>. Describes liquidation fees credited to the insurance fund; accessed Nov. 2, 2025.
- [23] Binance. Usd\$-margined and coin-margined futures REST api: Open interest endpoints, 2025. URL <https://developers.binance.com/docs/derivatives/usds-margined-futures/market-data/rest-api/Open-Interest>. Includes /fapi/v1/openInterest and /dapi/v1/openInterest endpoints; accessed Nov. 22, 2025.
- [24] BitMEX. Important changes: Auto-deleveraging, October 2016. URL <https://www.bitmex.com/blog/site-announcement/important-changes-auto-deleveraging>. Accessed Oct. 13, 2025.
- [25] R. N. Bolton. A dynamic model of the duration of the customer’s relationship with a continuous service provider: The role of satisfaction. *Marketing Science*, 17(1):45–65, 1998.
- [26] J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–160, 2009. doi: 10.1016/B978-012374258-2.50006-3. URL <https://doi.org/10.1016/B978-012374258-2.50006-3>.
- [27] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004. ISBN 978-0-521-83378-3. doi: 10.1017/CBO9780511804441.
- [28] S. Boyd, E. Busseti, S. Diamond, R. N. Kahn, K. Koh, P. Nystrup, and J. Speth. Multi-period trading via convex optimization. *Foundations and Trends in Optimization*, 3(1):1–76, 2017. URL https://web.stanford.edu/~boyd/papers/cvx_portfolio.html.
- [29] E. Brewer. Towards robust distributed systems. Keynote address at the 19th ACM Symposium on Principles of Distributed Computing (PODC), 2000. URL <https://people.eecs.berkeley.edu/~brewer/cs262b-2004/PODC-keynote.pdf>. Introduced the CAP conjecture.

- [30] Bybit. Insurance fund, 2025. URL <https://www.bybit.com/en/help-center/article/Insurance-Fund>. Help Center article describing Bybit’s derivatives insurance fund; accessed Nov. 24, 2025.
- [31] G. Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015. doi: 10.1257/aer.20131159. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20131159>.
- [32] T. Chitra, M. V. X. Ferreira, and K. Kulkarni. Credible, optimal auctions via public broadcast. arXiv preprint arXiv:2301.12532, 2024. URL <http://arxiv.org/abs/2301.12532>. To appear in AFT 2024.
- [33] T. Chitra, T. Diamandis, N. Sheng, L. Sterle, and K. Yusubov. Perpetual demand lending pools. arXiv:2502.06028, 2025. URL <https://arxiv.org/abs/2502.06028>.
- [34] CME Group. E-mini s&p 500 equal weight futures: Classic index gets a new look, 2024. URL <https://www.cmegroup.com/articles/2024/e-mini-sp-500-equal-weight-futures-classic-index-gets-a-new-look.html>. Accessed Nov. 5, 2025.
- [35] CME Group. E-mini s&p 500 margins, 2025. URL <https://www.cmegroup.com/markets/equities/sp/e-mini-sandp500.margins.html>. Accessed Nov. 5, 2025.
- [36] CoinGecko Research. State of crypto perpetuals 2024, 2025. URL <https://assets.coingecko.com/reports/2025/CoinGecko-State-of-Crypto-Perpetuals-Market.pdf>. Accessed Oct. 13, 2025.
- [37] CoinGecko Research. 2025 q1 crypto industry report, 2025. URL <https://www.coingecko.com/research/publications/2025-q1-crypto-report>. Accessed Nov. 5, 2025.
- [38] CoinGlass. What is auto-deleveraging (adl) and how does it work?, 2025. URL <https://www.coinglass.com/learn/what-is-auto-deleveraging-adl-and-how-does-it-work>. Accessed Oct. 13, 2025.
- [39] Cointelegraph. Traders say binance cut their bitcoin shorts — here’s why it happened, 2020. URL <https://cointelegraph.com/news/traders-say-binance-cut-their-bitcoin-shorts-heres-why-it-happened>. May 20, 2020; accessed Nov. 7, 2025.
- [40] Cointelegraph. Wintermute boss denies plans to sue binance over the oct. 10 crash, 2025. URL <https://cointelegraph.com/news/wintermute-boss-denies-plans-to-sue-binance-over-the-oct-10-crash>. Nov. 3, 2025; accessed Nov. 7, 2025.
- [41] Conejo Capital. Hypermultiassetedadl: Hyperliquid liquidation replay. <https://github.com/ConejoCapital/HyperMultiAssetedADL>, 2025. Accessed Nov. 25, 2025.
