

Dynamic Posted-Price Mechanisms for the Blockchain Transaction Fee Market*

Matheus V. X. Ferreira^{†1}, Daniel J. Moroz², David C. Parkes², and Mitchell Stern^{‡3}

¹Computer Science, Princeton University

²Computer Science, Harvard University

³EECS, University of California, Berkeley

November 17, 2021

Abstract

In recent years, prominent blockchain systems such as Bitcoin and Ethereum have experienced explosive growth in transaction volume, leading to frequent surges in demand for limited block space and causing transaction fees to fluctuate by orders of magnitude. Existing systems sell space using first-price auctions [29]; however, users find it difficult to estimate how much they need to bid in order to get their transactions accepted onto the chain. If they bid too low, their transactions can have long confirmation times. If they bid too high, they pay larger fees than necessary.

In light of these issues, new transaction fee mechanisms have been proposed, most notably EIP-1559 [4], aiming to provide better usability. EIP-1559 is a history-dependent mechanism that relies on block utilization to adjust a base fee. We propose an alternative design—a *dynamic posted-price mechanism*—which uses not only block utilization but also observable bids from past blocks to compute a posted price for subsequent blocks. We show its potential to reduce price volatility by providing examples for which the prices of EIP-1559 are unstable while the prices of the proposed mechanism are stable. More generally, whenever the demand for the blockchain stabilizes, we ask if our mechanism is able to converge to a stable state. Our main result provides sufficient conditions in a probabilistic setting for which the proposed mechanism is approximately welfare optimal and the prices are stable. Our main technical contribution towards establishing stability is an iterative algorithm that, given oracle access to a Lipschitz continuous and strictly concave function f , converges to a fixed point of f .

*The authors thank Cemil Dibek, Vikram V. Ramaswamy, and Tim Roughgarden for helpful discussions.

[†]Contact: mvxf@cs.princeton.edu

[‡]Part of this work was completed while visiting Harvard University.

Contents

1	Introduction	3
1.1	Brief technical overview	5
1.2	Related work	7
2	Model	7
2.1	Preliminaries	8
2.2	Decentralized multi-round auction game	9
2.3	Mechanism design	10
2.4	Illustrative mechanisms	12
2.5	Discussion	14
3	Dynamic Posted-Price Mechanisms	14
4	Asymptotically Stable Update Rules via Fixed Point Theory	16
5	Welfare-Based Dynamic Posted-Price Mechanism	17
5.1	Asymptotic stability of the WDPP mechanism	18
6	Utilization-Based Dynamic Posted-Price Mechanism	19
7	The Truncated Welfare-Based Dynamic Posted-Price Mechanism	21
7.1	Stability	23
7.2	Welfare guarantee at equilibrium	25
8	Experimental Results	25
8.1	Excess demand	26
9	Conclusion	26
9.1	Future work	28
A	Other Experimental Results	32
A.1	Under demand	32
A.2	Demand shocks	33
A.3	Instability of EIP-1559 utilization-based update rule	34
B	Existing Proposals of Blockchain Transaction Fee Mechanisms	35
B.1	EIP-1559	35
B.2	Monopolistic price auction	36
B.3	Random sampling optimal price	37
B.4	Generalized second-price auction	37
C	Omitted Proofs from Section 5	37

D Omitted Proofs from Section 7	39
D.1 Omitted proofs from Section 7.1	40
D.2 Omitted proofs from Section 7.2	45

1 Introduction

Due to the explosive interest in blockchains such as Ethereum and Bitcoin, allocating scarce block space has become a significant problem. The high competition for block space has made transaction fee auctions a multibillion-dollar venture for miners [18]. The status quo, first adopted by Bitcoin [29], is to auction block space via a first-price auction where miners select transactions with the highest fees to be included in a block.

Since the demand for block space varies with time, users need to continually change their bids to maximize their chances of getting a transaction included and, at the same time, minimize the amount they pay. This challenge has resulted in the proposal of new mechanisms [25, 2, 4] that pay special attention to the benefits of mechanisms that are *incentive compatible* (or strategy-proof, promoting truthful bids). A particular challenge for designing transaction fee auctions is that the miners are also auctioneers. While the blockchain can commit to implementing a payment rule, miners have complete control over the transactions included in a block (the *winner determination*). Moreover, there is information asymmetry, where the blockchain only sees the transactions that are included in blocks and not, for example, the excess demand.

Designing strategy-proof auctions that are at the same time robust to strategic auctioneers is a challenging task and closely related to the problem of designing credible strategy-proof auctions [1, 15]. In a *credible auction*, the auctioneer is allowed to implement any “safe deviation”—a deviation that cannot be detected by bidders—from the promised auction rule. Yet, in expectation, the auctioneer should achieve no more revenue than from implementing the promised auction. Akbarpour and Li [1] show that a family of ascending price auctions are the unique revenue-optimal, credible, and strategy-proof auctions for a computationally unbounded auctioneer. Ferreira and Weinberg [15] propose the *Deferred Revelation Auction* (DRA), which is credible for an auctioneer that is computationally bounded.¹ However, it is not known how to implement the ascending-price auction or DRA on blockchains efficiently. Ideally, only information about winning transactions would be stored on the blockchain, whereas a naive implementation of the ascending-price auction or DRA would require information on all transactions. The challenge is that blockchain space is a limited and expensive resource. Each communication on the blockchain is stored in blocks that propagate to the whole network.

Another possible direction would be that of static posted-price mechanisms [7], which are strategy-proof and credible. However, to extract good revenue and welfare, the posted

¹Ferreira and Weinberg [15] also showed that DRA is credible only under standard auction theoretical assumptions on the distribution of bidder valuations. Essaidi et al. [12] proposes the Ascending Deferred Revelation Auction (ADRA) which improves DRA by removing the distributional assumption. As a trade-off, ADRA requires a constant number of rounds (on expectation) to terminate rather than two rounds and bidders must deposit dynamic collateral—the collateral increases the longer a bidder participates in the mechanism—while DRA requires constant collateral.

price must be calibrated according to the distribution of user valuations. This is a challenge when user demand is highly volatile, making static posted-price mechanisms ill-suited to transaction fee markets.

In this paper, we study the design of *dynamic posted-price mechanisms*, where the price offered can adapt based only on information about winning bids from previous auctions. It is essential that the price adaptation only uses the value of winning bids and not all bids due to information asymmetry and the need for information efficiency. We seek to understand whether this ability to use adaptive prices enables approximately welfare-optimal credible mechanisms that provide a good user experience (incentive alignment).

