FirstSwap: Pre-Consensus Trade Approvals for Decentralized Market Design

William Diamond[1], Fatima Elsheimy[2], Dahlia Malkhi[3], Geoff Ramseyer[4], and Srivatsan Sridhar[5⋆]

[1] University of Wisconsin–Madison
[2] Yale University
[3] University of California, Santa Barbara
[4] Stellar
[5] University of California, Berkeley

November 6, 2026

Abstract. Decentralized exchanges typically give traders no execution guarantee until a block is finalized. In the meantime, external prices move, traders strategically delay submission, and block producers can influence inclusion and ordering. We introduce FirstSwap, a DEX built on a constant-function market maker pricing curve that gives traders an early pre-approval before consensus, guaranteeing inclusion and execution within their submitted price limit, with final ordering determined by limit prices rather than by a block producer.

We demonstrate the benefits of trade pre-approvals in an economic model combining auction theory with CEX-DEX arbitrage. When the CEX price changes gradually, FirstSwap rewards earlier traders with better prices, discouraging strategic delay and yielding higher social welfare and gains to AMM liquidity providers than a batch auction. When the CEX price jumps abruptly, FirstSwap yields higher welfare and smaller AMM losses than fair ordering.

The central challenge is delivering fast trade pre-approvals without consensus. We introduce a decentralized, scaled admissibility protocol under which each validator reserves only a fraction of the available liquidity at every price. A trader obtains approval in a single round of messages with a quorum of validators. Somewhat surprisingly, despite Byzantine validators who may approve trades arbitrarily, we show that all quorum-certified trades are jointly admissible and execute within their limits. Our protocol extends naturally to a setting where buy and sell orders can be pre-approved against each other before a block is finalized.

1 Introduction

This paper introduces FirstSwap, a decentralized exchange operated by a network of *validators* that finalizes blocks of trades epoch by epoch. FirstSwap

⋆ We thank Malcolm Granville for insights in the early stages of this project.

incentivizes early submission of trade requests by granting pre-approval certificates promptly upon arrival, without waiting for end-of-epoch consensus. Pre-approvals are processed through a single direct interaction between each trader and a quorum of validators, yielding a fast, consensus-free procedure. We develop an economic model to analyze FirstSwap's pre-approval mechanism as a market design. Early pre-approvals yield higher social welfare and gains to AMM liquidity providers than the end-of-block confirmation typical of blockchains. FirstSwap naturally mitigates two prevalent forms of strategic exploitation. First, MEV is curtailed because trade execution depends only on block contents and not on their ordering, and because pre-approvals must come from a quorum of validators and not at the discretion of a single block builder. Second, CEX–DEX arbitrage exploiting a passive AMM is reduced because traders must race for early pre-approval instead of waiting for arbitrage opportunities.

FirstSwap consists of two key ingredients, each of potential interest on its own: **First**, we propose a pre-approval mechanism which constitutes a new form of auction, where both arrival order and price matter. Traders post trade requests with limit prices; each trade request is either pre-approved or rejected based on a price-contingent capacity without being assigned a final position in the block. At epoch end, a final block is produced with an intra-execution ordering, guaranteed to include all pre-approved trade requests settling within their price limits. We develop an economic model to analyze this mechanism and demonstrate that early pre-approvals yield higher social and AMM profitability compared to other proposed AMM mechanisms. Figure 1 illustrates our findings.

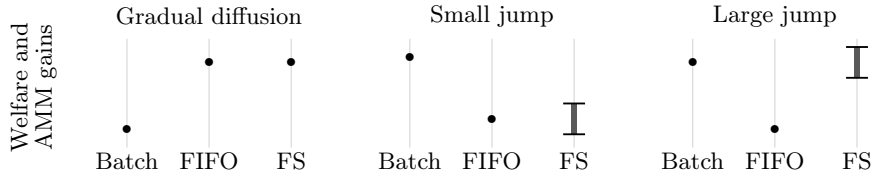

Fig. 1. Qualitative comparison of exchange mechanisms (Batch = batch auction, FIFO = first-in-first-out, FS = FirstSwap). Both overall welfare and AMM gains from trade follow the same trends, under the conditions of Equation (11) and Lemma 4. Higher positions indicate higher outcomes within each scenario; vertical distances are not to scale. The ranges indicate that the comparison of FS with FIFO (small jump) or Batch (large jump) is ambiguous without further assumptions.

Second, we propose a fast, consensus-free implementation for a network of validators to collectively operate the pre-approval system. Each pre-approval is obtained through a single direct exchange between the trader and a quorum of validators, with each validator deciding whether to pre-approve a trade request based solely on local information, requiring no coordination among validators. Somewhat surprisingly, the protocol guarantees that at the end of a block epoch, all pre-approvals obtained in this manner can settle within their price limits. Furthermore, Byzantine validators cannot prevent pre-approved trade requests from settling, and if they attempt to hide pre-approved trade requests, not only

do they fail but their action is punishable. A straightforward extension of our protocol allows for simultaneous buy and sell orders to be handled simultaneously, where validators can pre-approve additional buys when they have seen sell trade reqests that could execute against them and conversely. AMM liquidity providers can also deposit and withdraw funds at the end of the block after orders are executed without impacting our analysis.

1.1 Overview

The Exchange Design. Assume for now that the protocol is run by a single validator. FirstSwap allocates trades against a pool of fungible assets for sale. The exchange employs a continuous pricing mechanism identical to that of constant-function market makers (CFMMs). The price per unit $p(x)$ increases as more assets are purchased, where x is the quantity already bought.

The pre-approval mechanism operates epoch by epoch. During an epoch, trade requests in the form $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$ arrive and are admitted to FirstSwap promptly upon arrival. A trader's limit price limit constrains the total quantity that may be consumed before and including their trade. For each trade request, the pre-approval system determines whether it can be satisfied based on the trades pre-approved so far. If so, unlike existing blockchain-based platforms, the system returns a pre-approval certificate immediately. Pre-approved trades are then gathered into a finalized block for execution. FirstSwap guarantees that all pre-approved trade requests can be executed within their price limits, and the final ordering of trades within the finalized block is computed deterministically.

A key notion in the pre-approval mechanism is *admissibility*: A set T of trades of the form $(\mathsf{amt}, \mathsf{limit})$ is admissible if for every limit ℓ in T, the sum total x of amounts in trades whose limit is $\leq \ell$ satisfies $p(x) \leq \ell$. Note that if a set is admissible, then executing trades in ascending order of limits allows all trade requests to settle within the requested price limits. An incoming trade request tx may be pre-approved if adding the trade request to the set A of previously pre-approved trades preserves admissibility; that is, if $\mathsf{A} \cup \{\mathsf{tx}\}$ is admissible.

For illustration, consider the special case where all trade requests are for a single unit and the price function is simply $p(x) = x$, so the price per unit equals the number of units already consumed. The set $\mathsf{A} = \{3, 1, 3, 5\}$ (writing only limits for each trade) is admissible, because for every $\mathsf{limit} = \ell$, the number of trade requests in A with limit at most ℓ is $\leq \ell$; indeed, an execution in sorted order $1, 3, 3, 5$ satisfies all price limits.

Now consider an incoming trade request $\mathsf{tx} = 4$. The set $\mathsf{A} \cup \{\mathsf{tx}\} = \{3, 1, 3, 5, 4\}$ remains admissible, so the new trade request can be pre-approved, and the new execution order is $1, 3, 3, 4, 5$. A different incoming trade request may be $\mathsf{tx}' = 2$. The set $\mathsf{A} \cup \{\mathsf{tx}'\} = \{3, 1, 3, 5, 2\}$ is not admissible: for $\ell = 3$, there are 4 trade requests whose limits are at most 3, hence $p\big(\{\mathsf{tx} \in \mathsf{A} \cup \{2\} \mid \mathsf{tx} \leq 3\}\big) = 4 > 3$. Indeed, no sequenced execution of $T \cup \{\mathsf{tx}'\}$ may satisfy all price limits. Thus, tx' should be denied pre-approval.

Economic Incentives and Market Efficiency in FirstSwap. A trade in FirstSwap has a greater chance of executing when it is sent earlier and when

it offers a higher limit price, since trades are processed in the order received and more high-limit-price trades can be approved. Traders rationally offer prices strictly below their asset valuations, to ensure a positive gain from an executed trade, but high-valuation traders offer higher prices and also submit trades earlier. FirstSwap uses both trade time and limit price to ensure that assets go to traders who value them the most. The sorting method in FirstSwap lives between the extremes of two other MEV-resistant AMM mechanisms. A first-in-first-out fair ordering protocol as in [11] sets a trade's execution price purely based on arrival time, while a batch auction as in [8, 9] approves trades only based on limit price and ignores arrival time entirely. FirstSwap, unlike the other two mechanisms, incentivizes traders to bid strictly less than their valuations. This lack of honest bidding in FirstSwap actually increases economic efficiency in a model with the realistic possibility of CEX-DEX arbitrage by reducing the volume of arbitrage trades.

We illustrate the economic value of FirstSwap relative to other market designs in an economic model that combines auction theory with dynamic CEX-DEX arbitrage trading. Our model is the first, to our knowledge, that allows us to compare the social welfare of different AMM trading mechanisms where both the limit price and time of a trade are chosen strategically. This makes our model ideal for examining market design responses to the CEX-DEX arbitrage problem. By rewarding traders for trading fast instead of waiting for a better trading opportunity, FirstSwap yields higher utilitarian welfare and gains to AMM liquidity providers than an end-of-epoch batch auction. This is in contrast to [8], which argues in favor of batch auctions in a setting where there is no motive for traders to wait. While a fair ordering protocol based on [22, 11] provides similar improvements over a batch auction, FirstSwap provides higher welfare and gains to AMM liquidity providers than fair ordering after a large, sudden jump in the CEX price.

Fast, Consensus-free Pre-approvals. Granting pre-approvals one by one requires solving the consensus problem repeatedly among the network of validators. FirstSwap avoids solving consensus for each trade request to speed up the pre-approval step. We develop in Section 3 a pre-approval protocol consisting of a single exchange between the trader and a quorum of validators that requires no coordination among validators.

In a nutshell, consensus-free pre-approval works as follows. The trader sends a trade request to the network of validators. Each validator considers only the trade requests it has pre-approved locally and returns a signed response. A trade request that has pre-approvals from a quorum of q distinct validators is pre-approval-certified. In order to guarantee that all pre-approval-certified trade requests form an admissible set without relying on consensus, each validator pre-approves only a certain fraction of the allowed capacity at each price. For n validators, f dishonest validators, and quorum size $q = n-f$, the pre-approval fraction is $\eta := \frac{q-f}{n-f}$. For the conventional BFT quorum sizes, $n = 3f+1$, $q = 2f+1$ and $\eta = \frac{f+1}{2f+1} \approx \frac{1}{2}$; choosing larger quorum systems, e.g., $q = \lceil \frac{4n+1}{5} \rceil$ can yield $\eta \approx \frac{4}{5}$. Importantly, modified admissibility is checked locally by each validator

without any coordination among validators. Nevertheless, we show that all trade requests which are pre-approved individually by validators in this manner can be settled within their price limits.

At the end of each epoch, validators run a single BFT consensus instance, each contributing a *mini-block* of its own pre-approved trades. The consensus output is formed from at least $n-f$ mini-blocks and orders them into a final block, guaranteeing inclusion of every trade with a pre-approval certificate.

2 Fast Pre-Approvals as a Market Design Tool

To analyze the economic properties, we first describe the exchange mechanism, as if run by a single validator, and assume messages are delivered in the order they are sent. Theorem 1 shows that the same economic allocations would also occur in a range of distributed implementations. Section 3 presents one Byzantine fault-tolerant protocol under partial synchrony.

A *trade* is a tuple $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$, where $\mathsf{amt} \in \mathbb{R}_+$ is the amount of assets the trader wants to buy, and $\mathsf{limit} \in \mathbb{R}_+$ is the maximum price per unit. As in a constant-function market maker (CFMM) [21], we consider a monotonically increasing price function p such that if x units of assets have already been bought from the exchange, the price per unit paid for buying the next infinitesimal quantity is $p(x)$.

The exchange operates in periods of time called epochs. In each epoch, traders send their trades to the exchange. The exchange approves the trade if adding this trade to the set of all trades approved so far keeps the set *admissible*.

Definition 1 (Admissibility of trades). *Let $S = \{\mathsf{tx}_i = (\mathsf{amt}_i, \mathsf{limit}_i)\}_{i=1}^m$ be a set of trades. S is admissible if for all i, $p\left(\sum_{j:\mathsf{limit}_j \le \mathsf{limit}_i} \mathsf{amt}_j\right) \le \mathsf{limit}_i$.*

Let $S = \{\mathsf{tx}_i = (\mathsf{amt}_i, \mathsf{limit}_i)\}_{i=1}^m$ be the set of approved trades in the epoch. The validator executes each trade tx_i at the price $p\left(\sum_{j:\mathsf{limit}_j \le \mathsf{limit}_i} \mathsf{amt}_j\right)$. This is equivalent to ordering the approved trades by their limit price, then executing them at the price determined by the pricing curve. Admissibility ensures that all trades can be charged prices consistent with the pricing function while satisfying their limit prices. Ordering by limit prices is without loss of generality due to Lemma 1. Thus, while we could choose any such ordering, ordering by limit prices naturally favors trades with higher valuations.

Lemma 1. *Let $S = \{\mathsf{tx}_i = (\mathsf{amt}_i, \mathsf{limit}_i)\}_{i=1}^m$ be a set of trades. There exists an ordering $(\mathsf{tx}_i = (\mathsf{amt}_i, \mathsf{limit}_i))_{i=1}^m$ on S such that for all i, $p\left(\sum_{j \le i} \mathsf{amt}_j\right) \le \mathsf{limit}_i$ if and only if S is admissible.*

2.1 Economic Model

The model has a continuum of infinitesimal traders and time is measured continuously. At time t, trader i has value $\theta_{i,t} = p_{\mathrm{ex},t} + \varepsilon_i$ for holding the asset provided

by the AMM, where $p_{\text{ex},t}$ is a publicly observed "CEX price" and $\varepsilon_i \sim F_\varepsilon$ is drawn i.i.d. across traders and is fixed over time. The distribution F_ε has unit mass and is normalized to mean 0, so that $p_{\text{ex},t}$ reflects the valuation of the average trader. We write $\varepsilon(r) = F_\varepsilon^{-1}(1-r)$ for the value of ε with a fraction r of the population with a higher valuation for $r \in (0,1)$. We make the simple choice that only the CEX price evolves over time, since only arbitrage trades are time-sensitive enough to meaningfully vary in value within a single epoch.

We analyze the FirstSwap protocol first when the CEX price evolves gradually as a diffusion,

$$\text{Diffusion model:} \qquad p_{\text{ex},t} = p_{\text{ex},0} + \sigma_{\text{ex}} W_t, \qquad T_{\text{blk}} \sim \text{Exp}(\lambda_{\text{blk}}),$$

where W_t is Brownian motion and the exponentially distributed random time T_{blk} at which the epoch ends and transactions execute is independent of W_t. This Brownian-price/Poisson-epoch timing follows the approach of [27, 28] to tractably study loss-versus-rebalancing against an AMM. We adapt their approach in a model with rational traders with differing values who can strategically wait to trade. We also examine the case when the CEX price experiences a single large jump Δp_{ex} and $p_{\text{ex},t} = p_{\text{ex},0} + \Delta p_{\text{ex}}, \forall\, t > 0$.
Economic incentives for traders differ markedly in these settings. With gradual information, traders will strategically wait to observe the evolution of $p_{\text{ex},t}$ before trading, while they will race each other for good price execution [8] after a jump. Each active trader may buy once. Passive LPs commit ex ante to providing AMM liquidity at a price $p(x)$, where x is cumulative inventory sold.

Measures of Economic Value We compare our trading mechanisms using two criteria: social welfare and AMM gains from trade. In the continuum model, social welfare at the end of the epoch, T_{blk}, is

$$\text{Social welfare:} \qquad x_0 p_{\text{ex},T_{\text{blk}}} + \int q_i(\varepsilon_i)\varepsilon_i \, dF_\varepsilon(\varepsilon_i), \tag{1}$$

where x_0 is total asset supply and $q_i(\varepsilon_i)$ is the final quantity allocated to a trader with private value ε_i. For fixed x_0, the common-value component $x_0 p_{\text{ex},T_{\text{blk}}}$ is identical across allocations. Thus welfare differences come from allocating more of the asset to traders with higher private values. We treat active traders and passive AMM liquidity providers as ex ante drawn from the same population; because passive liquidity must be treated identically by the mechanism, its allocation is uncorrelated with ε_i. Hence, increasing the covariance between active-trader holdings and ε_i increases overall social welfare.