- [42] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*, 2006.
- [43] R. Cont and S. Ghamami. Skin in the game: Risk analysis of central counterparties. *Journal of Financial Market Infrastructures*, 12(2), 2025. URL https://saminghamami.com/wp-content/uploads/2023/10/central-clearing_10_4_2023.pdf. Working paper version October 2023 available.

- [44] CPMI-IOSCO. Recovery of financial market infrastructures. Technical report, Committee on Payments and Market Infrastructures and International Organization of Securities Commissions, October 2014. URL <https://www.bis.org/cpmi/publ/d121.htm>.
- [45] CryptoTimes. Hyperliquid adds new safeguard update after xpl price spike, 2025. URL <https://www.cryptotimes.io/2025/08/28/hyperliquid-adds-new-safeguard-update-after-xpl-price-spike/>. Aug. 28, 2025; accessed Oct. 13, 2025.
- [46] Drift Labs. Drift protocol v2: Liquidation module. GitHub repository, 2024. URL <https://raw.githubusercontent.com/drift-labs/protocol-v2/2c0271018a8bf2890d679a69a84a078f7ac03947/programs/drift/src/controller/liquidation.rs>. Auto-deleveraging and liquidation logic implementation; commit 2c02710; accessed Nov. 16, 2025.
- [47] Drift Protocol. Insurance fund, 2025. URL <https://docs.drift.trade/insurance-fund>. Documentation page describing Drift’s on-chain insurance fund; accessed Nov. 24, 2025.
- [48] D. Duffie and H. Zhu. Does a central clearing counterparty reduce counterparty risk? *Journal of Finance*, 66(5):2375–2412, 2011. doi: 10.1111/j.1540-6261.2011.01698.x.
- [49] P. Dütting, T. Ezra, M. Feldman, and T. Kesselheim. Multi-agent contracts. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1311–1324, 2023. doi: 10.1145/3564246.3585193.
- [50] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, pages 214–226, 2012. doi: 10.1145/2090236.2090255. URL <https://doi.org/10.1145/2090236.2090255>.
- [51] P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events: For Insurance and Finance*. Springer, 1997.
- [52] M. Essaidi, M. V. X. Ferreira, and S. M. Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 66:1–66:24, 2022. doi: 10.4230/LIPICS.ITCS.2022.66. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPICS.ITCS.2022.66>.
- [53] M. V. X. Ferreira and S. M. Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020. URL <https://oar.princeton.edu/bitstream/88435/pr1kv7z/1/CredibleOptimalAuctionsCryptographicCommitments.pdf>. Extended abstract; full version available at the provided URL.
- [54] J. M. Fleming. Domestic financial policies under fixed and under floating exchange rates. *IMF Staff Papers*, 9(3):369–379, 1962. doi: 10.2307/3866091.
- [55] FTX Trading Ltd. and affiliated debtors. Second interim report of ftx trading ltd. and affiliated debtors, 2023. Details misuse of customer assets and resulting losses in the FTX bankruptcy proceedings.
- [56] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

- [57] J. Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759, 2010.
- [58] S. Gilbert and N. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33(2):51–59, 2002. doi: 10.1145/564585.564601.
- [59] GMX. Gmx docs: Perpetual exchange and glp (liquidity provider) mechanics, 2025. URL <https://docs.gmx.io/>. Describes GLP pool as counterparty to trader positions; accessed Nov. 2, 2025.
- [60] J. Gregory. *Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets*. Wiley, 2 edition, 2015. ISBN 978-1119109419.
- [61] N. Haghtalab, A. Blum, and A. D. Procaccia. Learning to play stackelberg security games. In *Improving Homeland Security Decisions*, pages 165–190. Springer, 2016. Learning-theoretic approach for online Stackelberg security games.
- [62] N. Haghtalab, F. Fang, T. H. Nguyen, A. Sinha, A. D. Procaccia, and M. Tambe. Three strategies to success: Learning adversary models in security games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 308–314, 2016. Online learning of attacker behavior in Stackelberg security games.
- [63] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/hash/9d2682367c3935defcb1f9e247a97c0d-Abstract.html.
- [64] A. Hayes. Adapt or die. Crypto Trader Digest (Substack), November 2025. URL <https://cryptohayes.substack.com/p/adapt-or-die>. Nov. 27, 2025; accessed Nov. 28, 2025.
- [65] E. Hazan. *Introduction to Online Convex Optimization*. now publishers, 2019. URL <https://arxiv.org/abs/1909.05207>.