Towards that direction, the recent *EIP-1559 proposal* of Buterin et al. [4] changed the status quo, first-price auction on Ethereum to a hybrid between a first-price auction and a dynamic posted-price mechanism. In EIP-1559, each block has a fixed maximum capacity m and a posted price $q > 0$ that is a deterministic function of the previous block’s utilization (described formally in Appendix B). Every submitted transaction includes a *bid*, $b > 0$, as well as a *bid cap*, $c > 0$, which is the amount a user is willing to spend to have their transaction included. If a transaction’s bid cap c is greater than $b + q$, a miner may choose to include it in the block. A target utilization is set to $m/2$. The protocol increases the price in the next auction if block utilization is above the target and decreases the price if utilization is below the target. A winner pays $b + q$ to the blockchain, while miners receive only b as revenue (q is said to be “burned” since no one receives it). Introducing a price floor via dynamic posted prices is likely to make payments more predictable, and this aspect of the design received good reception from users in a community outreach report [34]. Moreover, awarding only b to miners makes EIP-1559 robust to collusion between miners and bidders [36].

When the posted price q is above the market clearing price, EIP-1559’s mechanism reduces to a dynamic posted-price mechanism similar to the mechanism we study: users will bid $b = 0$ and report a bid cap c that equals their valuation for getting a transaction included in the next block.² A challenge, though, is that EIP-1559 still requires bidders to compete in a first-price auction when the posted price q is below the market clearing price—the price where the demand meets the supply for block space.

Roughgarden [36] formalizes EIP-1559 as a decentralized auction game and shows that EIP-1559 is: incentive compatible for myopic bidders whenever q is above the market clearing price; incentive compatible for myopic miners; and resistant to collusion between miners with bidders.³ Roughgarden [36] also proposes a *tipless mechanism* that modifies EIP-1559 so that the bid b is hard-coded into the protocol. This change sacrifices collusion resistance in favor of the mechanism becoming incentive compatible for myopic bidders even when q is below the market clearing price. Our work is complementary. We focus on designing dynamic posted-price mechanisms that provably converge to a price equilibrium. We propose alternative price update rules that improve upon EIP-1559 and converge to an equilibrium in instances where EIP-1559 is unstable. Our main observation is to design price update rules that use observable bids from prior blocks rather than rely simply on block utilization. It is natural to ask if our mechanisms converge to a *good* price that approximates the market-clearing

²Here, we assume there is no reserve price. EIP-1559 also supports a hard-coded reserve price in its first-price auction.

³Myopic bidders and miners maximize their immediate utility from the next block.

price. If the equilibrium price is too high, few bidders can afford to pay for block access, resulting in low welfare. Thus we also show that our mechanisms are approximately welfare optimal.

The main difference between the mechanisms studied here and the EIP-1559 proposal is how they handle the case where the posted price q is below the market-clearing price (i.e., when the number of bidders willing to pay q exceeds the supply for block space). While EIP-1559 forces users to compete in a first-price auction, we retain the simple dynamic posted-price incentives. Crucially, we need to ensure that even in the case of excess demand, the mechanism remains robust to manipulation by blockchain auctions, i.e., miners will choose to randomize the allocation to users who bid above the posted price.

Our analysis assumes that a miner is myopic about the decisions made regarding which transactions to include in a block whenever they are chosen to create a new block. This is a reasonable assumption for proof-of-work systems in which a new miner is, in effect, chosen at random in each time step to create a block. This creates uncertainty as to when a miner will next be “chosen” and makes it reasonable that miners are myopic and seek to maximize immediate utility when given the opportunity.⁴

1.1 Brief technical overview

We formalize the *dynamic posted-price* (DPP) mechanism (Definition 3.2) as a posted-price mechanism with a *price update rule* and an *initial price* q_0 . In each time step t , the mechanism announces a *price* q_t and the miner whose turn it is can allocate block space to at most m users from those with bids $b_i \geq q_t$ (henceforth, we refer to users as *bidders*). Given the winning bids at time step t , the update rule computes the subsequent price q_{t+1} . We seek a *good* posted price z , which is a price that is above the market clearing price (the number of bidders willing to pay z is at most m), but not too large so that the resulting welfare is approximately optimal (it can approximate the welfare of selling to the highest m bidders).

How to pick the block size? It is not immediately obvious why we should settle on auctioning blocks of constant size m . While allowing for larger blocks permits the blockchain to process more transactions per second, it also results in higher network propagation delay for new blocks. In the Ethereum blockchain, miners converge to an ideal block size via voting: each block is allowed to vote to increase m , decrease m , or keep m unchanged for the subsequent block [30]. As a simplifying assumption, we assume stakeholders agree on a block size m that is time invariant.

Why use a price update rule? Clearly, no single posted price can be good (in the sense defined above) if user demand changes over time. Given this, we set the goal as that of converging to a good posted price when demand is not changing. That is, the blockchain goes through epochs where the demand is stable and epochs where the demand is changing (e.g., during the launch of a new DeFi project). Although a posted price might not be good when demand is changing, once the demand stabilizes, we want the price to converge to a

⁴This assumption is somewhat stronger in proof-of-stake (PoS) blockchains where miners only commit to a block at the moment that the block is revealed (instead of committing before the block is created). This significantly increases the ways in which miners can deviate from the intended protocol [16]. However, miner myopia is still a reasonable assumption in PoS systems if we assume miners do not withhold blocks.

good price.

A dynamic posted-price mechanism is *asymptotically stable with respect to the distribution F over bidder valuations and miners' randomness* if, starting from any initial price, the sequence of expected posted prices converges to an *equilibrium price* z ; that is, if z is the current posted price, then the expected value of next posted price is also z . An iterative algorithm that uses price update rules (like ours) allows us to naturally converge to a new equilibrium price whenever the distribution over bidder valuations changes. A mechanism that converges to a posted price z is *approximately welfare optimal with respect to distribution F over bidder valuations* if selling at this limit price z gives a constant fraction of the optimal welfare (attained when selling to the top m bidders).

A challenge is how to achieve incentive compatibility for bidders together with robustness against manipulation by miners, all at the same time as allowing for adaptive prices. One concern, for example, is how to determine the winning bids when there is excess demand, i.e., when there are more than m bids with $b_i \geq q_t$. A naive approach would suggest that miners use a *maximum value* (MV) allocation rule (Definition 2.15) and serve the highest bidders. However, this is not incentive-compatible for bidders. A bidder with value $v_i \geq q_t$ would have an incentive to bid $b_i = \infty$ to maximize their chances of receiving block space. To fix this issue, we propose the *random maximal* (RM) allocation rule (Definition 2.16): whenever more than m bidders bid $b_i \geq q_t$, the miner must serve a uniformly random set with m bidders among those with $b_i \geq q_t$.⁵ This removes incentives for bidders to overbid, since any bid above q_t has an equal chance of being served.