Our second metric is AMM gains relative to passively holding the asset:

$$\text{AMM gains:} \qquad (x_0 - x_{\text{LP}})\,(\bar{p} - p_{\text{ex},T_{\text{blk}}})\,, \tag{2}$$

where x_{LP} is the inventory remaining with passive liquidity providers and $\bar{p}$ is the average execution price received by the AMM across trades executed. Because AMM liquidity is passive, strategic traders can leave the AMM with less inventory precisely when $p_{\text{ex},T_{\text{blk}}}$ is high. This results in "loss-versus-rebalancing"

or LVR, where the passive AMM liquidity experiences losses after large price movements. An ideal protocol therefore induces traders with high private valuations ε_i to buy without inducing a large negative correlation between the volume $x_0 - x_{\text{LP}}$ executed and price movements $\bar{p} - p_{\text{ex},T_{\text{blk}}}$.

The trader's problem with gradual price evolution A trader chooses both when to submit a trade and the limit price attached to that trade. Let τ be the stopping time at which the trade occurs and let $p_{\text{bid},\tau}$ denote the limit price submitted at that time. If a mass x_τ of traders has already been approved to trade, the marginal AMM price for our trader is $p(x_\tau)$. Thus a trader with private value ε has expected utility

$$\text{Utility:} \qquad \sup_{\tau, p_{\text{bid},\tau}} \mathbb{E}\left[\mathbf{1}\{\tau < T_{\text{blk}}\}\mathbf{1}\{p_{\text{bid},\tau} \geq p(x_\tau)\}\left(p_{\text{ex},\tau} + \varepsilon - p(x_\tau)\right)\right]. \tag{3}$$

We assume in our gradual diffusion model that messages arrive at validators in exactly the same order as they are sent, so the bid will fill precisely if $p_{\text{bid},\tau} \geq p(x_\tau)$. Since the execution price is $p(x_\tau)$ and does not depend on the trader's bid, whenever the trader wishes to trade it is without loss to submit the smallest limit price that executes, $p_{\text{bid},\tau} = p(x_\tau)$. Thus, the economically substantive decision is when to trade. We look for a monotone equilibrium in which a trader with valuation $\varepsilon(r)$ chooses a time to trade that is strictly increasing in r. Given the AMM pricing function $p(.)$, a price $p(r)$ is available to a trade that fills after a mass r of trades has previously occurred, so $p_{\text{bid},\tau} = p(r)$ in a monotone equilibrium. We show in Appendix A that each trader optimally follows a simple "memoryless" trading strategy: a trader of type r chooses to submit when the CEX price $p_{\text{ex},t}$ first reaches some threshold $p^*_{\text{ex}}(r)$ regardless of other history.

Let $\gamma_{\text{diff}} \equiv \frac{\sqrt{2\lambda_{\text{blk}}}}{\sigma_{\text{ex}}}$. The probability that Brownian motion reaches $p^*_{\text{ex}}(r)$ before the epoch ends is $e^{-\gamma_{\text{diff}}[p^*_{\text{ex}}(r) - p_{\text{ex},0}]}$ (the Brownian first-passage probability formula). Hence the expected utility of a trader whose valuation $\varepsilon(r)$ of the asset is less than that of a fraction r of the population (i.e. $\Pr\{\varepsilon_i > \varepsilon(r)\} = r$) is

$$U(r) = e^{-\gamma_{\text{diff}}[p^*_{\text{ex}}(r) - p_{\text{ex},0}]}\left[p^*_{\text{ex}}(r) + \varepsilon(r) - p(r)\right]. \tag{4}$$

A particularly simple equilibrium obtains when $\kappa \equiv -\frac{p'(r)}{\varepsilon'(r)} > 0$ is constant. In this case, we show in Appendix A that the optimal threshold is

$$p^*_{\text{ex}}(r) = p(r) - \varepsilon(r) + \frac{1}{\gamma_{\text{diff}}(1+\kappa)}. \tag{5}$$

Traders optimally wait until their contemporaneous valuation exceeds the AMM price by a constant $\frac{1}{\gamma_{\text{diff}}(1+\kappa)}$. For our diffusion analysis, we assume that no trader has a value of $\varepsilon(r)$ so high that $p(r) - \varepsilon(r)$ is below the starting CEX price $p_{\text{ex},0}$, which prevents an initial discrete spike in trading volume.

Welfare and AMM gains under FirstSwap and a batch auction. Given the rational bidding strategy in Equation (5), type r trades when some value of $p_{\text{ex},t}$ exceeds $p^*_{\text{ex}}(r)$ before the epoch ends. This has probability (using

the Brownian first-passage probability)

$$\Pr_{\mathrm{FS}}^{\mathrm{diff}}(r) = \exp\left(-\frac{1}{1+\kappa}\right)\exp\{-\gamma_{\mathrm{diff}}[p(r)-\varepsilon(r)-p_{\mathrm{ex},0}]\}. \tag{6}$$

We show in Appendix A that expected welfare $\mathcal{W}_{\mathrm{FS}}^{\mathrm{diff}}$ and AMM gains $\mathcal{R}_{\mathrm{FS}}^{\mathrm{diff}}$ are respectively given by

$$\mathcal{W}_{\mathrm{FS}}^{\mathrm{diff}} = \int_0^1 \varepsilon(r)\Pr_{\mathrm{FS}}^{\mathrm{diff}}(r)\,dr, \tag{7}$$

$$\mathcal{R}_{\mathrm{FS}}^{\mathrm{diff}} = \int_0^1 \left[\varepsilon(r) - \frac{1}{\gamma_{\mathrm{diff}}(1+\kappa)}\right]\Pr_{\mathrm{FS}}^{\mathrm{diff}}(r)\,dr. \tag{8}$$

We compare these metrics to what would happen in a pure batch auction that executes at the end of the epoch. Here, all bidders will strategically wait until the epoch end and bid based on the final value of $p_{\mathrm{ex},T_{\mathrm{blk}}}$, but (since every trader is infinitely small) they will choose a limit price equal to their true honest values. Rank r trades if $p_{\mathrm{ex},T_{\mathrm{blk}}} \geq p(r)-\varepsilon(r)$. Since $p_{\mathrm{ex},T_{\mathrm{blk}}} - p_{\mathrm{ex},0}$ is a Brownian motion at an exponential stopping time, it follows a Laplace distribution, so the probability of this trade is

$$\Pr_{\mathrm{batch}}^{\mathrm{diff}}(r) = \frac{1}{2}\exp\{-\gamma_{\mathrm{diff}}[p(r)-\varepsilon(r)-p_{\mathrm{ex},0}]\}. \tag{9}$$

If r_{batch} is the rank of the trader choosing a limit price equal to the batch auction's clearing price, the uniform batch price is $p_{\mathrm{batch}}^{\mathrm{diff}} = p(r_{\mathrm{batch}}) = p_{\mathrm{ex},T_{\mathrm{blk}}} + \varepsilon(r_{\mathrm{batch}})$. Because $p'(r) = -\kappa\varepsilon'(r)$, $p(r)+\kappa\varepsilon(r)$ is constant. Conditional on rank r executing, the terminal public price exceeds $p(r)-\varepsilon(r)$ by an exponentially distributed amount with mean $\frac{1}{\gamma_{\mathrm{diff}}}$, so $\mathbb{E}[p_{\mathrm{batch}}^{\mathrm{diff}} - p_{\mathrm{ex},T_{\mathrm{blk}}} \mid r \text{ executes}] = \varepsilon(r) - \frac{1}{\gamma_{\mathrm{diff}}(1+\kappa)}$. Utilitarian welfare $\mathcal{W}_{\mathrm{batch}}^{\mathrm{diff}}$ and AMM gains $\mathcal{R}_{\mathrm{batch}}^{\mathrm{diff}}$ are given by

$$\mathcal{W}_{\mathrm{FS}}^{\mathrm{diff}} = 2e^{-1/(1+\kappa)}\mathcal{W}_{\mathrm{batch}}^{\mathrm{diff}} \qquad \text{and} \qquad \mathcal{R}_{\mathrm{FS}}^{\mathrm{diff}} = 2e^{-1/(1+\kappa)}\mathcal{R}_{\mathrm{batch}}^{\mathrm{diff}}. \tag{10}$$

So long as the mean gains from trade are nonnegative for the AMM, FirstSwap dominates a batch auction for both welfare and AMM gains iff

$$\frac{1}{1+\kappa} < \log 2 \quad \Longleftrightarrow \quad \kappa > \frac{1}{\log 2} - 1 \approx 0.4427. \tag{11}$$

A large value of κ is the natural assumption in financial markets, since the common-value variation $p_{\mathrm{ex},t}$ in an asset's value should be much larger than variation in agents' private valuations ε. We note that a FIFO/fair ordering protocol would yield exactly the same allocation as FS under gradual price diffusion, because the FS equilibrium is in fact providing price execution that is strictly better for earlier trades and no trades are rejected in equilibrium.

Optimal Trading Strategies with a Jump. Next, we analyze our protocol when the price $p_{\mathrm{ex},t}$ experiences a single large jump instead of gradually changing over time, taking on a single value

$$p_{\mathrm{ex,jump}} \equiv p_{\mathrm{ex},0} + \Delta p_{\mathrm{ex}}.$$

In this setting, many traders will submit their orders at once, and we model the arrival times of these orders as uniformly random. We examine two competitor mechanisms to FirstSwap: the batch auction we examined previously and the incentive-compatible protocol of Chan, Wu, and Shi [11] that relies on fair ordering of transactions that we will call FIFO. In our model, each trader chooses its limit price not knowing the speed with which the trade will be received by validators. A rational trader maximizes its expected utility across possible realizations where the trade arrives early or late.

Let $p^{\text{jump}}_{\text{bid,FS}}(r)$ denote the equilibrium bid of a trader with private valuation $\varepsilon(r)$ and $\Pr_{\text{fill}}(p^{\text{jump}}_{\text{bid,FS}}(r))$ the probability at equilibrium that the trade is approved. The optimal bid $p^{\text{jump}}_{\text{bid,FS}}(r)$ of type r maximizes over the choice of p_{bid}

$$[p_{\text{ex,jump}} + \varepsilon(r) - p_{\text{bid}}] \Pr_{\text{fill}}(p_{\text{bid}}). \tag{12}$$

In a "regular" equilibrium where limit prices are strictly decreasing in the rank r and $\Pr_{\text{fill}}(p_{\text{bid}})$ is smooth, the limit for any trader choosing $\Pr_{\text{fill}}(p_{\text{bid}}) \in (0,1)$ satisfies a first-order condition

$$p^{\text{jump}}_{\text{bid,FS}}(r) = p_{\text{ex,jump}} + \varepsilon(r) - \frac{\Pr_{\text{fill}}(p^{\text{jump}}_{\text{bid,FS}}(r))}{\Pr'_{\text{fill}}(p^{\text{jump}}_{\text{bid,FS}}(r))}. \tag{13}$$

In the FirstSwap equilibrium, every trader's limit price is strictly less than its true valuation of the asset. Increasing your limit raises the odds of the trade being executed, but a trader with a higher limit pays more if the trade executes. In a batch auction or FIFO, a small trader's limit price has no impact on the price at which it can trade and is only used to determine if the trader would accept the price it is offered. Because traders should accept a trade precisely when they pay less than their true valuation, traders rationally choose a limit price equal to their true valuation in both mechanisms. Choosing a limit price equal to the trader's true valuation $p^{\text{jump}}_{\text{bid,batch}}(r) = p_{\text{ex,jump}} + \varepsilon(r)$ is a dominant strategy for both of these mechanisms.

Welfare and Revenue Next, we compare welfare and AMM gains from trade across our three mechanisms: $\mathsf{m} \in \{FIFO, FS, \textit{batch auction}\}$. After a realized jump Δp_{ex}, welfare and gains have the following simple expressions, where we denote $\varepsilon_{virt}(r) = \varepsilon(r) + r\varepsilon'(r)$ following Myerson's [29] similar concept of a virtual value used to study revenue in auctions.

Lemma 2 (Allocation representation after a jump). *When all traders are rational, welfare and AMM gains from trade are respectively given by*

$$\mathcal{W}^{\text{jump}}_{\mathsf{m}} = \int_0^1 \varepsilon(r) \Pr{}^{\text{jump}}_{\mathsf{m}}(r)\, dr; \quad \mathcal{R}^{\text{jump}}_{\mathsf{m}} = \int_0^1 \varepsilon_{virt}(r) \Pr{}^{\text{jump}}_{\mathsf{m}}(r)\, dr. \tag{14}$$

The proof is in Appendix B.

We use this lemma to compare welfare and revenue across our three mechanisms. Under conditions we state in Lemma 3, the equilibria induced by these mechanisms can be ordered with a "single crossing condition." When this condition

holds, there is a cutoff r^* where one mechanism has a higher probability of trade execution for $r > r^*$ and the other for $r < r^*$.

Lemma 3 (Single-crossing welfare and revenue decompositions). *Consider two mechanisms* $\mathsf{m}, \mathsf{n} \in \{\text{FS}, \text{FIFO}, \text{batch}\}$*, with a difference in trade probabilities* $d_{mn}(r) = \Pr_{\mathsf{m}}^{\text{jump}}(r) - \Pr_{\mathsf{n}}^{\text{jump}}(r)$ *and volumes* $\Delta Vol_{\mathsf{mn}}^{\text{jump}} = \int_0^1 (\Pr_{\mathsf{m}}^{\text{jump}}(r) - \Pr_{\mathsf{n}}^{\text{jump}}(r))\, dr$*. If the two mechanisms satisfy the single-crossing condition with crossing point* r^**, welfare and AMM gains satisfy*

$$\mathcal{W}_{\mathsf{m}}^{\text{jump}} - \mathcal{W}_{\mathsf{n}}^{\text{jump}} = \underbrace{\int_0^1 [\varepsilon(r) - \varepsilon(r^*)] d_{mn}(r)\, dr}_{\textit{selection effect}} + \underbrace{\varepsilon(r^*) \Delta Vol_{\mathsf{mn}}^{\text{jump}}}_{\textit{volume effect}}, \tag{15}$$

$$\mathcal{R}_{\mathsf{m}}^{\text{jump}} - \mathcal{R}_{\mathsf{n}}^{\text{jump}} = \underbrace{\int_0^1 [(\varepsilon_{virt}(r) - \varepsilon_{virt}(r^*)] d_{mn}(r)\, dr}_{\textit{selection effect}} + \underbrace{\varepsilon_{virt}(r^*) \Delta Vol_{\mathsf{mn}}^{\text{jump}}}_{\textit{volume effect}}.$$

The batch equilibrium always has the single crossing property relative to FIFO or FirstSwap. Sufficient conditions for FirstSwap to have the single-crossing property relative to FIFO are: 1. F_ε *has a smooth nondecreasing hazard rate* $f_\varepsilon(e)/(1 - F_\varepsilon(e))$*; and 2.* p *is* C^3 *with* $p', p'' > 0$*, and* $1 \le \mathcal{C}(x) := \frac{p'(x)p'''(x)}{p''(x)^2} \le 2$ *is a nondecreasing function.*

In the above decomposition, the "selection effect" term always yields higher welfare and AMM gains for a batch auction than for FirstSwap or FIFO, since a batch auction perfectly allocates the asset to agents with the highest ε values. Because the FirstSwap equilibrium features limit prices strictly below agents' true valuations, it increases the probability that a given price will result in an executed trade. Under the conditions in the lemma, this results in a positive selection effect for FirstSwap relative to FIFO, since high-valuation agents will rationally choose limit prices at which they pay to increase their chances of trading.