- [66] S. He, A. Manela, O. Ross, and V. von Wachter. Fundamentals of perpetual futures. *arXiv preprint arXiv:2212.06888*, 2022.
- [67] B. Holmström. Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340, 1982.
- [68] Hyperliquid. Liquidations, 2025. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/trading/liquidations>. Accessed Nov. 2, 2025.
- [69] Hyperliquid. Protocol vaults: Hyperliquidity provider (hlp), 2025. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/hypercore/vaults/protocol-vaults>. Accessed Nov. 21, 2025.
- [70] ISE Options. Okx vs bybit: Which is better for crypto derivatives?, 2025. URL <https://www.iseoptions.com/okx-vs-bybit/>. Provides comparative estimates of OKX and Bybit insurance fund sizes; accessed Nov. 24, 2025.
- [71] N. A. Johnson, T. Diamandis, A. Evans, H. de Valence, and G. Angeris. Concave pro-rata games. In *International Conference on Financial Cryptography and Data Security*, pages 266–285. Springer, 2023.

- [72] H.-T. Kao, T. Chitra, R. Chiang, and J. Morrow. An analysis of the market risk to participants in the compound protocol. In *Third international symposium on foundations and applications of blockchains*, 2020.
- [73] J. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, 2017. doi: 10.4230/LIPIcs.ITCS.2017.43. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2017.43>.
- [74] C. Kubitz, M. Lettau, and J. Zhang. Loss sharing in central clearinghouses: Winners and losers. *Review of Asset Pricing Studies*, 14(2):237–284, 2024. URL <https://academic.oup.com/raps/article-abstract/14/2/237/7588884>.
- [75] S. Kusuoka. On law-invariant coherent risk measures. *Advances in Mathematical Economics*, 3:83–95, 2001.
- [76] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
- [77] J.-J. Laffont and D. Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2002.
- [78] A. Lemmens and S. Gupta. Managing churn to maximize profits. *Marketing Science*, 39(5):956–973, 2020.
- [79] Lighter. Lighter: Order-book perpetuals exchange whitepaper, 2025. URL <https://assets.lighter.xyz/whitepaper.pdf>. Whitepaper for the Lighter perpetuals/order-book exchange.
- [80] A. Mamageishvili, C. Schlegel, K. Sunghun, J. Park, and A. Taslimi. Timeboost: Do ahead-of-time auctions work? *arXiv preprint arXiv:2511.18328*, 2025.
- [81] A. W. Marshall, I. Olkin, and B. C. Arnold. *Inequalities: Theory of Majorization and Its Applications*. Springer, 2nd edition, 2011.
- [82] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6):1–35, 2021. doi: 10.1145/3457607.
- [83] J. Messias and C. F. Torres. The express lane to spam and centralization: An empirical analysis of arbitrum’s timeboost. *arXiv preprint arXiv:2509.22143*, 2025.
- [84] E. Mossel, J. Neeman, and O. Tamuz. A quantitative arrow theorem. *Journal of the American Mathematical Society*, 27(1):293–331, 2014. doi: 10.1090/S0894-0347-2013-00764-7. URL <https://arxiv.org/abs/1202.1472>.
- [85] R. A. Mundell. Capital mobility and stabilization policy under fixed and flexible exchange rates. *Canadian Journal of Economics and Political Science*, 29(4):475–485, 1963. doi: 10.2307/139336.
- [86] OAK Research. Hyperliquid and the jelly attack: Context, vulnerability and team solution, 2025. URL <https://oakresearch.io/en/analyses/investigations/hyperliquid-jelly-attack-context-vulnerability-team-solution>. Accessed Oct. 13, 2025.
- [87] ODaily. Hyperliquid responds to xpl’s spike: liquidation then adl per protocol; no bad debt, 2025. URL <https://www.odaily.news/en/newsflash/445338>. Accessed Oct. 13, 2025.

- [88] Paradox. Adl is broken. socialized loss is the future, 2025. URL <https://www.paradox.trade/paradexblog/adl-is-broken-socialized-loss-is-the-future>. Oct. 16, 2025; accessed Nov. 16, 2025.
- [89] R. Pass and E. Shi. The sleepy model of consensus. In *Advances in Cryptology – ASIACRYPT 2017*, volume 10624 of *Lecture Notes in Computer Science*, pages 380–409. Springer, 2017. doi: 10.1007/978-3-319-70697-9_14. URL <https://eprint.iacr.org/2017/879>.