However, a remaining challenge is to design a price update rule that provides asymptotic stability while motivating a miner to follow the RM allocation. For example, the *Welfare-Based Dynamic Posted-Price* (WDPP) mechanism (Section 5) converges exponentially quickly to a price equilibrium when the revenue curve, $\text{REV}(q) = q \cdot \bar{D}(q)$ (where q is the price and $\bar{D}(q)$ is the expected number of slots sold at price q), is Lipschitz continuous. However, under the WDPP mechanism, miners strictly prefer to implement the MV allocation rule over the RM allocation rule (even if instructed otherwise). This drives the WDPP mechanism to not be incentive compatible for bidders—from the discussion above, bidders can benefit from overbidding.

We also study the *Utilization-Based Dynamic Posted-Price* (UDPP) mechanism in Section 6, similar to the tipless mechanism [36] and based on EIP-1559's price update rule. Still, UDPP differs from EIP-1559 in that it is purely a posted-price mechanism and not a hybrid between a first-price and posted-price mechanism (since it only requests bidders to report their budget). Moreover, it differs from the tipless mechanism by adopting a randomized allocation. Although UDPP is incentive compatible for bidders and protects the blockchain from deviations from revenue-maximizing miners, we show that posted prices from the UDPP mechanism can be unstable even when demand is smaller than supply (Proposition 7.3).

To combine the desirable properties of the WDPP and UDPP mechanism (and avoid their limitations), our main result is the design and analysis of the *Truncated Welfare-Based Dynamic Posted-Price* (TWDPP) mechanism, Section 7, which is incentive compatible for myopic bidders, protects the blockchain from deviation from revenue-maximizing miners

⁵This protocol requires only private randomization and not trusted, public randomness, which is an expensive resource on blockchains [16].

(Theorem 3), and, under natural assumptions, converges to an equilibrium price (Theorem 4) in which the mechanism provably obtains a worst-case $1/4$ of the optimal welfare (Theorem 5). To prove stability, our main technical contribution is an iterative algorithm that, given oracle access to a Lipschitz continuous and strictly concave function f , converges to a fixed point of f (Lemma 7.6). Throughout the paper, we give intuitions for the theoretical results while deferring most proofs to the appendix.

In Section 8, we simulate the WDPP, UDPP, and TWDPP mechanisms. We find that the TWDPP mechanism achieves the strongest welfare of the group and obtains at least $4/5$ of the optimal welfare at the equilibrium price (outperforming the worst-case $1/4$ bound).

1.2 Related work

We have surveyed most of the relevant work above, with Roughgarden [36] being the closest to ours. Leonardos et al. [26] provides complementary stability conditions for the UDPP mechanism in the non-atomic model where the number of bidders is $n = m$ and $m \rightarrow \infty$. Our results do not make assumptions about n and m , but require Lipschitz and concavity conditions on a quantity derived from the distribution of bidder valuations for any n and m (Theorem 2) whereas their results do not.

We provide a detailed discussion of alternative transaction fee mechanism proposals in Appendix B. Lavi et al. [25] proposed the monopolistic price auction in a setting with unlimited block space where all transactions in a block pay the lowest transaction bid. Yao [40] answers several of the conjectures they introduce and shows that the monopolistic price auction is approximately incentive compatible under high demand ($n \rightarrow \infty$).

Google’s adoption of the *generalized second-price auction* (GSP) [11] was a similar response to bid cycles that had plagued first-price auctions [10]. Unlike the blockchain setting, Google had the commitment power to implement a second-price auction. Basu et al. [2] proposed a modified GSP auction as a transaction fee mechanism, and show that it is incentive compatible for bidders but like [25] requires the demand $n \rightarrow \infty$.

Dynamic mechanism design has studied the design of mechanisms in a setting where demand, preferences, and supply may change over time [33, 24, 19, 31]. To the best of our knowledge, this earlier literature has not considered the problems of credibility and information asymmetry that are a defining feature concerning the design of transaction fee mechanisms. Another body of related work is that of *distributed algorithmic mechanism design* (DAMD), which was introduced in the context of the Border Gateway Protocols (BGP) for Internet routing by Feigenbaum and Shenker [14]; this was later extended to formalize the concept of the “faithfulness” of distributed mechanisms, seeking to bring the intended rules into an *ex post* Nash equilibrium [39, 32]. In contrast to transaction fee mechanism design, this earlier work made critical use of redundant communication paths so that no one participant had a monopoly on information [27].

2 Model

We study a setting where there are m identical *block slots* to allocate at each discrete time step. A set of bidders arrives in each step, and one miner is chosen randomly to become

active. We assume miners are myopic and seek to maximize their immediate utility.

As we will explain, the miners can determine the allocation and can choose to deviate from the intended rules of a mechanism, but the miners cannot modify the payments conditioned on the allocation. These payments will be those of the dynamic, posted-price mechanism, and moreover, the prices will be updated across auctions according to the intended price update rule.

2.1 Preliminaries

The set of *active bidders* $M_t \subset \mathbb{N}_+$ denotes the bidders who arrive at time t . Bidder $i \in \cup_{t=1}^{\infty} M_t$ has a *private value* $v_i \in \mathbb{R}_+$. We consider a stylized setting where bidders are myopic—bidders are only interested in receiving a slot at time t and they do not participate at time $> t$.⁶

For each $i \in M_t$, let $b_i \in \mathbb{R}_+$ denote bidder i 's bid value. We assume lexicographic tie-breaking in the case of bidders with identical bid values. For a set of active bidders M , we define $M(q) := \{i \in M : b_i \geq q\}$. The *allocation* $B_t \subseteq \mathbb{N}$ is a set of at most m bids that receive a slot at time t .

We now define the allocation rule and payment rule for the family of dynamic mechanisms.

Definition 2.1 (Allocation Rule). An *allocation rule* for a dynamic mechanism is a vector-valued function $\vec{x}$ from history $B_1, \dots, B_{t-1}$ and active bidders M_t to an indicator $x_i(B_1, \dots, B_{t-1}, M_t) \in \{0, 1\}$ for each active bidder $i \in M_t$. The indicator is 1 if and only if bidder $i \in M_t$ receives a slot at time t . An allocation rule is *feasible* if $\sum_{i \in M_t} x_i(B_1, \dots, B_{t-1}, M_t) \leq m$. An allocation rule is *deterministic* if $x_i(B_1, \dots, B_{t-1}, M_t)$ is a deterministic variable, and *randomized* otherwise.