Lemma 4 (Jump size, volume, and welfare rankings).
(i) Let $\Delta p_{\text{mid}} = p_{\text{ex,jump}} - p(1/2)$*. For any* Δp_{mid}*,* $Vol_{\text{FIFO}}^{\text{jump}} \ge \max\{Vol_{\text{batch}}^{\text{jump}}, Vol_{\text{FS}}^{\text{jump}}\}$*. Batch's single crossing with either mechanism is at* $\varepsilon^* = \varepsilon(Vol_{\text{batch}}^{\text{jump}}) = p(Vol_{\text{batch}}^{\text{jump}}) - p_{\text{ex,jump}}$*. Batch has higher welfare and AMM gains than FIFO for* $\Delta p_{\text{mid}} \ge 0$*.*
(ii) Small jumps. Under additional regularity conditions in Appendix D, if $0 \le \Delta p_{\text{mid}} < (\frac{1}{2f_\varepsilon(0)} + p'(1/2)/4)\sqrt{[p(1/2) - p(1/4)]/(24p'(1/2))}$*, and* $p(1/2) - p(1/4) \ge 2F_\varepsilon^{-1}(3/4)$ *then batch has higher welfare and AMM gains than FirstSwap.*
(iii) Large jumps. If $p_{\text{ex,jump}} \ge p(1) + \frac{1}{2}p'(1)$*, then FirstSwap has higher welfare and AMM gains than FIFO. The proofs are in Appendices C and D.*

Equivalent allocations with many validators. Our economic analysis has assumed that orders submitted at the same time are received by a single validator in random order. We show in the appendix that the equilibrium of any of our protocols is equivalent in terms of the probability of trade, welfare

and AMM gains to one with many validators if every validator's view is used in selecting which transactions to execute at the end of the epoch. For FS, our equivalence result applies to any version of our protocol satisfying three conditions: (i) accepted orders execute in ascending bid order at marginal AMM prices, without exceeding their submitted limits; (ii) against the equilibrium bids of other traders, bids arbitrarily close to $p(0)$ from above have positive execution probability; and (iii) some finite bid guarantees execution. For FS, this result is akin to the revenue equivalence theorem of Myerson [29] for standard auctions, showing that within a certain family of mechanisms implementation details do not impact the economic allocation of resources. The proof is in Appendix E.

Theorem 1 (Validator-count equivalence). *For each of the mechanisms batch, FIFO, and FirstSwap, any regular monotone equilibrium in a protocol with multiple validators yields the same allocation, welfare, and AMM revenue as in the one-validator case. For batch and FIFO, the bidding strategies are also identical, but the optimal bidding strategies for FirstSwap can vary with the number of validators despite yielding the same allocation.*

3 The FirstSwap Decentralized Protocol

In this section, we describe the decentralized FirstSwap protocol, a distributed implementation of the abstract FirstSwap design presented in Section 2 over a network of n validators. Due to space limitations, a glossary, algorithms in pseudocode and several (easy) proofs are deferred to Appendix F.

The fault model. We assume at most $f \leq \lfloor (n-1)/3 \rfloor$ validators are Byzantine; typically $f = \lfloor (n-1)/3 \rfloor$, though other thresholds may be used. The numbers of all traders and Byzantine traders are unknown. Honest parties follow the protocol while Byzantine parties may deviate arbitrarily. We assume unforgeable digital signatures and work in the standard partially synchronous model [16]: there exists an unknown Global Stabilization Time (GST) and a known bound $\Delta_{\text{net}} > 0$ such that after GST any message sent at time t is received by $\max(t, \text{GST}) + \Delta_{\text{net}}$. Safety holds at all times; liveness is guaranteed only after GST.

Protocol structure. The protocol operates in fixed-duration periods called *epochs*; unless noted otherwise, we focus on a single epoch. Within each epoch the protocol proceeds in two stages: a consensus-free *pre-approval* stage, followed by a finalization stage that runs a single BFT consensus instance. For liveness, we assume the epoch is long enough for traders to submit trade requests and receive validator responses, followed by a period sufficient for one consensus instance.

Scaled admissibility. The key technical challenge is ensuring that all independently pre-approved requests are jointly admissible, without any coordination among validators. To this end, each validator pre-approves only a scaled fraction η of the available capacity at each price level. For n validators, up to f dishonest, and quorum size $q := n-f$, define the approval fraction

$$\eta := \frac{q-f}{n-f} \ .$$

For conventional BFT parameters ($n = 3f+1$, $q = 2f+1$), $\eta = \frac{f+1}{2f+1} \approx \frac{1}{2}$; larger quorums (e.g. $q = \lceil \frac{4n+1}{5} \rceil$) yield $\eta \approx \frac{4}{5}$. Recall that a set S of trade requests is admissible if they can all be satisfied simultaneously within their price limits (Definition 1). A scaled η-admissible criterion is defined as follows:

Definition 2 (η-admissible). *The set $S = \{(\mathsf{amt}_i, \mathsf{limit}_i)\}$ is* η-admissible *if the scaled set $\{(\mathsf{amt}_i/\eta, \mathsf{limit}_i)\}$ is admissible.*

Since $\eta < 1$, this is a strictly stronger requirement than admissibility, and the slack it introduces is precisely what allows each validator to enforce admissibility locally and independently, without coordinating with other validators, while still guaranteeing that all trades with approval certificates are jointly admissible.

Pre-Approval Protocol. During the pre-approval stage, traders submit trade requests $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$ to each validator and collect signed responses. A *pre-approval* $\sigma_v(\mathsf{tx})$ from validator v is a signed message $\langle \mathsf{APPROVE}, \mathsf{tx} \rangle$. A collection of pre-approvals for tx from q distinct validators constitutes an *approval certificate* $\pi(\mathsf{tx})$, where q is a Byzantine quorum threshold chosen so that any two quorums intersect in at least one honest validator (e.g. $q = n-f$ in standard BFT settings). A trader that obtains $\pi(\mathsf{tx})$ may forward it to validators. Each validator v maintains a local set A_v of trade requests it has pre-approved and a set Π_v of approval certificates it has seen. Both sets are empty at the start of an epoch. A validator stores its first response to each tx and replays it on retransmission. Validator v pre-approves tx if and only if $\mathsf{A}_v \cup \{\mathsf{tx}\}$ is η-admissible (Definition 2). If so, v returns $\sigma_v(\mathsf{tx})$ and updates $\mathsf{A}_v := \mathsf{A}_v \cup \{(\mathsf{tx}, \sigma_v(\mathsf{tx}))\}$. Scaled admissibility is checked locally by each validator without any coordination between validators. All trade requests holding approval certificates, each admitted individually by validators in the above decentralized manner, can be satisfied within their price limits, without relying on consensus.

Finalization Protocol. Throughout the epoch, traders and validators forward signed pre-approvals and approval certificates they observe to all validators. At the end of the epoch, each validator v forms and signs a *mini-block* $B_v = (\mathsf{A}_v, \Pi_v)$. An honest validator includes in A_v all pre-approvals it signed before the finalization cutoff. It also includes in Π_v every valid certificate received before this cutoff. Anyone can verify the validator did not surpass its scaled pre-approval capacity. More formally:

Definition 3 (Well-formed mini-block). *A mini-block is* well-formed *if: (i) $\{\mathsf{tx} : \mathsf{tx} \in \mathsf{A}_v\}$ is η-admissible and each pre-approval is correctly signed by v, and (ii) every $\pi(\mathsf{tx}) \in \Pi_v$ is correctly formed.*

Each validator submits its signed mini-block as input to a single MVBA consensus instance. The MVBA protocol must satisfy an *external validity* condition such that a decided value B satisfies the predicate Valid below:

Definition 4 (Valid). $\mathsf{Valid}(\mathcal{B})$ *holds if (i) $\mathcal{B}$ contains well-formed mini-blocks from at least $M_B \geq n-f$ distinct validators, and (ii) for every $\pi(\mathsf{tx})$ appearing in $\mathcal{B}$ and every validator v whose mini-block $B_v = (\mathsf{A}_v, *)$ is in $\mathcal{B}$, if v's signature appears in $\pi(\mathsf{tx})$ then the same signed pre-approval for tx must appear in A_v.*

External validity provides three assurances: each A_v must satisfy η-admissibility to be counted; a trade request pre-approved by $n-f$ honest validators appears in at least M_B-f included mini-blocks; and a validator that attempts to hide a signed pre-approval it issued may be excluded from $\mathcal{B}$ via the cross-check.

In addition to external validity, the MVBA must uphold standard consensus properties: *agreement* (no two honest validators decide different values), and *termination* (every honest validator eventually decides after GST, provided all honest validators submit a valid input). Beyond these properties, the MVBA protocol is treated as a black box. Once the MVBA decides $\mathcal{B}$, an ordered set of trades is induced deterministically:

Pre-approved inclusion. Every trade request tx with signatures from at least ηM_B distinct validators in $\mathcal{B}$ is included.

AC inclusion. Every trade request tx for which a valid $\pi(\mathsf{tx})$ appears in $\mathcal{B}$ is included, unless already covered above.

Execution ordering. All included trade requests are ordered in non-decreasing order of limit, with ties broken deterministically by trade request identifier. Every included trade is guaranteed to be satisfiable within its price limit.

Security Analysis. FirstSwap guarantees the following desired properties:

f-Agreement. If at most f validators are Byzantine, all honest validators output the same final block at each epoch. (Proof deferred to Appendix F.)

f-Price validity. Suppose at most f validators are Byzantine. Let T be the ordered set of trades in the final block. Then T is admissible (Definition 1).

(f,η)-Pre-approval liveness. Let S_{any} be the set of trade requests sent to any honest validator during the epoch, and $S_{\text{all}} \subseteq S_{\text{any}}$ the set sent to all honest validators. If S_{any} is η-admissible, then under synchrony every trade request in S_{all} obtains an approval certificate.

f-Pre-approval inclusion. If either (i) a trade request tx has pre-approvals from q honest validators, or (ii) an $\pi(\mathsf{tx})$ is received by $f+1$ honest validators, then tx must be included in the final block.

Accountable inclusion. There exists a function J such that if tx obtains an approval certificate and is not included in the epoch's final block, then there exists at least one validator v and evidence w with $v \in J(w)$. Moreover, for all honest validators v and all w, $v \notin J(w)$. (Proof deferred to Appendix F.)

Validity. Let V^* be the set of creators of the mini-blocks in the decided proposal $\mathcal{B}$, and let $M_B := |V^*|$. External validity gives $M_B \geq n - f$ and ensures that every $B_v = (\mathsf{A}_v, \Pi_v)$ for $v \in V^*$ is well-formed. Let S be the set of trade requests included by the two inclusion rules of the protocol. We use the notation $\mathsf{limit}(\mathsf{tx})$ and $\mathsf{amt}(\mathsf{tx})$ to denote the limit and amount of a trade request tx.

Lemma 5 (Representation of included requests). *Every trade request in S appears in A_v for at least ηM_B validators $v \in V^*$.*

Proof. If a request is included by the pre-approved inclusion rule, the claim follows from its threshold ηM_B. Otherwise, its approval certificate has q signers, at least $q + M_B - n$ of whom belong to V^*. Dividing this lower bound by the number of selected validators gives

$$\frac{q + M_B - n}{M_B} = 1 - \frac{n - q}{M_B} \geq 1 - \frac{n - q}{n - f} = \frac{q - f}{n - f} \geq \eta.$$

The first inequality uses $n - q \geq 0$ and $M_B \geq n - f$; the last uses the assumed bound on η. Thus at least ηM_B selected validators signed the request. The external-validity cross-check requires its signed approval to appear in A_v for every such signer. □

Lemma 6 (Admissibility of included requests). *The set S is admissible.*

Proof. Fix a limit ℓ occurring in S and write $X(\ell) = \sum_{\mathsf{tx} \in S : \mathsf{limit}(\mathsf{tx}) \leq \ell} \mathsf{amt}(\mathsf{tx})$. For each $v \in V^*$, similarly write $X_v(\ell) = \sum_{\mathsf{tx} \in \mathsf{A}_v : \mathsf{limit}(\mathsf{tx}) \leq \ell} \mathsf{amt}(\mathsf{tx})$. By Lemma 5, each request counted in $X(\ell)$ appears in at least ηM_B selected approved sets. Counting its amount in each gives

$$\eta M_B X(\ell) \leq \sum_{v \in V^*} X_v(\ell). \tag{16}$$

For every v with $X_v(\ell) > 0$, let ℓ_v be its largest limit at most ℓ among the requests counted by $X_v(\ell)$. Since A_v is η-admissible, $p(X_v(\ell)/\eta) \leq \ell_v \leq \ell$. Let $Y = \max_{v \in V^*} X_v(\ell)/\eta$, treating an empty prefix as zero. Then $p(Y) \leq \ell$. Dividing Equation (16) by $\eta M_B > 0$ gives

$$X(\ell) \leq \frac{1}{\eta M_B} \sum_{v \in V^*} X_v(\ell) \leq \frac{1}{\eta M_B} \sum_{v \in V^*} \eta Y = Y. \tag{17}$$

Monotonicity of p gives $p(X(\ell)) \leq p(Y) \leq \ell$. This holds for every limit in S, which is exactly admissibility. □

Corollary 1. *The protocol satisfies f-Price validity.*

Liveness. The following proofs use the synchrony and epoch-duration assumptions from above. In particular, messages sent during the reservation period arrive before its cutoff, and MVBA termination produces a final block.

Theorem 2. *The protocol satisfies (f, η)-pre-approval liveness.*

Proof. Assume S_{any} is η-admissible. Immediately before an honest validator v processes $\mathsf{tx} \in S_{\text{any}}$, the set $\mathsf{A}_v \cup \{\mathsf{tx}\}$ is a subset of S_{any}. Removing requests only decreases every prefix amount, so this subset remains η-admissible. Validator v therefore signs an approval for tx. Every request in S_{all} is approved by every honest validator. There are at least $n - f = q$ honest validators, and synchrony delivers their responses during the epoch. Thus the trader obtains an AC. □

Theorem 3. *The protocol satisfies f-pre-approval inclusion.*

Proof. $n \geq 4f+1$, $q = n - f$, and $\eta = \frac{q-2f}{n-f}$, First suppose tx has approvals from q honest validators. Exactly $n - M_B$ validators lie outside V^*, so at least $q + M_B - n \geq \eta M_B$ of these honest approvers belong to V^*, by the bound in the proof of Lemma 5. Their complete approved sets contain their signed approvals for tx, meeting the pre-approved inclusion threshold. Now suppose an AC for tx is received before the cutoff by $f+1$ honest validators. Since $n - M_B \leq f$, at least one of these validators belongs to V^*. The honest mini-block construction places the AC in that validator's Π_v, so the AC inclusion rule includes tx. □

Corollary 2. *For any n, f, q, η such that $n \geq 3f+1$, $q = n - f$ and $\eta \leq \frac{q-f}{n-f}$, the FirstSwap protocol with parameters n, f, q, η satisfies f-agreement, f-price validity, (f, η)-pre-approval liveness, f-pre-approval inclusion, and accountable inclusion.*

Stronger inclusion guarantees for pre-approved trades. By parameterizing the protocol with $n = 4f+1$, $q = n - f$, and $\eta = \frac{q-2f}{n-f} = \frac{f+1}{3f+1} \approx \frac{1}{3}$, we can strengthen pre-approval inclusion: if a trade request tx is pre-approved by q validators (not necessarily all honest), then tx is guaranteed to be included in the final block. Moreover, with these parameters, accountable inclusion is also strengthened: at least $f+1$ validators (not just one) is caught with evidence if a pre-approved trade is not included in the final block. The proofs of these strengthened properties are in Appendix F. All other security properties are preserved because these parameters satisfy the conditions of Corollary 2.

4 Maximal Fill Volume

The pre-approval liveness guarantee of the FirstSwap protocol in Section 3 requires the set of trade requests S_{any} sent to any honest validator to be η-admissible. For example, let $p(x) = x$ and suppose two traders submit requests $\mathsf{tx}_1 = (\mathsf{amt}_1 = 1, \mathsf{limit}_1 = 5)$ and $\mathsf{tx}_2 = (\mathsf{amt}_2 = 2, \mathsf{limit}_2 = 5)$ to all honest validators. For $\eta = 1/2$, the set $S_{\text{any}} = \{\mathsf{tx}_1, \mathsf{tx}_2\}$ is *not* η-admissible, so the liveness guarantee does not apply. No trade is guaranteed to be included. Yet each request is individually η-admissible, as is the full amount of tx_1 combined with half the amount of tx_2, suggesting that the protocol leaves some liveness unexploited. Moreover, a malicious trader can spam the system by submitting a large request to just one honest validator. This makes S_{any} non-η-admissible, leaving the pre-approval liveness property unable to guarantee inclusion of either the honest requests or the spammer's request. In contrast, a standard blockchain-based DEX would include either the honest requests or the spammer's request in a block: if the spammer displaced the honest requests, it would at least pay fees for its own. If no trade requests are pre-approved, no trade requests are guaranteed to be included in the block.