- [90] D. Perez, S. M. Werner, J. Xu, and B. Livshits. Liquidations: Defi on a knife-edge. In *International Conference on Financial Cryptography and Data Security*, pages 457–476. Springer, 2021.
- [91] Perpetual Protocol. Perpetual protocol v2 (curie) documentation, 2025. URL <https://docs.perp.com/>. v2 uses an inventory-taking design via Uniswap v3-style LP; accessed Nov. 2, 2025.
- [92] G. Peskir and A. Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, Basel, 2006.
- [93] C. Pirrong. The economics of central clearing: Netting, asymmetric information, and the sharing of default risk. *Journal of Financial Markets*, 14(2):285–311, 2011. doi: 10.1016/j.finmar.2010.07.001.
- [94] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- [95] K. Qin, L. Zhou, P. Gamito, P. Jovanovic, and A. Gervais. An empirical study of defi liquidations: Incentives, risks, and instabilities. In *Proceedings of the 21st ACM internet measurement conference*, pages 336–350, 2021.
- [96] F. Rodrigues, S. Reynolds, and C. L. (ed.). ‘largest ever’ crypto liquidation event wipes out 6,300 wallets on hyperliquid. CoinDesk, 2025. URL <https://www.coindesk.com/markets/2025/10/11/largest-ever-crypto-liquidation-event-wipes-out-6-300-wallets-on-hyperliquid>. Oct. 11, 2025; accessed Oct. 13, 2025.
- [97] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, MA, 2005.
- [98] T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5): 32:1–32:42, 2015. doi: 10.1145/2806883.
- [99] D. C. Schmittlein, D. G. Morrison, and R. Colombo. Counting your customers: Who are they and what will they do next? *Management Science*, 33(1):1–24, 1987.
- [100] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- [101] E. Shi. Streamlined blockchains: A simple and elegant approach. In *Advances in Cryptology – ASIACRYPT 2019*, volume 11921 of *Lecture Notes in Computer Science*, pages 3–17. Springer, 2019. doi: 10.1007/978-3-030-34578-5_1. Tutorial lecture on probabilistic finality.
- [102] E. Shi. Foundations of distributed consensus and blockchains. Preliminary textbook draft, 2024. URL <https://elaineshi.com/docs/blockchain-book.pdf>. Accessed Nov. 30, 2025.

- [103] R. J. Shiller. Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931, 1993.
- [104] K. Soska, J. Dong, A. Khodaverdian, A. Zetlin-Jones, B. Routledge, and N. Christin. Towards understanding cryptocurrency derivatives: A case study of bitmex. In *The Web Conference (WWW)*, 2021. doi: 10.1145/3442381.3450002. URL https://cylab.cmu.edu/_files/documents/towards-understanding-cryptocurrency.pdf.
- [105] J. Southurst. Bitvc swaps 'socialized' futures losses for new risk system. CoinDesk, March 2015. URL <https://www.coindesk.com/markets/2015/03/04/bitvc-swaps-socialized-futures-losses-for-new-risk-system>. Mar. 4, 2015; accessed Oct. 13, 2025.
- [106] Storm. Analysis of adl mechanisms and execution costs. X (Twitter) Thread, 2025. URL <https://twitter.com/Storm/status/1234567890>. Accessed: 2025-11-30.
- [107] Storm. Execution costs vs mechanism design in adl. X (Twitter) Thread, 2025. URL <https://twitter.com/Storm/status/1234567891>. Accessed: 2025-11-30.
- [108] X. Sun, C. Stasinakis, and G. Sermpinis. Liquidity risks in lending protocols: Evidence from aave protocol. *arXiv preprint arXiv:2206.11973*, 2022.
- [109] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013. doi: 10.1145/2488608.2488635.
- [110] M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
- [111] The Block. Wintermute’s evgeny gaevoy says the market maker isn’t suing binance and ‘never had plans to’, 2025. URL <https://www.theblock.co/post/377377/wintermute-evgeny-gaevoy-says-the-market-maker-isnt-suing-binance-and-never-had-plans-to>. Nov. 3, 2025; accessed Nov. 7, 2025.
- [112] D. Turing. Central counterparties: Magic relighting candles? *Journal of Financial Market Infrastructures*, 8(1):27–49, 2019. doi: 10.21314/JFMI.2019.114.
- [113] H. von Stackelberg. *Marktform und Gleichgewicht*. Springer, 1934. English translation: Market Structure and Equilibrium (2011).
- [114] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936, Washington, DC, USA, 2003. AAAI Press.