Definition 2.2 (Payment Rule). A *payment rule* for a dynamic mechanism is a vector-valued function $\vec{p}$ from history $B_1, \dots, B_t$ to a payment $p_i(B_1, \dots, B_t)$ for each bidder $i \in B_t$.

All bids are public to all miners, but are not available to the blockchain, i.e. the infrastructure determining payments and how prices are updated across auctions. Rather, the blockchain is only privy to information about the values of bids that are actually assigned a slot. Hence the payment rule can only depend on information disclosed in a block.

To distinguish the role of miners and the blockchain, we refer to mechanism $(\vec{x}, \vec{p})$ as the *intended allocation rule* $\vec{x}$, to be implemented by miners, and the *intended payment rule* $\vec{p}$, to be implemented by the *pricing mechanism* (e.g., the blockchain). Although the pricing mechanism, which is provided by the blockchain, can faithfully commit to implement a payment rule, miners are unable to faithfully commit to implement the allocation rule. The pricing mechanism only learns the bids in B_t and does not learn the bids in $M_t \setminus B_t$. Thus, miners can try to change the allocation in their favor, including the effect this might have on future payments. As an example, a miner can impersonate a set of *fake bidders*, $F_t \subset \mathbb{N}_+$, at time t , and each fake bidder $i \in F_t$ can submit a *fake bid* $b_i \in \mathbb{R}$ and be allocated a slot (see Example 2.4).

⁶Huberman et al. [21] considers the status quo first-price auction in the more general setting where bidders are interested in reducing their confirmation time rather than being included in the subsequent block.

We assume that bidders have *quasilinear utility*. That is, the utility of bidder $i \in M_t$ is

$$u_i(v_i, b_i) := x_i(B_1, \dots, B_{t-1}, M_t) \cdot (v_i - p_i(B_1, \dots, B_t)). \quad (1)$$

We only consider mechanisms that transfer all payments received from bidders to the active miner.⁷ The miner's utility is a function u_0 from allocation B_t to the revenue of the miner (net any payments from fake bids):

$$u_0(B_t) := \sum_{i \in B_t \setminus F_t} p_i(B_1, \dots, B_t). \quad (2)$$

2.2 Decentralized multi-round auction game

We now define the sequence of steps that models this transaction fee market.

Definition 2.3 (Decentralized multi-round auction game). The *decentralized multi-round auction game* is a multi-round game between multiple miners and the pricing mechanism. We are given an allocation rule $\vec{x}$, a payment rule $\vec{p}$, and initially all miners are *inactive* and the time step is $t = 1$. The game proceeds as follows:

- (1) All miners observe M_t and all bids $b_i \in M_t$.
- (2) One miner is chosen at random to become *active*.
- (3) The active miner chooses and announces an allocation

$$B_t := \{i \in M_t : x_i(B_1, \dots, B_{t-1}, M_t) = 1\}.$$

- (4) The pricing mechanism observes B_t and all bids $b_i \in B_t$.⁸
- (5) For each bidder $i \in B_t$, the pricing mechanism charges $p_i(B_1, \dots, B_t)$.
- (6) The active miner receives a payment $\sum_{i \in B_t} p_i(B_1, \dots, B_t)$.
- (7) The active miner becomes inactive. Increment t and return to step (1).

Unlike standard mechanism design [28], miners are responsible for choosing B_t , and yet at the same time, cannot commit to implementing a particular allocation rule $\vec{x}$; that is, an active miner can corrupt step (3) of the game in any way they want. Miners may even allocate slots to fake bidders, $i \in F_t$.

Because the pricing mechanism can only observe bids in B_t and cannot observe $M_t \setminus B_t$, the payments must depend only on $B_1, \dots, B_t$.

Blockchains can commit to executing honest code and can faithfully implement steps (5) and (6). Thus, miners can manipulate the allocation B_t of the intended mechanism, but not what the pricing mechanism computes once it learns B_t .

⁷Roughgarden [36] shows that not awarding miners the revenue of the auction, which is what EIP-1559 does, is important to avoid collusion of bidders with miners. All of our results remain untouched if one modifies Equation 2 to transfer $\delta \in (0, 1]$ fraction of the revenue to the active miner.

⁸The integrity of bid values can be enforced via digital signatures.

Example 2.4. Suppose there is a single slot for sale, and the mechanism $(\vec{x}, \vec{p})$ is a first-price auction (the highest bidder receives the slot and pays his bid). Suppose there are two active bidders at time t , $M_t = \{1, 2\}$, with identical values $v_1 = v_2 = 2$ and distinct bids $b_1 = 1, b_2 = 2$. While the values are private to bidders, the bids become known to miners at Step (1). At Step (3), the active miner is instructed by the allocation rule to allocate the slot to bidder 2, but can make other choices. For example, the active miner can allocate the slot to bidder 1 and receive utility \$1; allocate to bidder 2 and receive utility \$2; allocate the slot to neither bidder and receive utility \$0; or even allocate the slot to a fake bidder with some bid b and receive a utility of $v - v = \$0$ (i.e., they pay the mechanism at Step (5), but the mechanism returns the payment at Step (6)). In this example, the active miner's optimal action is to allocate the slot to bidder 2 and receive utility \$2. Thus, in this example, it is in the best interest for active miners to implement the intended mechanism.

2.3 Mechanism design

We study the stability and the social welfare of a mechanism in the *probabilistic setting*. For simplicity, assume n bidder values are drawn i.i.d. from a continuous *cumulative density function* F :

$$F(x) := \Pr[v \leq x]. \quad (3)$$

Although our welfare guarantees require the i.i.d. assumption, our conditions for stability do not. Instead, our conditions for stability will require Lipschitz and concavity conditions on quantities derived from the distribution where bidder valuations are drawn. We define $\vec{v} := \{v_1, v_2, \dots, v_n\} \in \mathbb{R}_+^n$ as an n -dimensional random variable, where $v_1, v_2, \dots, v_n$ are drawn i.i.d. from F .