We propose a simple extension that addresses these limitations by allowing partial fills in the final block. A *trade* is denoted by $\mathsf{tx} = (\mathsf{tx}, \mathsf{fill})$, where $\mathsf{tx} =$

$(\mathsf{amt}, \mathsf{limit})$ is the original *trade request* and $\mathsf{fill} \in [0, \mathsf{amt}]$ is the amount filled. The trader specifies the maximum amount amt he is willing to buy, and the protocol determines the actual amount bought fill. Validators continue to use the η-admissibility condition to pre-approve trade requests, so pre-approval liveness remains as it is. The final block fully fills all pre-approved requests (preserving pre-approval inclusion) and partially fills residual requests while maintaining admissibility (improving livness). The residual fill algorithm solves a fractional knapsack problem variant, maximizing total filled volume.

With partial fills, a set of trades is considered admissible based on their filled amounts fill rather than their requested amounts amt. η-admissibility continues to use the full amounts.

Definition 5 (Admissible trades). *Let* $T = \{(\mathsf{tx}_i = (\mathsf{amt}_i, \mathsf{limit}_i), \mathsf{fill}_i)\}_{i=1}^{m}$. *The set* T *is* admissible *if* $0 \leq \mathsf{fill}_i \leq \mathsf{amt}_i$ *for every* i *and, for every* i *with* $\mathsf{fill}_i > 0$, $p\left(\sum_{j:\mathsf{limit}_j \leq \mathsf{limit}_i} \mathsf{fill}_j\right) \leq \mathsf{limit}_i$.

We keep the pre-approval rule and the consensus validity conditions of Section 3. To make rejected requests available for residual inclusion, each honest validator v also records every valid request it receives before the finalization cut-off in a set $\mathcal{R}_v$, and includes this set in its signed mini-block $B_v = (\mathsf{A}_v, \Pi_v, \mathcal{R}_v)$. Well-formedness checks the validity and trader authorization of these additional requests. Recording does not reserve capacity: the η-admissibility check still applies only to A_v.

Let $\mathcal{B}$ be the collection of mini-blocks decided by consensus, V^* their distinct creators, and $M_B = |V^*| \geq n - f$. Let U contain every distinct request appearing in $\mathcal{B}$, including recorded but unapproved requests. Write G for the requests selected by either of the first two inclusion rules below. As in the main protocol, use $q = n - f$ and $\eta = (n - 2f)/(n - f)$. The final ordered block is constructed as follows.

Pre-approved inclusion. Every trade request $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$ with signatures from at least $M_B - f$ distinct validators in $\mathcal{B}$ is included with $\mathsf{fill} = \mathsf{amt}$.

AC inclusion. Every trade request $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$ for which a valid $\pi(\mathsf{tx})$ appears in $\mathcal{B}$ is included with $\mathsf{fill} = \mathsf{amt}$, unless already covered above.

Residual inclusion. Then consider every remaining trade request in U in a fixed deterministic order, and assign it the maximum fill amount, up to its requested amount, that preserves admissibility. Initially set the fills of all other requests to zero. At each step keep all other fills fixed, and check admissibility against every prefix in U, including prefixes ending at requests that currently have zero fill. Equivalently, assign the largest $x \in [0, \mathsf{amt}_i]$ for which replacing $\mathsf{fill}_i = 0$ by x keeps the complete fill vector admissible.

Execution ordering. Discard requests with zero fill and order all remaining requests in non-decreasing order of limit, breaking ties deterministically by request identifier.

Feasible fills and capacity. For a set of requests U, write

$$\mathcal{F}(U) = \left\{ \mathbf{fill} \in \prod_{i=1}^{|U|} [0, \mathsf{amt}_i] : \{(\mathsf{tx}_i, \mathsf{fill}_i)\}_{\mathsf{tx}_i \in U} \text{ is admissible} \right\}. \tag{18}$$

Define

$$R(U) = \max_{\mathbf{fill} \in \mathcal{F}(U)} \sum_i \mathsf{fill}_i, \qquad R(\emptyset) = 0.$$

Security analysis. The strengthened liveness property achieved by this modified protocol is stated in Definition 6 and proved in Theorem 4. It guarantees that the total filled volume of the final block is at least as large as the maximum feasible fill of all valid trade requests sent to all honest validators before the finalization cutoff. This holds even if the set of requests sent to honest validators is not η-admissible. In particular, even if a malicious trader spams some honest validators, either the honest trade requests or an equivalent volume of the spammer's trade requests will be filled in the final block. This restores a liveness guarantee similar to that of a standard blockchain-based DEX, in which the spammer would pay fees for the spammer's trade requests.

Price validity is preserved by construction: By Lemma 6, G is admissible when $\mathsf{fill} = \mathsf{amt}$ for every request. Assigning zero fills to the remaining requests changes no filled amount and preserves admissibility. Each residual step preserves admissibility by construction and leaves pre-approved fills unchanged. Hence the complete final fill vector is admissible. All other security properties of the original protocol remain unchanged.

Definition 6 (f-Maximal volume filled). *Suppose that the network is synchronous and at most f validators are Byzantine during an epoch. Let S be the set of valid trade requests sent to all honest validators in time to arrive before the finalization cutoff. Let T denote the trades in the final block of the epoch. The protocol satisfies f-Maximal volume filled if*

$$\sum_{(\mathsf{tx}_i, \mathsf{fill}_i) \in T} \mathsf{fill}_i \geq R(S). \tag{19}$$

Lemma 7 (Maximal fills maximize volume). *Starting from the full fills of G, the protocol's residual rule produces a feasible fill of U with total volume $R(U)$.*

Proof. By Lemma 6, the initial full fills of G are admissible; zero fills for the remaining requests and each subsequent residual step preserve admissibility. Moreover, no final fill can be increased while the other fills are fixed. Pre-approved requests are already full. For a residual request, any increase that were admissible at the end would also have been admissible when it was processed: all other fills were then no larger, and decreasing fills preserves admissibility. Such an increase would contradict the greedy rule.

Suppose, for a contradiction, that another admissible fill vector $\widehat{\mathbf{fill}}$ has larger total volume. Write $D = \sum_j (\widehat{\mathsf{fill}}_j - \mathsf{fill}_j) > 0$, and choose a request tx_i with largest limit among those for which $\widehat{\mathsf{fill}}_i > \mathsf{fill}_i$. For every $\ell \geq \mathsf{limit}_i$, requests with limit above ℓ have $\widehat{\mathsf{fill}}_j \leq \mathsf{fill}_j$, so

$$\sum_{j:\mathsf{limit}_j \leq \ell} (\widehat{\mathsf{fill}}_j - \mathsf{fill}_j) = D - \sum_{j:\mathsf{limit}_j > \ell} (\widehat{\mathsf{fill}}_j - \mathsf{fill}_j) \geq D. \tag{20}$$

Increase fill_i by $\varepsilon = \min\{D, \widehat{\mathsf{fill}}_i - \mathsf{fill}_i\} > 0$. The increased fill is at most $\widehat{\mathsf{fill}}_i \leq \mathsf{amt}_i$. Prefixes below limit_i are unchanged, and (20) shows that every prefix at or above limit_i has total fill at most the corresponding prefix of $\widehat{\mathbf{fill}}$.

For each such limit ℓ, the alternative prefix contains a positive fill because $\widehat{\mathsf{fill}}_i > 0$. Let ℓ' be the largest limit of any positive-filled request in that prefix. Admissibility of $\widehat{\mathbf{fill}}$ gives

$$p\left(\sum_{j:\mathsf{limit}_j \leq \ell} \widehat{\mathsf{fill}}_j\right) = p\left(\sum_{j:\mathsf{limit}_j \leq \ell'} \widehat{\mathsf{fill}}_j\right) \leq \ell' \leq \ell.$$

By monotonicity of p, the increased fill vector therefore remains admissible, including at limit_i. This contradicts the impossibility of increasing any final fill. No admissible vector has larger total volume, so the protocol attains $R(U)$. □

Theorem 4 (f-Maximal volume filled). *Under the synchrony and epoch-duration assumptions of Section 3, the extended protocol satisfies f-Maximal volume filled as defined in Definition 6.*

Proof. Let S be the set in the property definition. Every honest validator records each $\mathsf{tx} \in S$ and includes it in its mini-block. Since V^* contains at least $M_B - f \geq n - 2f > 0$ honest validators, every request in S appears in U. Thus $S \subseteq U$.

Any admissible fill of S extends to an admissible fill of U by assigning zero to $U \setminus S$. The added requests require no execution, and the filled amount in the prefix of every positive-filled request is unchanged. Hence its inequality in Definition 5 remains satisfied. It follows that $R(U) \geq R(S)$.

By Lemma 7, the protocol fills total volume $R(U)$. Discarding zero-filled requests changes no volume, so its final block T satisfies

$$\sum_{(\mathsf{tx}_i, \mathsf{fill}_i) \in T} \mathsf{fill}_i = R(U) \geq R(S).$$

□

5 Combining Buys and Sells

We extend our main protocol to allow for simultaneous buys and sells. Each Validator receives trade requests with both buys and sells and must decide when

it is able to pre-approve a request. We add $\mathsf{side} \in \{\mathsf{buy}, \mathsf{sell}\}$ to the existing trade request format: $\mathsf{tx} = (\mathsf{side}, \mathsf{amt}, \mathsf{limit})$. A naive way to handle both buys and sells would be to run two independent instances of the the protocol from Section 3, one for buys and one for sells. However, this approach approves would forego the opportunity of approving more buys when there are matching sells and conversely.

The challenge lies in determining when a validator can safely pre-approve additional buys based on the sell requests it has seen (and vice versa) without violating any limit prices. Towards this, we first propose an extended definition of admissibility for the bilateral setting with both buys and sells (Definition 7). Unlike the unilateral setting (Definition 1), bilateral admissibility involves two sets of trade requests: a set G of pre-approved trade requests that must be executed fully and a set $U \supseteq G$ of trade requests that can optionally be executed to fill enough liquidity to support the trades in G. We show a pre-approval rule and block finalization protocol (similar to those in Section 3) that ensures bilateral admissibility when G is the set of pre-approved trade requests and U is the set of all other trade requests reported in any finalized block that may or may not execute. Finally, we show in lemma 9 that when the set of trades in a block is bilaterally admissible, all pre-approved trades in G (and optionally some others in U) can be executed while satisfying all limit prices.

We extend our AMM curve $p(x)$ to allow for either positive or negative net purchases x and write $P(x) = \int_0^x p(u)\,du$. Let $c(\ell)$ be the inverse of $p(x)$, with $c(p(0)) = 0$ and both $c(.), p(.)$ continuous. Write $[z]_+ = \max\{z, 0\}$. At limit ℓ, the AMM offers supply $[c(\ell)]_+$ and demand $[-c(\ell)]_+$.

Definition 7 (Bilateral admissibility). *Let $G \subseteq U$ be a collection of trade requests. G is said to be bilaterally admissible in U if for every ℓ,*

$$\sum_{\substack{\mathsf{tx}_i \in G \ buys \\ \mathsf{limit}_i \le \ell}} \mathsf{amt}_i \le [c(\ell)]_+ + \sum_{\substack{\mathsf{tx}_j \in U \ sells \\ \mathsf{limit}_j \le \ell}} \mathsf{amt}_j, \tag{21}$$

$$\sum_{\substack{\mathsf{tx}_i \in G \ sells \\ \mathsf{limit}_i \ge \ell}} \mathsf{amt}_i \le [-c(\ell)]_+ + \sum_{\substack{\mathsf{tx}_j \in U \ buys \\ \mathsf{limit}_j \ge \ell}} \mathsf{amt}_j. \tag{22}$$

In our extended protocol, similar to Section 4, validators keep track of not only trade requests they pre-approved but also all trade requests that they have seen. An honest validator pre-approves only trade requests where it has sufficient capacity. As new trade requests are received that may increase capacity, it promotes old seen trade requests to pre-approved in the order received. When forming a mini-block, validators report separate lists of seen and pre-approved trade requests that can be used to build a final block. Similar to Section 4, trades included in the final block will have a parameter $\mathsf{fill}_i \in [0, \mathsf{amt}_i]$ denoting partial filling of trade requests. In addition, each trade will also be assigned price_i denoting the execution price of the trade. When we considered only buys, we had implcitly considered price_i to be the price determined by the AMM curve

after trades has been ordered by their limit prices. Our goal is to ensure that $\mathsf{price}_i \le \mathsf{limit}_i$ for all executed buys and $\mathsf{price}_i \ge \mathsf{limit}_i$ for all executed sells.

Modified Pre-Approval Protocol. Alongside the original pre-approval and certificate sets, validator v records trade requests $\mathcal{R}_v$ that it has seen. Before the finalization cutoff, it records and relays every valid trade request, including trade requests that fail the approval test. For each validator v, given the set of trade requests it has seen, define

$$\widehat{S}_v(\ell) = \sum_{\substack{\mathsf{tx}_i \in \mathcal{R}_v \text{ sells} \\ \mathsf{limit}_i \le \ell}} \mathsf{amt}_i, \qquad \widehat{B}_v(\ell) = \sum_{\substack{\mathsf{tx}_i \in \mathcal{R}_v \text{ buys} \\ \mathsf{limit}_i \ge \ell}} \mathsf{amt}_i.$$

Our validator's new rule for adding a trade request to its pre-approved set A_v is that it will pre-approve if doing so does not violate

$$\sum_{\substack{\mathsf{tx}_i \in \mathsf{A}_v \text{ buys} \\ \mathsf{limit}_i \le \ell}} \mathsf{amt}_i \le \eta \left([c(\ell)]_+ + \widehat{S}_v(\ell) \right) \quad \text{for every } \ell \tag{23}$$

$$\sum_{\substack{\mathsf{tx}_i \in \mathsf{A}_v \text{ sells} \\ \mathsf{limit}_i \ge \ell}} \mathsf{amt}_i \le \eta \left([-c(\ell)]_+ + \widehat{B}_v(\ell) \right) \quad \text{for every } \ell. \tag{24}$$

Note that this is the analogue of η-admissibility (Definition 2) for bilateral admissibility (Definition 7).

Modified Finalization Protocol. In our modified finalization protocol, each mini-block $B_v = (\mathsf{A}_v, \Pi_v, \mathcal{R}_v)$ additionally contains the recorded requests $\mathcal{R}_v$. Well-formedness verifies these requests and replaces the original η-admissibility by checking Equations (23) and (24). The quorum requirement and approval-certificate cross-check for block validity are as before. The inclusion and execution ordering steps are described below. First, if U is the set of all trade requests appearing in the block (including recorded and approved), note that all pre-approved transactions $G \subseteq U$ satisfy joint bilateral feasibility.

Lemma 8 (Admissibility of pre-approved requests). *Under* (23)–(24) *with $\eta = (q - f)/q$, the set G of trade requests with at least $M_B - f$ approvals among selected mini-blocks is bilaterally admissible.*

Proof. For each limit ℓ, let $X(\ell)$ be the total quantity of buys in G with limits at most ℓ, and let $X_v(\ell)$ be the corresponding quantity in A_v. Let $S_U(\ell)$ be the total quantity of sells in U with limits at most ℓ. Every trade request recorded as seen in a selected mini-block belongs to U, so $\widehat{S}_v(\ell) \le S_U(\ell)$ for every $v \in V^*$. Counting approvals and summing the local buy budgets gives

$$X(\ell) \le \frac{1}{M_B - f} \sum_{v \in V^*} X_v(\ell) \le \frac{\eta M_B}{M_B - f} \left([c(\ell)]_+ + S_U(\ell) \right) \le [c(\ell)]_+ + S_U(\ell).$$

This demonstrates Equation (21) for buys and an analogous argument demonstrates Equation (22) for sells. □

Next, we show how to compute the amounts filled for each trade request fill_i, given a block with trades U and pre-approved trades $G \subseteq U$ that are bilaterally admissible. Lemma 9 shows that all trades in G are completely filled, every filled trade can be assigned a price that satisfies its limit, and the AMM receives the correct revenue. This is because the bilateral admissibility ensures that the trades in U supply enough demand or supply to augment the AMM's own demand and supply to meet the supply or demand of the pre-approved trades in G. It also preserves monotonicity: a trade with a higher limit price will not be executed at a lower price than a trade with a lower limit price (monotonicity was implicit in Section 3 by ordering trades by limit prices).

Our protocol works in two steps. First, we compute the quantity of buys minus the quantity of sells in the set of pre-approved trades G. We then compute a price ℓ^* where executing all buys in $U \setminus G$ with limit prices below ℓ^* and all sells above ℓ^* (with potential partial fills at ℓ^*) allows all trades in G to execute. This results in a total set U^* of trades that actually execute. If U^* has more buys than sells so that the AMM is a net seller, sells in U^* are effectively added to the AMM's liquidity supply. Buys are executed in order of limit price, where each buy is greedily given the best possible price provided by the remaining sell liquidity as in our main protocol. Sells execute at their limit price. If U^* had more sells, we reverse this procedure so buys fill at their limit price and sells are given price improvement.

This pricing preserves the property in our main protocol that a trade's execution price is an increasing function of its limit price. Because requesting an attractive limit price lowers a trade request's chance of approval, this allows traders with a particularly high desire to trade to choose an unattractive limit price to increase their odds of acceptance. In our economic analysis, this allowed for a "selection effect" where traders who benefit the most from trading have the highest chance of their trades executing.

Input: Set of seen trade requests U and set of pre-approved transactions $G \subseteq U$, with G bilaterally admissible in U.

Output: Fill vector $\mathbf{fill} = (\mathsf{fill}_i)_{\mathsf{tx}_i \in U}$

Pre-approval inclusion. Set $\mathsf{fill}_i = \mathsf{amt}_i$ for every request in G and initialize all other fills to zero. Define the guaranteed net demand

$$g = \sum_{\mathsf{tx}_i \in G \text{ buys}} \mathsf{amt}_i - \sum_{\mathsf{tx}_i \in G \text{ sells}} \mathsf{amt}_i$$

and the right-continuous augmented liquidity curve

$$\widetilde{c}(\ell) = c(\ell) + \sum_{\substack{\mathsf{tx}_i \in U \setminus G \text{ sells} \\ \mathsf{limit}_i \le \ell}} \mathsf{amt}_i - \sum_{\substack{\mathsf{tx}_i \in U \setminus G \text{ buys} \\ \mathsf{limit}_i > \ell}} \mathsf{amt}_i. \tag{25}$$

Limit-matched-inclusion. Choose the cutoff $\ell^* = \inf\{\ell : \widetilde{c}(\ell) \ge g\}$. Among requests in $U \setminus G$, fully fill buys with $\mathsf{limit}_i > \ell^*$ and sells with $\mathsf{limit}_i < \ell^*$.

Leave buys with $\mathsf{limit}_i < \ell^*$ and sells with $\mathsf{limit}_i > \ell^*$ unfilled. Define the remaining imbalance

$$\Delta_{\ell^*} = g - \widetilde{c}(\ell^*) + \sum_{\substack{\mathsf{tx}_i \in U \setminus G \text{ sells} \\ \mathsf{limit}_i = \ell^*}} \mathsf{amt}_i.$$

If $\Delta_{\ell^*} > 0$, fill this quantity of sells with limit ℓ^*; if $\Delta_{\ell^*} < 0$, fill $-\Delta_{\ell^*}$ of buys with limit ℓ^*. Process these requests in a fixed order by identifier, assigning each at most amt_i until the required quantity is reached. Leave all other requests at the cutoff unfilled.

Output $\mathbf{fill} = (\mathsf{fill}_i)_{\mathsf{tx}_i \in U}$.

Lemma 9 (Limit-satisfying prices for bilateral admissibility). *If $G \subseteq U$ is bilaterally admissible and* **fill** *is as computed above, then there exists* **price** *such that:*

1. *all pre-approved trades are completely filled: for all $\mathsf{tx}_i \in G$, $\mathsf{fill}_i = \mathsf{amt}_i$;*
2. *all limits are respected: for all $\mathsf{tx}_i \in U$ such that $\mathsf{fill}_i > 0$, $\mathsf{side}_i = \mathsf{buy} \implies \mathsf{price}_i \le \mathsf{limit}_i$ and $\mathsf{side}_i = \mathsf{sell} \implies \mathsf{price}_i \ge \mathsf{limit}_i$;*
3. *the AMM's revenue is exactly $P(x)$ where $x := \sum_{\mathsf{tx}_i \in U\, buys} \mathsf{fill}_i - \sum_{\mathsf{tx}_i \in U\, sells} \mathsf{fill}_i$ is the net purchase quantity;*
4. *prices are monotonic: for all $\mathsf{tx}_i, \mathsf{tx}_j \in U$ such that $\mathsf{fill}_i > 0$ and $\mathsf{side}_i = \mathsf{side}_j$, $\mathsf{limit}_i \le \mathsf{limit}_j \implies \mathsf{price}_i \le \mathsf{price}_j$.*

Proof. We first justify that the cutoff ℓ^* exists and is finite. For this argument, let $a = \min(\{p(0)\} \cup \{\mathsf{limit}_i : \mathsf{tx}_i \in U\})$ and $b = \max(\{p(0)\} \cup \{\mathsf{limit}_i : \mathsf{tx}_i \in U\})$. Since $c(a) \le 0$ and every request has limit at least a, Equation (22) at a gives

$$\widetilde{c}(a^-) = c(a) - \sum_{\mathsf{tx}_i \in U \setminus G \text{ buys}} \mathsf{amt}_i \le g.$$

Similarly, $c(b) \ge 0$ and every request has limit at most b, so Equation (21) at b gives

$$g \le c(b) + \sum_{\mathsf{tx}_i \in U \setminus G \text{ sells}} \mathsf{amt}_i = \widetilde{c}(b).$$

For every $\ell < a$, strict monotonicity of $c = p^{-1}$ gives

$$\widetilde{c}(\ell) = c(\ell) - \sum_{\mathsf{tx}_i \in U \setminus G \text{ buys}} \mathsf{amt}_i < \widetilde{c}(a^-) \le g.$$

Thus the set defining the cutoff is nonempty and has no element below a, so its infimum lies in $[a, b]$. Second, we show that there are enough trades with limit exactly ℓ^* to fill the quantity $\pm\Delta_{\ell^*}$. The curve $\widetilde{c}$ is nondecreasing and right-continuous, with only upward jumps. Consequently,

$$\widetilde{c}(\ell^*-) \le g \le \widetilde{c}(\ell^*).$$

Here the left-hand value includes buys at the cutoff and excludes sells at the cutoff. Consequently,

$$-\sum_{\substack{\mathsf{tx}_i \in U \setminus G \text{ buys} \\ \mathsf{limit}_i = \ell^*}} \mathsf{amt}_i \le \Delta_{\ell^*} \le \sum_{\substack{\mathsf{tx}_i \in U \setminus G \text{ sells} \\ \mathsf{limit}_i = \ell^*}} \mathsf{amt}_i.$$

There is therefore enough volume at the cutoff to fill $\pm\Delta_{\ell^*}$ while respecting every request's quantity bound.

Price construction. Let U^* be the set of executed requests. Their net asset purchase is

$$x = \sum_{\mathsf{tx}_i \in U^* \text{ buys}} \mathsf{fill}_i - \sum_{\mathsf{tx}_i \in U^* \text{ sells}} \mathsf{fill}_i = c(\ell^*).$$

Define the AMM's executed volume with marginal price at most ℓ by

$$F_{\mathrm{AMM}}(\ell) = \left| \int_0^x \mathbf{1}_{\{p(v) \le \ell\}}\, dv \right|.$$

We now define two distribution functions that represent the set of buy and sell limit prices that that will execute as actual trades, including the liquidity supplied by the AMM.

$$F_{U^*,\mathrm{buy}}(\ell) = \sum_{\substack{\mathsf{tx}_i \in U^* \text{ buys} \\ \mathsf{limit}_i \le \ell}} \mathsf{fill}_i + \mathbf{1}_{\{x<0\}} F_{\mathrm{AMM}}(\ell), \tag{26}$$

$$F_{U^*,\mathrm{sell}}(\ell) = \sum_{\substack{\mathsf{tx}_i \in U^* \text{ sells} \\ \mathsf{limit}_i \le \ell}} \mathsf{fill}_i + \mathbf{1}_{\{x>0\}} F_{\mathrm{AMM}}(\ell), \tag{27}$$

The total buy and sell quantities are equal, i.e., $F_{U^*,\mathrm{buy}}(\infty) = F_{U^*,\mathrm{sell}}(\infty) = M$, and all liquidity $F_{\mathrm{AMM}}(\ell)$ is only added on one side to meet the quantity imbalance in U^* between buys and sells.

Define the inverse functions $F^{-1}_{U^*,\mathrm{buy}}(u)$ with $F^{-1}_{U^*,\mathrm{sell}}(u)$ where $F^{-1}(u) = \inf\{\ell : F(\ell) \ge u\}$.

If $x \ge 0$, assign $\mathsf{price}_i = \mathsf{limit}_i$ for each sell in U^*. Sort all buys in U^* by limit prices, treating buys with a tied limit price as a single larger buy. For the i-th buy, let $I_i := \left(\sum_{j=1}^{i-1} \mathsf{fill}_j, \sum_{j=1}^{i} \mathsf{fill}_j\right]$ be the interval it occupies in the sorted order. Set the i-th buy's price as

$$\mathsf{price}_i = \frac{1}{\mathsf{fill}_i} \int_{I_i} F^{-1}_{U^*,\mathrm{sell}}(u)\, du.$$

If $x < 0$, assign $\mathsf{price}_i = \mathsf{limit}_i$ for each buy in U^*. Sort all sells in U^* by limit prices, also merging tied limit prices into a single sell. For the i-th sell, let $I_i := \left(\sum_{j=1}^{i-1} \mathsf{fill}_j, \sum_{j=1}^{i} \mathsf{fill}_j\right]$ be the interval it occupies in the sorted order. Set the i-th sell's price as

$$\mathsf{price}_i = \frac{1}{\mathsf{fill}_i} \int_{I_i} F^{-1}_{U^*,\mathrm{buy}}(u)\, du.$$

A buyer pays $\mathsf{fill}_i\mathsf{price}_i$ and a seller receives $\mathsf{fill}_i\mathsf{price}_i$.

Limit validity. If we can show that $F^{-1}_{U^*,sell}(u) \le F^{-1}_{U^*,\mathrm{buy}}(u)$, then executing at the above price satisfies limits. This is because the i-th buy's price is exactly limit_i when $x < 0$ and when $x \ge 0$,

$$\mathsf{price}_i = \frac{1}{\mathsf{fill}_i}\int_{I_i} F^{-1}_{U^*,\mathrm{sell}}(u)\,du \le \frac{1}{\mathsf{fill}_i}\int_{I_i} F^{-1}_{U^*,\mathrm{buy}}(u)\,du = \mathsf{limit}_i,$$

Similarly, when $x < 0$, for the i-th sell,

$$\mathsf{limit}_i = \frac{1}{\mathsf{fill}_i}\int_{I_i} F^{-1}_{U^*,\mathrm{sell}}(u)\,du \le \frac{1}{\mathsf{fill}_i}\int_{I_i} F^{-1}_{U^*,\mathrm{buy}}(u)\,du = \mathsf{price}_i.$$

We prove the equivalent statement that $F_{U^*,\mathrm{buy}}(\ell) \le F_{U^*,\mathrm{sell}}(\ell)$ pointwise. For $\ell < \ell^*$, executed buys with limits at most ℓ are exactly the corresponding buys in G, and all available sells with limits at most ℓ in U fill completely. Since $c(\ell) \le x = c(\ell^*)$, the definition of F_{AMM} gives

$$\mathbf{1}_{\{x<0\}}F_{\mathrm{AMM}}(\ell) = 0, \qquad \mathbf{1}_{\{x>0\}}F_{\mathrm{AMM}}(\ell) = [c(\ell)]_+.$$

Thus inequality (21) gives $F_{U^*,\mathrm{buy}}(\ell) \le F_{U^*,\mathrm{sell}}(\ell)$. For $\ell \ge \ell^*$, we instead have $c(\ell) \ge x$, so

$$\begin{aligned}\mathbf{1}_{\{x>0\}}\big(|x| - F_{\mathrm{AMM}}(\ell)\big) &= 0,\\ \mathbf{1}_{\{x<0\}}\big(|x| - F_{\mathrm{AMM}}(\ell)\big) &= [-c(\ell)]_+.\end{aligned}$$

Equal total mass and the cutoff rule for filling trades on each side of ℓ^* therefore give

$$\begin{aligned}F_{U^*,\mathrm{buy}}(\ell) - F_{U^*,\mathrm{sell}}(\ell) = &\sum_{\substack{\mathsf{tx}_j\in G \text{ sells}\\ \mathsf{limit}_j>\ell}} \mathsf{amt}_j - \sum_{\substack{\mathsf{tx}_i\in U \text{ buys}\\ \mathsf{limit}_i>\ell}} \mathsf{amt}_i\\ &- \mathbf{1}_{\{x<0\}}\big(|x| - F_{\mathrm{AMM}}(\ell)\big) \le 0,\end{aligned}$$

where the inequality follows from (22) applied just above ℓ. Thus $F_{U^*,\mathrm{buy}} \le F_{U^*,\mathrm{sell}}$ everywhere.

AMM payment and price monotonicity. If $x \ge 0$, total buyer payments are

$$\int_0^M F^{-1}_{U^*,\mathrm{sell}}(u)\,du = \sum_{\mathsf{tx}_i\in U^* \text{ sells}} \mathsf{fill}_i\mathsf{limit}_i + P(x),$$

since the sell distribution includes the AMM's supply at marginal prices $p(v)$ for $v \in [0, x]$. After paying trader sellers at their limits, the AMM retains $P(x)$. If $x < 0$, total seller receipts are

$$\int_0^M F^{-1}_{U^*,\mathrm{buy}}(u)\,du = \sum_{\mathsf{tx}_i\in U^* \text{ buys}} \mathsf{fill}_i\mathsf{limit}_i - P(x),$$

since the AMM's buy orders have total value $\int_x^0 p(v)\,dv = -P(x)$. Subtracting these receipts from trader buyers' payments again gives a net AMM receipt

of $P(x)$. Finally, both $F^{-1}_{U^*,\text{buy}}$ and $F^{-1}_{U^*,\text{sell}}$ are nondecreasing. Requests on one side execute at their limits; on the other side, execution prices are averages of the corresponding inverse function over consecutive intervals I_i. An earlier interval's average cannot exceed a later interval's average, so execution prices are nondecreasing on each side in the prescribed order of limits and identifiers. □

Security properties. Price validity and pre-approval inclusion are implied by Lemma 9. Pre-approval liveness holds with η-admissibility replaced by Equations (23) and (24). Agreement and accountable inclusion remain unchanged.

LP deposits and withdrawals at block end. A standard extension to our protocol would allow passive LPs to deposit and withdraw liquidity in the AMM. We impose that all deposits and withdrawals are handled at the end of a block, after regular trade requests have been executed. Validators pre-approve all valid LP deposit and withdrawal requests, and such a pre-approval in even a single selected mini-block ensures execution of a valid request. After all trades settle, process every distinct valid LP request. Let $R = (R_A, R_C)$ be the post-trade reserves in the AMM and $\Lambda > 0$ the outstanding LP shares before processing these requests. All conversions use this same snapshot:

- A withdrawal burns its locked w shares and receives $(w/\Lambda)R$.
- A deposit with locked asset caps (d_A, d_C) mints

$$m = \Lambda \min\{d_A/R_A,\ d_C/R_C\},$$

  contributes $(m/\Lambda)R$, and receives its unused escrow back. Any deposit that does not have sufficient inventory in both assets does not execute, and any withdrawal without enough shares does not execute.

6 Related Work

AMM arbitrage and loss-versus-rebalancing. Constant-function market makers (CFMMs) enable decentralized asset exchange through deterministic pricing curves [21, 1, 2]. Because AMM prices change only when trades occur, movements in external market prices create arbitrage opportunities against stale AMM prices, resulting in *loss-versus-rebalancing* (LVR) [28]. Prior work studies the effects of fees and arbitrage on these losses and on liquidity provision [27, 10], while recent empirical work highlights the role of large price movements in CEX–DEX arbitrage [25]. FirstSwap keeps the underlying CFMM pricing rule but changes when traders receive an execution commitment. Traders can obtain a pre-approval while the block is still being formed rather than waiting until its end. Our economic analysis studies how this earlier commitment reduces the incentive to wait for a large CEX–DEX price discrepancy before trading.

MEV mitigation and transaction ordering. A growing literature studies how execution and ordering rules can reduce maximal extractable value (MEV), including front-running, sandwich attacks, and arbitrage [14, 33]. One approach

is to reduce the importance of ordering. Batch execution [9] executes trades together at a uniform end-of-batch price, while SPEEDEX [31] and Groundhog [32] avoid imposing an intra-block order through commutative state updates; these systems still rely on consensus to determine end-of-batch outcomes. A complementary line seeks fairer sequencing: order fairness [22] and enforceable sequencing rules [17]. For AMMs specifically, CLVR [26] reduces intra-block price movements, while Chan, Wu, and Shi [11] design an incentive-compatible mechanism requiring fair ordering. Diamandis and Angeris [15] show, however, that fair ordering need not yield economically efficient allocations. Moreover, Condorcet cycles limit fair ordering. Practical systems such as MEV Blocker [12] and SUAVE [18] instead use order-flow auctions and alternative block-building mechanisms to reduce or redistribute ordering-based MEV. FirstSwap deterministically executes pre-approved trades, without the complex task of fairly ordering transactions by arrival time.