Definition 2.5 (Social Welfare). For a mechanism $(\vec{x}, \vec{p})$ with history $B_1, B_2, \dots, B_{t-1}$, let $M_t = [n] = \{1, 2, \dots, n\}$ be a set of bidders with values $\vec{v} = \{v_1, v_2, \dots, v_n\}$ drawn i.i.d. from F . The expected *social welfare* at time t with respect to allocation rule $\vec{x}$ is

$$\text{WELFARE}_t(\vec{x}) := \mathbb{E}_{\vec{v} \sim F^n} \left[\sum_{i=1}^n v_i \cdot x_i(B_1, \dots, B_{t-1}, M_t) \middle| B_1, \dots, B_{t-1} \right]. \quad (4)$$

We drop the subscript t and write $\text{WELFARE}(\vec{x})$ when the time step is clear from context. A mechanism is *efficient* if it extracts the maximum possible social welfare. Let $\vec{x}^*(\vec{v})$ be the welfare-maximizing feasible allocation for a particular value profile $\vec{v}$, with $\sum_{i=1}^n x_i^*(\vec{v}) = m$. The *optimal expected social welfare* is:

$$\text{OPT} := \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \right]. \quad (5)$$

We will consider a mechanism as having a *good user experience* if it is *ex post Nash incentive compatible (IC)* and *individually rational (IR)* for bidders—bidders have the incentive to bid their valuation and pay at most their bid.

Definition 2.6 (Ex Post Nash Equilibrium). A *strategy* for bidder i is a function b_i^* from value v_i to bid $b_i^*(v_i)$. A strategy profile $b^*(\cdot)$ is an *ex post Nash equilibrium* (EPNE) for a

mechanism $(\vec{x}, \vec{p})$ if for all histories $B_1, \dots, B_{t-1}$, for all active bidders M_t , for all bidders $i \in M_t$, for all values v_i , bidding $b_i^*(v_i)$ maximizes the utility (1) of bidder i conditioned on all $j \neq i \in M_t$ following strategy $b_j^*(v_j)$.

Definition 2.7 (Ex Post Incentive Compatible). A mechanism $(\vec{x}, \vec{p})$ is *ex post incentive compatible* (IC) for bidders if it has a *truthful ex post* Nash equilibrium where all real bidders bid their true value. That is, $b_i^*(v_i) = v_i$ for all $i \in \cup_{t=1}^{\infty} M_t$.

Note dominant-strategy IC rather than *ex post* IC is too strong for the blockchain setting because the pool of bids is public and, in principle, bidders could choose to condition their strategies on the reports of others. Yet *ex post* IC is an appropriate and typical solution concept for dynamic settings [6, 3], and provides a useful strategic simplification and improvement in ease of use: it is optimal for a bidder to report truthfully, no matter its actual value and the values of others, as long as others are also truthful.

Definition 2.8 (Individual Rationality (IR)). A mechanism $(\vec{x}, \vec{p})$ is *individually rational* if for all histories $B_1, B_2, \dots$, bidders pay at most their bids: $p_i(B_1, B_2, \dots) \leq b_i$ for all bidders i .

Definition 2.9 (Dominant Strategy Incentive Compatible (DSIC) for Myopic Miners). A mechanism $(\vec{x}, \vec{p})$ is *DSIC for myopic miners* if for all histories $B_1, \dots, B_{t-1}$, for all active bidders M_t , the active miner at time t maximizes their utility (2) by implementing the allocation rule $\vec{x}$. That is, the miner sets $F = \emptyset$ and chooses $B_t = \{i \in M_t : x_i(B_1, \dots, B_{t-1}, M_t) = 1\}$.

To ensure the bidder's myopia is an adequate assumption, we will aim to design dynamic mechanisms that are stable. To be concrete: whenever the demand for blocks is at a stable state, we will require the mechanism to reach a stable state, at which point a bidder will receive the same expected utility regardless of the time they submit a bid. For simplicity, we first give a notion of stability suitable for the deterministic setting and in Section 3, we provide a more general definition for the probabilistic setting.

Definition 2.10 (Deterministic Stability). Assume the same set M of bidders with bid profile $\vec{b}$ arrives at all time steps, i.e. $M_t = M$ for all t . The mechanism $(\vec{x}, \vec{p})$ is *stable* with respect to the deterministic demand M if for all bidders $i \in M$, the allocation rule $x_i(B_1, B_2, \dots, B_{t-1}, M_t)$ converges to $x_i^*(\vec{b})$ and the payment rule $p_i(B_1, B_2, \dots, B_t)$ converges to $p_i^*(\vec{b})$ as $t \rightarrow \infty$.

Note if mechanism $(\vec{x}, \vec{p})$ is stable, then the dynamic mechanism converges to a static mechanism $(\vec{x}^*, \vec{p}^*)$. If $\vec{b}$ is the bid profile of all bidders, then the utility of bidder i with value v_i is $(v_i - p_i^*(\vec{b})) \cdot x_i^*(\vec{b})$ which is independent of the time they choose to participate which motivates bidder's myopia.

Below we summarize the desiderata for the design of transaction fee mechanisms for decentralized blockchains:

- **Communication Complexity.** The mechanism requires $O(m)$ communication between bidders and the blockchain.⁹

⁹The first-price, second-price, and posted-price mechanisms have communication complexity $O(m)$. As a distinction, the communication cost between miners and the blockchain on known implementations of credible, strategyproof, optimal auctions would scale with n , the number of bidders, not the number of slots.

Static mechanisms	IC for bidders	DSIC for miners
First-Price (2.12)	No	Yes
Second-Price (2.13)	Yes	No
DPP (2.17)	No	Yes
RPP (2.17)	Yes	Yes
EIP-1559 (B.1)	Sometimes	Yes
Monopolistic Price (B.2)	Approximately	Yes
RSOP (B.3)	Yes	No
Modified GSP (B.4)	Approximately	Approximately

Table 1: Bidder payments and miner revenue under various static mechanisms. q_t denotes a posted-price. At time t , M_t denote the set of active bidders and B_t denotes the allocation.

- **Ex Post Incentive Compatible for Bidders.** The mechanism has a truthful *ex post* Nash equilibrium for bidders.
- **Individually Rational.** If bidders receive a slot, they pay at most their bid; otherwise, they pay nothing.
- **DSIC for Myopic Miners.** In each period, it is a weakly dominant strategy for miners to implement the intended allocation rule.
- **Stable.** The mechanism converges to a stable state.
- **Approximately Welfare Optimal.** If in each period, n i.i.d. bidders are drawn from distribution F , then the mechanism converges to an equilibrium price.¹⁰ At the equilibrium price, the expected welfare is $\Omega(OPT)$, where OPT is the optimal expected social welfare.

2.4 Illustrative mechanisms

Next, to motivate the design of dynamic mechanisms, we highlight undesirable properties from traditional static mechanisms. Table 1 summarizes the properties of existing proposals. We also provide further discussion in Appendix B.

Definition 2.11 (Static mechanism; Dynamic mechanism). A mechanism $(\vec{x}, \vec{p})$ is *static* if its allocation and payment rule depend only on the most recent bids and allocation. That is, for all histories $B_1, \dots, B_{t-1}$, we have $x_i(B_1, \dots, B_{t-1}, M_t) = x_i(M_t)$ and $p_i(B_1, \dots, B_t) = p_i(B_t)$. A mechanism is *dynamic* if it is not static.

Definition 2.12 (First-price auction). A *first-price* (FP) auction is a static mechanism where the top m bidders receive a slot and pay their bid.

Nakamoto [29] proposed the first-price auction as the default transaction fee mechanism for Bitcoin.

¹⁰We consider asymptotic stability with respect to the distribution F over bidder valuations as our formal definition for a stable equilibrium (Definition 3.5).

An important property of the first-price auction is that it is DSIC for myopic miners. Miners can only lose revenue by replacing a real bid $i \in M_t$ with a fake bid $j \in F_t$. As a downside, the first-price auction is known to provide a bad user experience as a transaction fee mechanism because it has no *ex post* Nash equilibrium. A bidder wants to bid the minimum possible that would allow her to win a slot, but that requires the bidder to monitor all other bids and adapt her strategy.

Definition 2.13 (Second-price auction). A *second-price* (SP) auction is a static mechanism where the top m bidders receive a slot and pay the $(m + 1)$ -st highest bid.

Unlike the first-price auction, truthful bidding $b_i(v_i) = v_i$ is an *ex post* Nash equilibrium for the second-price auction, making it easy for bidders to participate. However, to implement its payment rule, miners must commit to truthfully reporting the value of $(m + 1)$ -st highest bid with the mechanism (note the payment rule depends not only on the allocation B_t but also on unallocated transactions $M_t \setminus B_t$). Thus miners are better off by reporting a fake bid that equals the m -th highest bid, and the second-price auction is not DSIC for myopic miners.

Next, we consider posted-price mechanisms and two possible allocation rules.

Definition 2.14 (Maximal Allocation). For a set of active bidders M , we say $S \subseteq M$ is a *maximal allocation* with respect to M if the following is true:

- S contains at most m bidders.
- If S contains less than m bidders, $S = M$.

Let 2^M denote the *collection of maximal allocations* with respect to M :

$$2^M := \{B \subseteq M : |B| \leq m \text{ is a maximal allocation}\}. \quad (6)$$

Definition 2.15 (Maximum value allocation rule). The *maximum value* (MV) allocation rule $\vec{x}^{MV}(M)$ selects a maximal allocation $B \in 2^M$ that maximizes bid values: $B = \arg \max_{B' \in 2^M} \sum_{i \in B'} b_i$.

Let $Unif(\{B_1, \dots, B_\ell\})$ denotes the uniformly random distribution over $\{B_1, \dots, B_\ell\}$.

Definition 2.16 (Random maximal allocation rule). The *random maximal* (RM) allocation rule $\vec{x}^{RM}(M)$ samples a maximal allocation $B \in 2^M$ from the distribution $Unif(2^M)$.

Definition 2.17 (Posted-price mechanism). A *posted-price* (PP) mechanism with posted price q is a static mechanism $(\vec{x}, \vec{p})$ where each bidder that receives a slot pays q and all other bidders pay nothing. We refer to $(\vec{x}(q), \vec{p}(q))$ as the posted-price mechanism with posted price q . A posted-price mechanism with posted price q is a *randomized posted-price* (RPP) mechanism if for all M , $\vec{x}(M) = \vec{x}^{RM}(M(q))$ is the RM allocation rule on bids $M(q) = \{i \in M : b_i \geq q\}$. A posted-price mechanism with posted price q is a *deterministic posted-price* (DPP) mechanism if for all M , $\vec{x}(M) = \vec{x}^{MV}(M(q))$ is the MV allocation rule on bids $M(q)$.

Posted-price mechanisms are DSIC for myopic miners since they commit to a fixed-price that miners cannot influence, and thus, the best a miner can do is accept a maximal allocation of bids priced at or above the price q .

Depending on how the posted-price mechanism allocates slots, the mechanism can also be incentive compatible for bidders. The posted-price mechanism with the RM allocation rule is *ex post* IC since any bidder with $b_i \geq q$ has an equal chance to receive a slot irrespective of its bid. On the other hand, the posted-price mechanism with the MV allocation rule is not *ex post* IC. Whenever more than m bidders have value $v_i \geq q$, there exists a bidder with value $v_i \geq q$ that will not receive a slot. If that bidder bids $b_i = \infty$ instead of v_i , they receive the slot and pay $q \leq v_i$.

2.5 Discussion

The first-price, second-price, and posted-price mechanisms are individually rational and require little communication with the blockchain. This communication efficiency is a first-order concern for decentralized blockchains because any information stored in the blockchain must propagate through the whole network. In particular, these mechanisms require miners to report $O(m)$ bids to the mechanism to implement the payment rule—the communication scales as the number of slots, not the number of bidders.

Posted-price mechanisms are also naturally robust to miners’ strategic behavior and are incentive compatible for bidders provided the posted price is high enough (so that the supply is greater than the demand at the posted price). While first-price auctions are not incentive compatible for bidders, they are resilient to manipulation by miners. Finally, second-price auctions are incentive compatible for bidders but not resilient to manipulation by miners.

In terms of social welfare, the second-price auction is efficient since truthful bidding is a dominant strategy and the highest bidders receive block slots. In theory, first-price auctions are also efficient assuming bids are drawn i.i.d. from some distribution F and under the assumption bidders follow a Bayes-Nash equilibrium. However, bids in a first-price auction do not typically converge to an equilibrium in repeated settings, as was observed for paid search engines [10], and the current blockchain transaction fee market is an example of the kind of instability that can occur. Sequential posted-price mechanisms—where the auctioneer visits bidders to make a take-or-leave offer—can obtain $1/2$ of the optimal welfare, but require prior knowledge on the distribution F and the demand n [7, 23].

3 Dynamic Posted-Price Mechanisms

The rationale behind a dynamic posted-price mechanism is that even though the demand curve is unknown and dynamic over time, we can hope to learn a proxy for the market-clearing price from transactions included in preceding blocks. We must be careful, though, since we must align the mechanism with miners’ incentives to implement the allocation rule faithfully.