Consensus-free fast paths and pre-confirmations. Several systems bypass full consensus for latency-sensitive operations. Prior work shows that payments and single-owner transactions can be confirmed by validator quorums without consensus [19, 3, 24, 4, 6, 34], while hybrid protocols optimistically avoid consensus and fall back to it upon conflicts [20, 13]. Recent work shows that some shared-object operations can also avoid consensus. Stingray [35] uses local approvals against scaled capacity bounds, CryptoConcurrency [37] avoids consensus when concurrent transfers are jointly feasible, and Zef [5] explores consensus-free decentralized assets and swaps. FirstSwap applies similar ideas to AMMs, where trades compete for liquidity and independently approved trades may not all be feasible. We use limit prices together with a scaled admissibility rule to pre-approve trades while maintaining admissibility. Recent work on pre-confirmations [7] makes commitments before block confirmation for censorship resistance. FirstSwap obtains similar guarantees directly from a quorum of validators and provides a different economic rationale for their usefulness.

Acknowledgments

SS was supported in part by the AFOSR Award FA9550-24-1-0156, and research grants from the Bakar Fund, Peder Sather Funds, the Stellar Development Foundation, Supra Inc., Category Labs, Input Output Global, and Byzantine Research Inc.

References

1. Adams, H.: Uniswap whitepaper (2018), `https://hackmd.io/@HaydenAdams/HJ9jLsfTz`, whitepaper
2. Angeris, G., Chitra, T.: Improved price oracles: Constant function market makers. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT). pp. 80–91. ACM (2020). `https://doi.org/10.1145/3419614.3423251`

3. Auvolat, A., Frey, D., Raynal, M., Taïani, F.: Money transfer made simple: a specification, a generic algorithm, and its proof. Bull. EATCS **132** (2020)
4. Baudet, M., Danezis, G., Sonnino, A.: FastPay: High-performance byzantine fault tolerant settlement. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT). pp. 163–177. ACM (2020). https://doi.org/10.1145/3419614.3423249
5. Baudet, M., Sonnino, A., Król, M.: Low-latency, scalable, DeFi with Zef. arXiv preprint arXiv:2201.05073 (2022). https://doi.org/10.48550/arXiv.2201.05073
6. Blackshear, S., Chursin, A., Danezis, G., Kichidis, A., Kokoris-Kogias, L., Li, X., Logan, M., Menon, A., Nowacki, T., Sonnino, A., Williams, B., Zhang, L.: Sui lutris: A blockchain combining broadcast and consensus. In: Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security. pp. 2606–2620. ACM (2024). https://doi.org/10.1145/3658644.3670286
7. Botha, Q., Kyriacou, D., McMenamin, C., Oshitani, L., Stouka, A.P.: SoK: Preconfirmations. In: 2026 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–19. IEEE (2026). https://doi.org/10.1109/ICBC67748.2026.11575526
8. Budish, E., Cramton, P., Shim, J.: The high-frequency trading arms race: Frequent batch auctions as a market design response. The Quarterly Journal of Economics **130**(4), 1547–1621 (2015). https://doi.org/10.1093/qje/qjv027
9. Canidio, A., Fritsch, R.: Batching trades on automated market makers. In: 5th Conference on Advances in Financial Technologies (AFT 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 282, pp. 24:1–24:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2023). https://doi.org/10.4230/LIPIcs.AFT.2023.24
10. Capponi, A., Jia, R.: Liquidity provision on blockchain-based decentralized exchanges. The Review of Financial Studies **38**(10), 3040–3085 (2025). https://doi.org/10.1093/rfs/hhaf046
11. Chan, T.H.H., Wu, K., Shi, E.: Mechanism design for automated market makers. In: 7th Conference on Advances in Financial Technologies (AFT 2025). Leibniz International Proceedings in Informatics (LIPIcs), vol. 354, pp. 7:1–7:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2025). https://doi.org/10.4230/LIPIcs.AFT.2025.7
12. CoW Protocol, Beaver Build, Agnostic Relay: MEV blocker: Order-flow auction for user protection. Technical specification (2023), https://mevblocker.io
13. Cowling, J.A., Myers, D.S., Liskov, B., Rodrigues, R., Shrira, L.: HQ replication: A hybrid quorum protocol for byzantine fault tolerance. In: OSDI. pp. 177–190. USENIX Association (2006)
14. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 910–927. IEEE (2020). https://doi.org/10.1109/SP40000.2020.00040
15. Diamandis, T., Angeris, G.: A note on the welfare gap in fair ordering (2023), https://arxiv.org/abs/2303.15239
16. Dwork, C., Lynch, N., Stockmeyer, L.: Consensus in the presence of partial synchrony. Journal of the ACM (JACM) **35**(2), 288–323 (1988)
17. Ferreira, M.V.X., Parkes, D.C.: Credible decentralized exchange design via verifiable sequencing rules. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC). pp. 723–736. ACM (2023). https://doi.org/10.

1145/3564246.3585233
18. Flashbots: SUAVE: Single unifying auction for value expression. Blog post and technical specification (2023), https://writings.flashbots.net/the-future-of-mev-is-suave
19. Guerraoui, R., Kuznetsov, P., Monti, M., Pavlovic, M., Seredinschi, D.: The consensus number of a cryptocurrency. In: PODC. pp. 307–316. ACM (2019)
20. Hendricks, J., Sinnamohideen, S., Ganger, G.R., Reiter, M.K.: Zzyzx: Scalable fault tolerance through byzantine locking. In: Proceedings of the 2010 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2010, Chicago, IL, USA, June 28 - July 1 2010. pp. 363–372. IEEE Computer Society (2010). https://doi.org/10.1109/DSN.2010.5544297, https://doi.org/10.1109/DSN.2010.5544297
21. Hertzog, E., Benartzi, G., Benartzi, G.: Bancor protocol: Continuous liquidity and asynchronous price discovery for tokens through their smart contracts (Apr 2017), https://resources.bancor.network/pages/BancorProtocolWhitepaper.pdf, whitepaper, Draft Version 0.77
22. Kelkar, M., Zhang, F., Goldfeder, S., Juels, A.: Order-fairness for byzantine consensus. In: Advances in Cryptology – CRYPTO 2020. Lecture Notes in Computer Science, vol. 12172, pp. 451–480. Springer (2020). https://doi.org/10.1007/978-3-030-56877-1_16, https://doi.org/10.1007/978-3-030-56877-1_16
23. Kemeny, J.G.: Mathematics without numbers. Daedalus **88**(4), 577–591 (1959), https://www.jstor.org/stable/20026529
24. Kuznetsov, P., Pignolet, Y., Ponomarev, P., Tonkikh, A.: Permissionless and asynchronous asset transfer. Distributed Comput. **36**(3), 349–371 (2023)
25. Ladóczki, B., Rásonyi, M., Tapolcai, J.: Where does MEV really come from? revisiting CEX–DEX arbitrage on ethereum (2026), https://arxiv.org/abs/2604.15973
26. McLaughlin, R., Chemaya, N., Liu, D., Malkhi, D.: CLVR ordering of transactions on AMMs (2024), https://arxiv.org/abs/2408.02634
27. Milionis, J., Moallemi, C.C., Roughgarden, T.: Automated market making and arbitrage profits in the presence of fees. In: Financial Cryptography and Data Security. pp. 159–171. Springer (2024). https://doi.org/10.1007/978-3-031-78676-1_9, https://arxiv.org/abs/2305.14604
28. Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L.: Automated market making and loss-versus-rebalancing (2022), https://arxiv.org/abs/2208.06046
29. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research **6**(1), 58–73 (1981). https://doi.org/10.1287/moor.6.1.58
30. Ramseyer, G., Goel, A.: Fair ordering in replicated systems via streaming social choice. In: Mavronicolas, M., Qi, Q., Schoenebeck, G. (eds.) Web and Internet Economics (WINE 2024). Lecture Notes in Computer Science, vol. 15534, pp. 438–456. Springer (2026). https://doi.org/10.1007/978-3-032-08560-3_25, https://doi.org/10.1007/978-3-032-08560-3_25
31. Ramseyer, G., Goel, A., Mazières, D.: Speedex: A scalable, parallelizable, and economically efficient decentralized exchange. In: 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23). pp. 849–875 (2023)
32. Ramseyer, G., Mazières, D.: Groundhog: Linearly-scalable smart contracting via commutative transaction semantics (2024)
33. Roughgarden, T.: Transaction fee mechanism design. In: Proceedings of the 22nd ACM Conference on Economics and Computation (EC). p. 792. ACM (2021). https://doi.org/10.1145/3465456.3467591

Table 1. Glossary of main notation used in Section 2.

Symbol / Term	Definition
The FirstSwap exchange	
$\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$	A trade request; amt: amount of assets the trader wishes to buy; limit: max. price per unit the trader is willing to pay.
$p(x), P(x), x_0$	Price per unit after x units have been bought from the pool; with cumulative payment $P(x) = \int_0^x p(u)\,du$, x_0 and initial pool capacity x_0.
S, admissible	Set of trade requests $S = \{(\mathsf{amt}_i, \mathsf{limit}_i)\}$, admissible if $\forall\, \mathsf{tx}_i \in S$: $p(\sum_{j:\mathsf{limit}_j \le \mathsf{limit}_i} \mathsf{amt}_j) \le \mathsf{limit}_i$.
Economic model	
$\mathsf{m} \in \{\mathrm{FS}, \mathrm{FIFO}, \mathrm{batch}\}$	Trading mechanism: FirstSwap, fair ordering, or batch auction.
$p_{\mathrm{ex},t},\ p_{\mathrm{ex},0}$	CEX price at time t, and its initial value.
$\varepsilon_i,\ F_\varepsilon$	Trader i's fixed private asset valuation, with a mean-zero distribution having CDF F_ε.
$\theta_{i,t},\ r,\ \varepsilon(r)$	Total valuation $\theta_{i,t} = p_{\mathrm{ex},t} + \varepsilon_i$. Rank r is the fraction of higher-valuation traders; $\varepsilon(r) = F_\varepsilon^{-1}(1-r)$.
$q_i(\varepsilon_i),\ x_{\mathrm{LP}}$	Quantity allocated to trader i, inventory remaining in the pool.
$\tau,\ p_{\mathrm{bid},\tau}$	Submission time, submitted limit price.
$\mathrm{Pr}_{\mathrm{fill}}(p_{\mathrm{bid}}), \mathrm{Pr}_{\mathsf{m}}(r)$	Execution probability of a submitted order and of a trader of rank r under mechanism m.
$Vol_{\mathsf{m}}, \mathcal{W}_{\mathsf{m}},\ \mathcal{R}_{\mathsf{m}}$	Total volume: $\int_0^1 \mathrm{Pr}_{\mathsf{m}}(r)\,dr$, social welfare and AMM gains from trade for mechanism m.

34. Spiegelman, A., Sonnino, A., Malkhi, D., Keidar, I.: Linera: A Byzantine-fault tolerant protocol for low-latency total order (2023), https://arxiv.org/abs/2308.09534
35. Sridhar, S., Sonnino, A., Kokoris-Kogias, L.: Stingray: Fast concurrent transactions without consensus (2025), https://arxiv.org/abs/2501.06531
36. Tideman, T.N.: Independence of clones as a criterion for voting rules. Social Choice and Welfare **4**(3), 185–206 (1987). https://doi.org/10.1007/BF00433944, https://doi.org/10.1007/BF00433944
37. Tonkikh, A., Ponomarev, P., Kuznetsov, P., Pignolet, Y.A.: CryptoConcurrency: (almost) consensusless asset transfer with shared accounts. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. pp. 1556–1570. ACM (2023). https://doi.org/10.1145/3576915.3616587

A Brownian FirstSwap equilibrium and outcome formulas

This appendix proves the threshold formula (5) and derives the execution probability (6), welfare (7), and AMM gains (8) stated in Section 2.

1. First-order condition. A trader with private value $\varepsilon(r)$ chooses a cutoff $p^*_{\mathrm{ex}}(s)$, at which a mass s has traded and the AMM price is $p(s)$, to maximize

$$e^{-\gamma_{\mathrm{diff}}[p^*_{\mathrm{ex}}(s) - p_{\mathrm{ex},0}]}[p^*_{\mathrm{ex}}(s) + \varepsilon(r) - p(s)]. \tag{28}$$

Differentiating with respect to s and evaluating at $s = r$ gives

$$\gamma_{\mathrm{diff}}[p^*_{\mathrm{ex}}(r) + \varepsilon(r) - p(r)] = 1 - \frac{p'(r)}{p^{*\prime}_{\mathrm{ex}}(r)}. \tag{29}$$

2. Constant markup. For $p^*_{\text{ex}}(r) = p(r) - \varepsilon(r) + c$ and $p'(r) = -\kappa\varepsilon'(r)$, we have $p^{*\prime}_{\text{ex}}(r) = -(1+\kappa)\varepsilon'(r)$. The FOC becomes

$$\gamma_{\text{diff}} c = 1 - \frac{\kappa}{1+\kappa} \quad \Longrightarrow \quad c = \frac{1}{\gamma_{\text{diff}}(1+\kappa)}, \tag{30}$$

which gives (5). Because the derivative of the trader's objective function in Equation (28) only crosses zero once, this choice of trading time is globally optimal.

Because total asset supply is fixed and passive liquidity is allocated independently of private values, the mechanism-dependent component of expected welfare is

$$\mathcal{W}^{\text{diff}}_{\text{FS}} = \int_0^1 \varepsilon(r)\,\Pr^{\text{diff}}_{\text{FS}}(r)\,dr. \tag{31}$$

For AMM gains from trade, let τ_r be the first time $p_{\text{ex},t}$ reaches $p^*_{\text{ex}}(r)$. Conditional on $\tau_r < T_{\text{blk}}$, exponential memorylessness and the strong Markov property imply

$$\mathbb{E}[p_{\text{ex},T_{\text{blk}}} \mid \tau_r < T_{\text{blk}}] = p^*_{\text{ex}}(r). \tag{32}$$

Thus the expected AMM markout on a filled type-r trade is $p(r) - p^*_{\text{ex}}(r)$, giving

$$\mathcal{R}^{\text{diff}}_{\text{FS}} = \int_0^1 [p(r) - p^*_{\text{ex}}(r)]\,\Pr^{\text{diff}}_{\text{FS}}(r)\,dr. \tag{33}$$

B Allocation representation after a jump

This appendix proves the allocation representation in Lemma 2 of Section 2.

Proof (Proof of Lemma 2). Fix a mechanism $\mathsf{m} \in \{\text{FS}, \text{FIFO}, \text{batch}\}$. Write e for private value, $a(e) := \Pr^{\text{jump}}_{\mathsf{m}}(1 - F_\varepsilon(e))$ for its interim allocation probability, and $t(e)$ for expected payment. The expected utility of a trader given its type e is

$$u(e) = (p_{\text{ex,jump}} + e)a(e) - t(e).$$

Because each trader could "mimic" the choices of any other trader and obtain their expected payment and fill probabilities $a(e'), t(e')$, we have the first-order condition $u'(e) = a(e)$. So long as the distribution F_ε has support on types who do not want to trade, $u(e) \to 0$ as $e \to -\infty$, and hence

$$t(e) - p_{\text{ex,jump}} a(e) = ea(e) - \int_{-\infty}^{e} a(s)\,ds.$$

Integrating over F_ε and applying Fubini gives

$$\mathcal{R}^{\text{jump}}_{\mathsf{m}} = \int \left[e - \frac{1 - F_\varepsilon(e)}{f_\varepsilon(e)}\right] a(e)\,dF_\varepsilon(e).$$

Since $r = 1 - F_\varepsilon(e)$ and $\varepsilon'(r) = -1/f_\varepsilon(\varepsilon(r))$, this is the stated rank formula. The welfare formula follows directly because, conditional on $p_{\text{ex,jump}}$, allocating the asset to type e contributes the mechanism-dependent surplus e. □

C Single crossing and jump-volume comparisons

This appendix proves the decompositions and the FirstSwap–FIFO single-crossing conditions of Lemma 3. The lemma statement appears in Section 2.

Proof of sufficient condition for single crossing in Lemma 3

Proof. We compare the slopes of the two fill probabilities. If their ratio crosses one only upward, the allocation difference first decreases and then may increase, giving the desired single crossing.