Definition 3.1 (Price Update Rule). A *price update rule* is a real-valued function T from positive q and feasible allocation B with respect to q to the subsequent price $T(q, B)$ where $b_i \geq q$ for all $i \in B$.

Remark: we use B to refer to the bidders that receive a slot. Given $i \in B$, the update rule has access to the bid b_i of bidder i .

Definition 3.2 (Dynamic Posted-Price Mechanism). A *dynamic posted-price* (DPP) mechanism $(\vec{x}, \vec{p}, T)$ is a dynamic mechanism endowed with a price update rule T . Let positive q_0 denote the initial posted price. Given history $B_1, B_2, \dots, B_{t-1}$ and active bidders M_t , define the posted price $q_t := T(q_{t-1}, B_{t-1})$ at time t . The mechanism $(\vec{x}, \vec{p}, T)$ is a *randomized dynamic posted-price* (RDPP) mechanism if, at time step t , it implements the RM allocation rule:

$$\vec{x}(B_1, B_2, \dots, B_{t-1}, M_t) := \vec{x}^{RM}(M_t(q_t)).$$

The mechanism $(\vec{x}, \vec{p}, T)$ is a *deterministic dynamic posted-price* (DDPP) mechanism if, at time step t , it implements the MV allocation rule:

$$\vec{x}(B_1, B_2, \dots, B_{t-1}, M_t) := \vec{x}^{MV}(M_t(q_t)).$$

We study the *stability* of the expected value of the random operator $T(q, B)$ when at each time step, n bidders with values $\vec{v} = (v_1, v_2, \dots, v_n)$ are drawn from distribution F^n , with the assumption that bidders are truthful (bid their value). We define the real-valued function E_T to be

$$E_T(q) := \mathbb{E}_{\vec{v} \sim F^n} [T(q, B)], \quad (7)$$

where the expectation taken is over $v_1, v_2, \dots, v_n$ and over the allocation

$$B = \{i \in M : x_i(q) = 1\}.$$

Before we formally define stability for a DPP mechanism, we define the fixed point iteration of a real-valued function. We use (X, d) to refer to a complete metric space where $X \subseteq \mathbb{R}$ and d is the Euclidean norm.

Definition 3.3 (n -iterate). The n -iterate of function $f : X \rightarrow X$, for integer $n \geq 0$, is the function

$$f^n(q) := \begin{cases} q & \text{for } n = 0, \\ f(f^{n-1}(q)) & \text{for } n \geq 1. \end{cases} \quad (8)$$

A point x is a *fixed point* for function $f : X \rightarrow X$ if and only if $f(x) = x$.

Definition 3.4 (Fixed-point iteration). Given a continuous function $f : X \rightarrow X$ and an initial point $q_0 \in X$, the *fixed-point iteration* of f starting from q_0 is the sequence of function evaluations

$$q_0, f(q_0), f^2(q_0), f^3(q_0), \dots \quad (9)$$

If the sequence converges to some q^* , then from the continuity of f , q^* is a fixed point of f .

Definition 3.5 (Asymptotic Stability). A DPP mechanism $(\vec{x}, \vec{p}, T)$ is *asymptotically stable* with respect to distribution F over n bidder valuations if, for *any* positive initial price q_0 , the fixed-point iteration of E_T starting from q_0 converges to the *equilibrium price* q^* . A DPP mechanism is *unstable* if it is not asymptotically stable.

Definition 3.6 (Approximately Welfare Optimal at Equilibrium). A DPP mechanism is *approximately welfare optimal* if it is asymptotically stable with equilibrium price q^* and the welfare at equilibrium $\text{WELFARE}(\vec{x}(q^*)) \geq \Omega(\text{OPT})$ is a constant approximation of the optimal welfare.

There are two main reasons why a DPP mechanism could be unstable: prices could grow unbounded, or prices could oscillate in a closed orbit. The proofs of Propositions 6.5 and 7.3 give examples of mechanisms that are unstable when the distribution F is a point mass.

The convergence of prices is also important in justifying our assumption that bidders are myopic and will immediately submit bids for the next auction. If prices were to oscillate, bidders might seek to time the moment they enter their bids, waiting for moments where prices are lower, driving more instability and uncertainty in the market.

4 Asymptotically Stable Update Rules via Fixed Point Theory

There is a rich theory on the study of fixed points [8, 20] for monotone decreasing operators and in many settings, the existence of a fixed point is guaranteed. In general, we do not expect a function $f : \mathbb{R} \rightarrow \mathbb{R}$ to have a fixed point, for example $f(x) = x + 1$. However, if a continuous function f is non-increasing, then f always intersects the line $y(x) = x$. Let z denote the x value corresponding to this intersection. Then z is a fixed point for function f , since $f(z) = y(z) = z$.

Definition 4.1 (Lipschitz continuity). A real-valued function $f : X \rightarrow \mathbb{R}$ is L -Lipschitz if for all $x, y \in X$, $|f(x) - f(y)| \leq L \cdot |x - y|$. We say f is *Lipschitz continuous* if there is a positive constant L such that f is L -Lipschitz. We say that f is a *constant function* if and only if f is 0-Lipschitz.

Definition 4.2 (Contractive mapping). A *contractive mapping* is a function $f : X \rightarrow X$ with the property that f is Lipschitz continuous with constant $0 \leq L < 1$.

Definition 4.3 (Monotone mixture). A *monotone mixture* f with kernel $g : X \rightarrow X$ is a function $f(x) = \alpha \cdot g(x) + (1 - \alpha) \cdot x$ where the kernel is nonconstant, nonincreasing and $\alpha \in [0, 1]$ is the convergence parameter.

Observation 4.4. If $f(x) = \alpha g(x) + (1 - \alpha)x$, then x is a fixed point for f if and only if x is a fixed point for g .

Proof. Assume $x = f(x)$, then $g(x) = x$. Conversely, assume $x = g(x)$, then $f(x) = x$. $\square$

Thus, for any continuous nonincreasing function g , we can use Observation 4.4 to construct a function f that has the same fixed point as g . Then, we can search for a fixed point for g by searching for a fixed point for f . For this, we use the following well known fact:

Lemma 4.5. If $f(x) = \alpha g(x) + (1 - \alpha)x$ is a monotone mixture with a nonconstant L -Lipschitz kernel g , then for all $0 \leq \alpha \leq \frac{1}{L+1}$, f is $(1 - \alpha)$ -Lipschitz. Additionally, if f maps X to itself, the restriction of f to X is a contractive mapping.