For prices z, define

$$w(z) = \frac{1}{p'(p^{-1}(z))}, \qquad \alpha(z) = -\frac{w'(z)}{w(z)} = \frac{p''(p^{-1}(z))}{p'(p^{-1}(z))^2}, \qquad \mathcal{H}(z) = \frac{1}{\alpha(z)}.$$

The AMM assumption gives

$$\mathcal{H}'(z) = 2 - \mathcal{C}(p^{-1}(z)) \in [0,1], \qquad \mathcal{H} \text{ is concave.} \tag{34}$$

On the common interior allocation range, write $\theta = p_{\text{ex,jump}} + e$, $p_{\text{bid}}(e) = p^{\text{jump}}_{\text{bid,FS}}(r(e))$, and $\delta = \theta - p_{\text{bid}}(e) > 0$. Here p_{bid} abbreviates $p_{\text{bid}}(e)$, and $p'_{\text{bid}}(e)$ denotes differentiation with respect to private value, not rank. The queue formulas and the differentiated first-price FOC give

$$\begin{aligned} \frac{d}{de}\Pr^{\text{jump}}_{\text{FIFO}}(r(e)) &= \frac{w(\theta)}{\overline{F}_\varepsilon(e)}, \qquad \frac{d}{de}\Pr^{\text{jump}}_{\text{FS}}(r(e)) = \frac{p'_{\text{bid}}(e)w(p_{\text{bid}})}{\overline{F}_\varepsilon(e)}, \\ p'_{\text{bid}}(e) &= \frac{1+\delta h_\varepsilon(e)}{2+\delta\alpha(p_{\text{bid}})}. \end{aligned} \tag{35}$$

Thus, defining the allocation difference and its slope ratio by

$$\begin{aligned} \Delta(e) &= \Pr^{\text{jump}}_{\text{FS}}(r(e)) - \Pr^{\text{jump}}_{\text{FIFO}}(r(e)), \\ \mathcal{D}(e) &= \frac{\frac{d}{de}\Pr^{\text{jump}}_{\text{FS}}(r(e))}{\frac{d}{de}\Pr^{\text{jump}}_{\text{FIFO}}(r(e))} = \frac{w(p_{\text{bid}})}{w(\theta)}\frac{1+\delta h_\varepsilon(e)}{2+\delta\alpha(p_{\text{bid}})}. \end{aligned}$$

we have $\Delta'(e) = w(\theta)[\mathcal{D}(e)-1]/\overline{F}_\varepsilon(e)$. All derivatives of Δ and $\mathcal{D}$ below are with respect to e.

We show that $\mathcal{D}' > 0$ whenever $\mathcal{D} = 1$. At such a point, set

$$a = \delta\alpha(p_{\text{bid}}) > 0, \qquad k = p'_{\text{bid}}(e) = \frac{w(\theta)}{w(p_{\text{bid}})}, \qquad N = 1+\delta h_\varepsilon(e) = k(2+a), \qquad \beta = \mathcal{H}'(p_{\text{bid}}) \in [0,1].$$

Since $N \geq 1$ and $\mathcal{H}' \leq 1$,

$$\frac{1}{2+a} \leq k = \exp\left\{-\int_{p_{\text{bid}}}^{\theta} \alpha(z)\,dz\right\} \leq \frac{1}{1+a}. \tag{36}$$

Differentiating $\mathcal{D}$ at this point yields

$$\delta N\frac{\mathcal{D}'}{\mathcal{D}} = (1-k)(2k-1) + \delta^2 h'_\varepsilon(e) - \delta^2\alpha'(p_{\text{bid}})k^2 + \delta N[\alpha(\theta) - \alpha(p_{\text{bid}})k]. \tag{37}$$

Here derivatives of α are with respect to price. Increasing hazard gives $h'_\varepsilon \geq 0$, while concavity of $\mathcal{H}$ gives $\alpha(\theta) \geq \alpha(p_{\text{bid}})/(1+\beta a)$ and $-\delta^2\alpha'(p_{\text{bid}}) = \beta a^2$. Therefore

$$\delta N \frac{\mathcal{D}'}{\mathcal{D}} \geq (1-k)(2k-1) + \beta a^2 k^2 + ak(2+a)\left[\frac{1}{1+\beta a} - k\right].$$

The right side decreases in β: its derivative is $a^2 k[k - (2+a)/(1+\beta a)^2] < 0$ by (36). Setting $\beta = 1$ hence gives

$$\delta N \frac{\mathcal{D}'}{\mathcal{D}} \geq L_a(k) := -1 + \left[3 + \frac{a(2+a)}{1+a}\right] k - 2(1+a)k^2.$$

This quadratic is concave, and at the endpoints of the permitted interval,

$$L_a\left(\frac{1}{2+a}\right) = \frac{a}{(2+a)^2(1+a)} > 0, \qquad L_a\left(\frac{1}{1+a}\right) = \frac{a}{(1+a)^2} > 0.$$

Thus $\mathcal{D}' > 0$ whenever $\mathcal{D} = 1$, so the slope ratio can cross one only upward.

□

D Small- and large-jump comparisons

This appendix proves Lemma 4(ii) and (iii) of Section 2. We first prove part (ii), without linking the AMM curve to the private-value quantiles. At $\Delta p_{\text{mid}} = 0$, batch fills exactly the positive-private-value traders. For small positive Δp_{mid}, its welfare loss is quadratic, while FirstSwap continues to execute a positive fraction of central negative types.

Theorem 5 (Batch dominance near efficient clearing). *Suppose $f_\varepsilon > 0$ is C^1 on $\mathbb{R}$, symmetric about zero, and log-concave, and $p \in C^2$ satisfies $p' > 0$ and $p'' \geq 0$ on an inventory domain containing $[0, 1/2]$, with $p(0)$ finite. Fix any equilibrium of the first-price jump game, and assume both mechanisms have finite terminal prices and inventories inside this domain. Set $\Delta p_{\text{mid}} = p_{\text{ex,jump}} - p(1/2)$, $s_p = p'(1/2)$, $\Delta_p = p(1/2) - p(1/4)$, and $s_\varepsilon = 1/(2f_\varepsilon(0))$. If $\Delta_p \geq 2F_\varepsilon^{-1}(3/4)$ and $0 \leq \Delta p_{\text{mid}} < (s_\varepsilon + s_p/4)\sqrt{\Delta_p/(24 s_p)}$, then batch has strictly higher welfare and AMM gains: $\mathcal{W}^{\text{jump}}_{\text{batch}} > \mathcal{W}^{\text{jump}}_{\text{FS}}$ and $\mathcal{R}^{\text{jump}}_{\text{batch}} > \mathcal{R}^{\text{jump}}_{\text{FS}}$.*

Proof. Put $a = F_\varepsilon^{-1}(3/4)$ and let $\mathcal{W}^* = \int_0^\infty e f_\varepsilon(e)\, de$, the welfare from filling exactly the positive-private-value traders. Symmetric log-concavity implies $f_\varepsilon(e) \leq f_\varepsilon(0)$ and a nondecreasing hazard rate $h_\varepsilon = f_\varepsilon/(1-F_\varepsilon)$. In particular,

$$a \geq \frac{1}{4f_\varepsilon(0)} = \frac{s_\varepsilon}{2}, \qquad f_\varepsilon(a) = \tfrac{1}{4}h_\varepsilon(a) \geq \tfrac{1}{4}h_\varepsilon(0) = \frac{f_\varepsilon(0)}{2}.$$

Since $\Delta_p \geq 2a \geq s_\varepsilon$ and $\Delta_p \leq s_p/4$, we have $s_\varepsilon \leq s_p/4$. The assumed upper bound on Δp_{mid} is therefore at most $(s_p/2)\sqrt{1/96} < s_p/4$. We will prove the bound

$$\mathcal{W}^{\text{jump}}_{\text{batch}} - \mathcal{W}^{\text{jump}}_{\text{FS}} \geq s_\varepsilon \left[\frac{\Delta_p}{96 s_p} - \frac{1}{4}\left(\frac{\Delta p_{\text{mid}}}{s_\varepsilon + s_p/4}\right)^2\right] > 0. \tag{38}$$

Batch's loss is quadratic. Write $-z$ for batch's marginal private value. Since $\Delta p_{\text{mid}} \geq 0$, we have $z \geq 0$, and symmetry and batch clearing give

$$x_{\text{batch}}^{\text{jump}} = F_\varepsilon(z), \qquad \Delta p_{\text{mid}} = z + p(F_\varepsilon(z)) - p(1/2) \geq z + s_p[F_\varepsilon(z) - 1/2].$$

For $\Delta p_{\text{mid}} \leq s_p/4$, this implies $F_\varepsilon(z) \leq 3/4$ and hence $z \leq a$. The density is nonincreasing on $[0, a]$, so $F_\varepsilon(z) - 1/2 \geq z f_\varepsilon(a) \geq z f_\varepsilon(0)/2$. Therefore

$$z \leq \frac{\Delta p_{\text{mid}}}{1 + s_p f_\varepsilon(0)/2},$$

and

$$\mathcal{W}^* - \mathcal{W}_{\text{batch}}^{\text{jump}} = \int_0^z e f_\varepsilon(e)\, de \leq \frac{f_\varepsilon(0) z^2}{2} \leq \frac{s_\varepsilon}{4} \left(\frac{\Delta p_{\text{mid}}}{s_\varepsilon + s_p/4} \right)^2. \tag{39}$$

FirstSwap has a uniform selection loss. A bid of $p(1/4)$ fills with probability at least $1/4$: the eligible population has mass at most one, so cumulative sales satisfy $x'(t) \leq 1$ and $x(t) \leq t$, and an order arriving before queue time $1/4$ finds the quote at or below $p(1/4)$. For every private value $e \in [-a, 0]$ this deviation gives

$$U_{\text{FS}}^{\text{jump}}(e) \geq \frac{\Delta p_{\text{mid}} + e + \Delta_p}{4} \geq \frac{\Delta_p - a}{4} \geq \frac{\Delta_p}{8} > 0. \tag{40}$$

Let $p_{\text{bid}}(e)$ denote the equilibrium bid of type e and $\delta(e) := p_{\text{ex,jump}} + e - p_{\text{bid}}(e)$ its surplus per fill, so that $U_{\text{FS}}^{\text{jump}}(e) = \Pr_{\text{FS}}^{\text{jump}}(1 - F_\varepsilon(e))\, \delta(e)$. Since $U_{\text{FS}}^{\text{jump}}(e) > 0$, the equilibrium bid fills with positive probability, which requires $p_{\text{bid}}(e) \geq p(0)$ because the AMM quote never falls below $p(0)$. Hence

$$\delta(e) \leq p_{\text{ex,jump}} + e - p(0) \leq \Delta p_{\text{mid}} + \big[p(1/2) - p(0)\big] \leq \frac{s_p}{4} + \frac{s_p}{2} = \frac{3 s_p}{4},$$

using $e \leq 0$, the maintained restriction $\Delta p_{\text{mid}} \leq s_p/4$, and convexity ($p' \leq s_p$ on $[0, 1/2]$). Therefore

$$\Pr_{\text{FS}}^{\text{jump}}(1 - F_\varepsilon(e)) = \frac{U_{\text{FS}}^{\text{jump}}(e)}{\delta(e)} \geq \frac{\Delta_p/8}{3 s_p/4} = \frac{\Delta_p}{6 s_p}, \qquad -a \leq e \leq 0.$$

The argument uses only that type e's equilibrium payoff weakly exceeds the payoff from the deviation; no monotonicity or differentiability of the equilibrium is needed.

The central negative types have a uniformly positive private-value loss. By symmetry and $f_\varepsilon \leq f_\varepsilon(0)$,

$$\int_{-a}^0 (-e) f_\varepsilon(e)\, de = \int_0^a \left[\frac{1}{4} - (F_\varepsilon(u) - 1/2) \right] du \geq \int_0^{s_\varepsilon/2} \left[\frac{1}{4} - f_\varepsilon(0) u \right] du = \frac{s_\varepsilon}{16}.$$

Consequently,

$$\mathcal{W}^* - \mathcal{W}_{\text{FS}}^{\text{jump}} = \int_0^\infty e[1 - \Pr_{\text{FS}}^{\text{jump}}(1 - F_\varepsilon(e))] f_\varepsilon(e)\, de + \int_{-\infty}^0 (-e) \Pr_{\text{FS}}^{\text{jump}}(1 - F_\varepsilon(e)) f_\varepsilon(e)\, de \geq \frac{\Delta_p}{6 s_p} \cdot \frac{s_\varepsilon}{16} = \frac{s_\varepsilon\, \Delta_p}{96\, s_p}.$$

Subtracting (39) proves (38).

AMM gains are also higher. Let $\mathcal{U}_{\mathsf{m}}^{\mathrm{jump}} = \mathcal{W}_{\mathsf{m}}^{\mathrm{jump}} - \mathcal{R}_{\mathsf{m}}^{\mathrm{jump}}$ denote aggregate trader utility. If $x_{\mathrm{FS}}^{\mathrm{jump}} \le x_{\mathrm{batch}}^{\mathrm{jump}}$, the batch price $p_{\mathrm{batch}}^{\mathrm{jump}} = p(x_{\mathrm{batch}}^{\mathrm{jump}})$ is at least FirstSwap's terminal quote. An infinitesimal trader can therefore obtain its batch payoff under FirstSwap by bidding $p_{\mathrm{batch}}^{\mathrm{jump}}$, which guarantees execution, if its value is at least $p_{\mathrm{batch}}^{\mathrm{jump}}$, and by abstaining otherwise. Optimality implies $\mathcal{U}_{\mathrm{FS}}^{\mathrm{jump}} \ge \mathcal{U}_{\mathrm{batch}}^{\mathrm{jump}}$.

If $x_{\mathrm{FS}}^{\mathrm{jump}} \ge x_{\mathrm{batch}}^{\mathrm{jump}}$, Lemma 2 gives $\mathcal{U}_{\mathsf{m}}^{\mathrm{jump}} = \int_0^1 \omega(r)\,\mathrm{Pr}_{\mathsf{m}}^{\mathrm{jump}}(r)\,dr$, where $\omega(r) = -r\varepsilon'(r) = 1/h_\varepsilon(\varepsilon(r))$ is positive and nondecreasing because the hazard rate h_ε is nondecreasing. At a fixed volume this integral is minimized by allocating to the smallest ranks. Hence

$$\mathcal{U}_{\mathrm{FS}}^{\mathrm{jump}} \ge \int_0^{x_{\mathrm{FS}}^{\mathrm{jump}}} \omega(r)\,dr \ge \int_0^{x_{\mathrm{batch}}^{\mathrm{jump}}} \omega(r)\,dr = \mathcal{U}_{\mathrm{batch}}^{\mathrm{jump}}.$$

In either case,

$$\mathcal{R}_{\mathrm{batch}}^{\mathrm{jump}} - \mathcal{R}_{\mathrm{FS}}^{\mathrm{jump}} = \mathcal{W}_{\mathrm{batch}}^{\mathrm{jump}} - \mathcal{W}_{\mathrm{FS}}^{\mathrm{jump}} + \mathcal{U}_{\mathrm{FS}}^{\mathrm{jump}} - \mathcal{U}_{\mathrm{batch}}^{\mathrm{jump}} \ge \mathcal{W}_{\mathrm{batch}}^{\mathrm{jump}} - \mathcal{W}_{\mathrm{FS}}^{\mathrm{jump}} > 0.$$

□

The condition $\Delta_p \ge 2a$ says that the AMM's price rise from inventory $1/4$ to $1/2$ exceeds the interquartile range of private values. It guarantees that central negative types can profitably bid for early liquidity. The jump bound keeps batch's marginal private value close to zero, limiting its welfare loss to a quadratic in Δp_{mid}. The result applies to general convex AMMs, including constant-product curves with sufficient liquidity, and uses no bound on the AMM slope near inventory exhaustion. It also places no restriction on the FirstSwap equilibrium beyond existence: the welfare proof uses only that a filled order pays at least $p(0)$ and that no type can gain by switching to the bid $p(1/4)$; the AMM-gain comparison uses optimal bidding and the envelope identity of Lemma 2.

Proof (Proof of Lemma 4(iii)). Write $\overline{F}_\varepsilon = 1 - F_\varepsilon$ and let $c_{\mathrm{FS}}^{\mathrm{jump}}$ be FirstSwap's terminal eligibility cutoff. The limiting queue FOC gives

$$c_{\mathrm{FS}}^{\mathrm{jump}} = p(x_{\mathrm{FS}}^{\mathrm{jump}}) - p_{\mathrm{ex,jump}} + p'(x_{\mathrm{FS}}^{\mathrm{jump}})\overline{F}_\varepsilon(c_{\mathrm{FS}}^{\mathrm{jump}}).$$

If $c_{\mathrm{FS}}^{\mathrm{jump}} > 0$, symmetry gives $\overline{F}_\varepsilon(c_{\mathrm{FS}}^{\mathrm{jump}}) < 1/2$. Since $x_{\mathrm{FS}}^{\mathrm{jump}} \le 1$, convexity and the jump condition imply

$$c_{\mathrm{FS}}^{\mathrm{jump}} \le p(1) - p_{\mathrm{ex,jump}} + \frac{1}{2}p'(1) \le 0,$$

a contradiction. Thus FirstSwap fills every nonnegative-private-value trader with certainty. FIFO does too, because its terminal quote is $p(x_{\mathrm{FIFO}}^{\mathrm{jump}}) \le p(1) < p_{\mathrm{ex,jump}}$. The single crossing can therefore be chosen with $e^\star \le 0$.

FirstSwap bids do not exceed values, so its eligible mass at any common inventory is at most FIFO's, giving $x_{\mathrm{FS}}^{\mathrm{jump}} \le x_{\mathrm{FIFO}}^{\mathrm{jump}}$ by queue comparison. Both terms in the welfare decomposition of Lemma 3 are consequently nonnegative. For AMM gains, Lemma 2 replaces private value e by the virtual value $e - \overline{F}_\varepsilon(e)/f_\varepsilon(e)$. The increasing hazard rate makes this function strictly increasing, and its value at $e^\star \le 0$ is negative. The same centered selection–volume decomposition therefore gives nonnegative terms for AMM gains. Both rankings are strict if allocations differ on a set of positive measure. □

E Validator-count equivalence after a simultaneous jump

E.1 Fair ordering: anonymity, neutrality, and a uniform queue

Fix a simultaneous cohort of m labeled orders. Conditional on all submitted bids, quantities, and private values, suppose the n validators' receipt rankings are independent uniform permutations:

$$\pi_1, \dots, \pi_n \overset{\text{iid}}{\sim} \mathrm{Unif}(\mathfrak{S}_m).$$

Independent continuous, identically distributed receipt delays satisfy this assumption. We consider an aggregation rule $\mathcal{A}_n$ that returns a complete ordering and satisfies the following axioms.

1. *Transaction neutrality.* For every relabeling $\rho \in \mathfrak{S}_m$,

$$\mathcal{A}_n(\rho\pi_1, \dots, \rho\pi_n) \overset{d}{=} \rho\, \mathcal{A}_n(\pi_1, \dots, \pi_n).$$

   Here $\rho\pi$ relabels the entries of the ordered list π.
2. *Validator anonymity.* Permuting the n input rankings does not change the output distribution.
3. *Bid blindness and neutral resolution.* The aggregator uses only the receipt rankings and fresh independent randomness. Bids, values, quantities, and privileged transaction labels cannot affect priority; any ties or partial orders are resolved into a complete order in a transaction-neutral way.

Theorem 6 (Ideal fair ordering is equivalent to one random queue). *Under these assumptions, for every $n \ge 1$ and every conditional bid and value profile,*

$$\Pr\{\mathcal{A}_n(\pi_1, \dots, \pi_n) = \sigma\} = \frac{1}{m!}, \qquad \sigma \in \mathfrak{S}_m.$$

Consequently, applying the same serial AMM execution and payment rule to the aggregated queue gives exactly the same joint allocation and payment law as one validator with a uniform receipt queue.

Proof. Condition on a fixed bid and value profile. The joint receipt-ranking law is unchanged by a common relabeling ρ. Transaction neutrality therefore implies that the output law is also unchanged by ρ. For any two full rankings σ and

σ', there is a relabeling ρ with $\rho\sigma = \sigma'$. Thus every ranking has the same probability. Their probabilities sum to one, proving the displayed formula. Bid blindness ensures that this argument remains valid after conditioning on the traders' chosen bids. □

The stated axioms are satisfied, for example, by the Kemeny [23] and Ranked Pairs methods [36, 30] of aggregating preference orderings.

E.2 FirstSwap with bid-ordered AMM execution

We now specify conditions under which any monotone equilibrium of the FS protocol yields the same allocation as in the version with a single centralized validator. Let

$$\theta(r) = p_{\text{ex,jump}} + \varepsilon(r), \qquad D := 1 - F_\varepsilon(p(0) - p_{\text{ex,jump}}) \in (0, 1).$$

Proof. Entry, execution prices, and the incentive equation. Every execution price is at least $p(0)$, so values below $p(0)$ cannot profit from entry. Conversely, any value above $p(0)$ can bid strictly between $p(0)$ and its value. Arriving sufficiently early at all validators and receiving a sufficiently early tie-break among the highest signature scores has positive probability and gives a feasible fill at a price no larger than this bid. Thus a regular equilibrium cannot exclude an interval of values above $p(0)$. Participation is precisely $r < D$, up to the indifferent boundary type.

Because bids strictly decrease with r, lower-value active types execute first. Hence

$$X_n(r) = \int_r^D a_n(t)\, dt, \qquad p_n(r) = p(X_n(r)).$$

A trader of true rank r who submits rank t's bid receives utility

$$a_n(t)\,[\theta(r) - p(X_n(t))].$$

The first-order condition at $t = r$, together with $X_n' = -a_n$, gives the following system, suppressing the validator subscript:

$$\begin{aligned} X'(r) &= -a(r), \\ a'(r) &= -\frac{a(r)^2 p'(X(r))}{\theta(r) - p(X(r))}, \qquad a(0+) = 1, \quad X(D-) = 0. \end{aligned} \tag{41}$$

The boundary $X(D-) = 0$ follows from the definition of execution depth. A bid covering $p(2D)$ receives every signature and can be added to any feasible allocation, so it guarantees execution at a price at most $p(D)$. Comparing this deviation with an equilibrium payoff at most $a_n(r)[\theta(r) - p(0)]$, and using $\theta(r) \to \infty$ as $r \downarrow 0$, proves $a_n(0+) = 1$. We show that there is at most one solution to this system of differential equations with these boundary conditions.

Table 2. Glossary of main notation used in Section 3.

Symbol / Term	Definition
Scaled Admissibility	
$n; f; q$	Number of validators; Byzantine threshold; quorum size
η	Pre-approval fraction; $\eta := (q-f)/(n-f)$. Each validator pre-approves at most an η-fraction of available capacity at each price level.
η-admissible	admissible when every $\mathsf{tx}_i \in S$ has amt_i scaled by a factor $1/\eta$.
Pre-approvals and certificates	
$\sigma_v(\mathsf{tx})$	Pre-approval of tx by validator v: a signed message $\langle \mathsf{APPROVE}, \mathsf{tx} \rangle$.
$\pi(\mathsf{tx})$	Approval certificate for tx with pre-approvals from q distinct validators.
A_v	Local set of trade requests pre-approved by validator v in the current epoch; η-admissible by construction.
Π_v	Local set of approval certificates learned by validator v in the current epoch.
Protocol structure	
epoch	A fixed-duration period comprising a pre-approval phase followed by a consensus round to decide a final block.
B_v	Mini-block formed by validator v at epoch end: $B_v = (\mathsf{A}_v, \Pi_v)$.
$\mathcal{B}$	Final block output by MVBA: a set of at least $n-f$ well-formed mini-blocks.
Valid	External validity predicate for MVBA; holds when $\mathcal{B}$ contains well-formed mini-blocks from $M_B \geq n-f$ validators with consistent cross-signatures.
$J(w)$	Accountability function; returns the set of validators implicated by evidence w; never implicates an honest validator.

To prove uniqueness, set $Z = -a$. The system for (X, Z) is a so-called cooperative system with $\frac{\partial X'}{\partial Z}$ and $\frac{\partial Z'}{\partial X}$ both positive:

$$X' = Z, \qquad Z' = \frac{Z^2 p'(X)}{\theta - p(X)}, \qquad \frac{\partial}{\partial X}\frac{p'(X)}{\theta - p(X)} = \frac{p''(X)[\theta - p(X)] + p'(X)^2}{[\theta - p(X)]^2} > 0.$$

By the comparison theorem for cooperative systems, a larger initial condition $X(0+)$ and the same $Z(0) = -1$ give larger X and weakly larger Z pointwise. The difference between the two X paths is nondecreasing, so stays at least their strictly positive initial difference. They cannot both end at $X(D-) = 0$. There is therefore at most one solution satisfying $a(0+) = 1$ and $X(D-) = 0$.

Thus every regular separating equilibrium, whenever it exists, has

$$a_n(r) = a(r), \qquad X_n(r) = X(r), \qquad p_n(r) = p(X(r)), \quad 0 < r < D, \qquad (42)$$

where $p_n(r)$ is the expected price conditional on filling. □

F Supplemental Protocol Information: Glossary, Algorithms and Proofs

This appendix gives pseudocode for the decentralized protocol of Section 3 and overflow security proofs.

Algorithm 1 is run by traders, Algorithm 2 by validators during the pre-approval stage, and Algorithm 3 by validators at the end of the epoch. All parameters, the pricing function, and the notions of admissibility and η-admissibility are as defined in the main text. The pseudocode describes a single epoch e. All local state is reset at the start of each epoch, and approvals are signed together with the epoch number, so an approval cannot be reused in a later epoch.

Algorithm 1 Trader-side pre-approval protocol for trade request tx in epoch e

Require: Trade request $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$
Send tx to all validators
$\mathsf{Responses} \leftarrow \emptyset$
while the pre-approval stage of epoch e has not ended **do**
  Wait for a signed validator response, retransmitting tx to validators that have not responded
  if the response is a valid signature on $\langle \mathsf{APPROVE}, \mathsf{tx}, e \rangle$ from a validator not yet in $\mathsf{Responses}$ **then**
    Add the response to $\mathsf{Responses}$
    if $|\mathsf{Responses}| = q$ **then**
      Form $\pi(\mathsf{tx})$ from $\mathsf{Responses}$ and send it to all validators
      return $\pi(\mathsf{tx})$
    end if
  end if
end while
return $\perp$

Proof of Agreement

Theorem 7 (f-Agreement). *The protocol satisfies f-Agreement (defined in Section 3).*

Proof. By MVBA agreement, all honest validators decide the same output $\mathcal{B}$. They apply the same inclusion and execution rules, including the deterministic tie-breaking rule, and therefore output the same final block. □

Proof of Accountable Inclusion.

Theorem 8 (Accountable inclusion). *The protocol satisfies accountable inclusion.*

Proof. Let the evidence w consist of an AC for tx, the decided output $\mathcal{B}$, V^* the creators of its mini-blocks, and $M_B \leftarrow |V^*|$. Define $J(w)$ to contain every validator v such that $v \in V^*$, the AC contains v's approval of tx, and A_v omits that signed approval.

Suppose tx obtains an AC but is absent from the final block. The AC cannot appear in any selected Π_v, because that would include tx directly. Moreover, fewer than ηM_B selected approved sets contain tx, because otherwise the pre-approved inclusion rule would include it. The AC signer set and V^* intersect in at least $q + M_B - n \geq \eta M_B$ validators, by the bound in the proof of Lemma 5. Thus at least one signer in this intersection omitted its approval, so $J(w)$ is nonempty. An honest validator retains every approval it issues in A_v. Therefore no valid evidence causes J to identify an honest validator. □

Algorithm 2 Validator-side processing at validator v in epoch e

Upon receiving trade request $\mathsf{tx} = (\mathsf{amt}, \mathsf{limit})$:
if v has previously responded to tx **then**
  return previously signed response
end if
if the pre-approval stage of epoch e has ended **then**
  Ignore tx
end if

{Check η-admissibility of $\mathsf{A}_v \cup \{\mathsf{tx}\}$; only limits $\geq \mathsf{limit}$ are affected}
for all $\ell \in \{\mathsf{limit}\} \cup \{\mathsf{limit}_i : \mathsf{tx}_i \in \mathsf{A}_v,\ \mathsf{limit}_i > \mathsf{limit}\}$ **do**
  if $\ell < p((\sum_{\mathsf{tx}_j \in \mathsf{A}_v\,:\,\mathsf{limit}_j \leq \ell} \mathsf{amt}_j/\eta) + \mathsf{amt}/\eta)$ **then**
    return signed $\langle \mathsf{REJECT}, \mathsf{tx}, e\rangle$ {stored and replayed}
  end if
end for
$\sigma_v(\mathsf{tx}) \leftarrow$ signed $\langle \mathsf{APPROVE}, \mathsf{tx}, e\rangle$
$\mathsf{A}_v \leftarrow \mathsf{A}_v \cup \{(\mathsf{tx}, \sigma_v(\mathsf{tx}))\}$
return $\sigma_v(\mathsf{tx})$

Upon receiving $\pi(\mathsf{tx})$ before the end of the pre-approval stage:
if $\pi(\mathsf{tx})$ holds valid epoch-e approvals of tx from q distinct validators **then**
  $\Pi_v \leftarrow \Pi_v \cup \{\pi(\mathsf{tx})\}$
end if

Strengthened Inclusion Guarantees

Lemma 10 (f-Certified inclusion). *Let $n \geq 4f + 1$, $q = n - f$, and $\eta = (q - 2f)/(n - f)$. The protocol satisfies the following strengthened pre-approval inclusion property: if at most f validators are Byzantine and a trade request tx obtains valid pre-approvals for an epoch from q distinct validators, not necessarily all honest, then tx is included in that epoch's final block.*

Proof. At least $q - f$ approvers are honest. Since only $n - M_B$ validators lie outside V^*, at least $q - f - (n - M_B) = M_B - 2f$ honest approvers have selected mini-blocks, each containing its approval of tx. Moreover,

$$M_B - 2f - \eta M_B = \frac{2f\big(M_B - (n - f)\big)}{n - f} \geq 0.$$

Thus tx appears in at least ηM_B selected approved sets and is included by the pre-approved inclusion rule. □

Lemma 11 (Strengthened accountable inclusion). *Let $n \geq 4f + 1$, $q = n - f$, and $\eta = (q - 2f)/(n - f)$. There exists a publicly verifiable function J such that, if a trade request tx obtains an approval certificate but is omitted from a valid final block for that epoch, there is evidence w with $|J(w)| \geq f + 1$ (and hence at least f). Moreover, $v \notin J(w)$ for every honest validator v and every w.*

Algorithm 3 End-of-epoch block consensus at validator v

External validity predicate (Definition 4). $\mathsf{Valid}(\mathcal{B}) = 1$ if and only if:
(i) $\mathcal{B}$ contains well-formed mini-blocks from at least $n-f$ distinct validators, where $B_u = (\mathsf{A}_u, \Pi_u)$ is well-formed if A_u is η-admissible with every approval correctly signed by u for epoch e, and every $\pi(\mathsf{tx}) \in \Pi_u$ is correctly formed; and
(ii) for every $\pi(\mathsf{tx})$ in $\mathcal{B}$ and every B_u in $\mathcal{B}$, if u's signature appears in $\pi(\mathsf{tx})$, then the same signed approval of tx appears in A_u.

Protocol.
At the end of the epoch, form the mini-block $B_v \leftarrow (\mathsf{A}_v, \Pi_v)$ and sign it
Invoke MVBA with input B_v and external-validity predicate Valid
Let $\mathcal{B}$ be the decision by MVBA, V^* the creators of its mini-blocks, and $M_B \leftarrow |V^*|$

Pre-approved inclusion: include every tx with approval signatures from at least ηM_B distinct validators in V^*
AC inclusion: include every tx for which a valid $\pi(\mathsf{tx})$ appears in $\mathcal{B}$, unless already included
Execution ordering: order all included trade requests in non-decreasing order of limit, breaking ties deterministically by trade request identifier
Output the resulting ordered block for execution {optionally supplemented with non-pre-approved trades as in Section 4}

Proof. Use the certificate and authenticated decided mini-blocks as evidence w. Let $J(w)$ contain each selected certificate signer whose mini-block omits its signed approval, and set $J(w) = \emptyset$ for invalid evidence. At least $q + M_B - n = M_B - f$ certificate signers are selected. If tx is omitted, fewer than ηM_B selected approved sets contain it. Therefore

$$|J(w)| > M_B - f - \eta M_B = \frac{2fM_B}{n-f} - f \geq f.$$

Since $|J(w)|$ is an integer, at least $f + 1$ validators are identified. An honest validator includes every approval it issued for the epoch in its mini-block, so valid evidence never identifies an honest validator. □