Proof. Fix $x, y \in X$ and w.l.o.g. assume $x < y$, then $g(x) \geq g(y)$ is nonincreasing and

$$\begin{aligned}
|f(x) - f(y)| &= |\alpha(g(x) - g(y)) + (1 - \alpha)(x - y)| \\
&= |\alpha(g(x) - g(y)) - (1 - \alpha)(y - x)| \\
&\leq \max\{\alpha|g(x) - g(y)|, (1 - \alpha)|x - y|\} \\
&\leq \max\{\alpha L|x - y|, (1 - \alpha)|x - y|\} \\
&\leq (1 - \alpha)|x - y|.
\end{aligned}$$

The last step is from the assumption $\alpha \cdot L + \alpha \leq 1$. This proves f is $(1 - \alpha)$ -Lipschitz. For the “additionally” part, if f maps X to itself, the fact that g is non-constant implies $L > 0$, thus $\alpha < 1$ and the restriction of f to X is a contraction. $\square$

Lemma 4.6 (Banach Contraction Principle [37]). *If a real-valued function $f : X \rightarrow X$ is a contraction, then f has a unique fixed point x^* . To find x^* , start at any $x_0 \in X$, then $\lim_{n \rightarrow \infty} f^n(x_0) = x^*$.*

Thus, if we design a DPP mechanism $(\vec{x}, \vec{p}, T)$ such that for distribution F over bidder valuations the expected value E_T is a monotone mixture with an L -Lipschitz kernel and convergence parameter $\alpha \leq \frac{1}{L+1}$, we will derive that $(\vec{x}, \vec{p}, T)$ is asymptotically stable with respect to F . This will be the main idea behind the proof that the WDPP mechanism is stable (Theorem 1). The proof of stability of the UDPP mechanism and the TWDPP mechanism will be more involved because their kernels will not be monotone mixtures, but the proof will still rely on the Banach contraction principle.

5 Welfare-Based Dynamic Posted-Price Mechanism

This section describes our first proposal and gives sufficient conditions for our mechanism to be asymptotically stable and to obtain a 1/2-approximation of the optimal welfare.

To maximize social welfare, we must guarantee the highest-valued transactions are included in each block. The challenge is that whenever there are more than m bidders with bids $b_i \geq q_{t-1}$, the mechanism is not *ex post* IC for bidders when miners use the MV allocation rule. If we use the RM allocation rule, a low bid may receive a slot in place of a high bid. In this case, the mechanism must increase the posted price. A welfare-based approach uses bids of allocated bidders as a proxy for the degree of excess demand:

Definition 5.1 (Welfare-based update rule). The α -welfare-based update rule is the update rule:

$$T_W^\alpha(q, B) := \alpha \frac{1}{m} \sum_{i \in B} b_i + (1 - \alpha)q, \quad (10)$$

where $\alpha \in [0, 1]$ is the convergence parameter.

Observe that whenever $|B_t| = m$, $q_t = T_W^\alpha(q_{t-1}, B_{t-1}) > q_{t-1}$ unless $b_i = q_t$ for all $i \in B_{t-1}$. Thus the hope is that the mechanism converges to a price z at least as big as the market clearing price—the price where at most m bidders have value $v_i \geq z$.

Unfortunately, miners' myopia is an unreasonable assumption for a dynamic posted-price mechanism with the RM allocation rule endowed with the welfare-based update rule. To see this, note that if the current price q_{t-1} is smaller than the market-clearing price, there exists a price $q > q_{t-1}$ that also sells to m buyers. Thus under a demand surge, the active miner might have a strict preference for implementing the MV allocation rule instead of the RM allocation rule to drive prices up for the next time when they become the active miner. Once miners signal a preference for adding the highest bids, each bidder with $v_i \geq q_t$ would have a dominant strategy to bid $b_i > v_i$ to maximize the probability of being selected. Thus, the dynamic posted-price mechanism would not be *ex post* IC for myopic bidders.¹¹

Since a DPP mechanism with the welfare-based update rule and the RM allocation rule reduces to a DPP mechanism with the MV allocation rule, we study a DPP mechanism endowed with the welfare-based update rule and the MV allocation rule:

Definition 5.2 (Welfare-Based Dynamic Posted-Price mechanism). The *welfare-based dynamic posted-price* (WDPP) mechanism is a DDPP mechanism $(\vec{x}, \vec{p}, T_W^\alpha)$ endowed with the welfare-based update rule T_W^α .

The WDPP mechanism is not *ex post* IC. However, we show the WDPP mechanism is asymptotically stable assuming truthful bidding as a warm-up to studying *ex post* IC mechanisms in the following sections.

5.1 Asymptotic stability of the WDPP mechanism

We will now describe our main results for the WDPP mechanism. In particular, we show that by setting a sufficiently small α , the WDPP mechanism is asymptotically stable and approximately welfare optimal under the mild assumption the revenue curve is Lipschitz continuous.

Definition 5.3 (Demand Curve). The *demand at price q* is the random variable $N(q) := \sum_{i=1}^n \mathbb{1}_{v_i \geq q}$. A *demand curve* is a non-increasing real-valued function D mapping a price q to the expected number of buyers $D(q) := \mathbb{E}[N(q)]$ that would purchase at price q . The *demand curve with limited supply* is $\bar{D}(p) := \mathbb{E}[\min\{m, N(q)\}]$.

Definition 5.4 (Revenue Curve). A *revenue curve* is a real-valued function from price q to the expected revenue $\text{REV}(q) := q \cdot \bar{D}(q)$ when selling a supply of m slots at price q .

We can show the WDPP mechanism is asymptotically stable as long as the revenue curve $\text{REV}(q)$ is Lipschitz continuous and α is upper bounded by a function of the Lipschitz constant of $\text{REV}(q)$.

¹¹A possible counterargument to bidders overbidding comes from *Mempools*—global decentralized databases of pending transactions—where Bitcoin and other blockchains store bids until they are processed and added to the blockchain or removed after 48 hours. With Mempools, there is *carryover* from $M_t \setminus B_t$ to M_{t+1} (if not served, bidders participate in the next step). Thus if a bidder bids $b_i > v_i$ in the hope of increasing their chances of being included in B_t , there is a risk their bid will carry over to the next time step. Since there is uncertainty about future prices (since $M_{t+1}, M_{t+2}, \dots$ are unknown at time t), there is a risk bidder i could pay $b_i > v_i$ and receive a negative utility. It is hard to model carryover, and incorporating Mempools into the model is an interesting avenue for future work. Here, we assume each bidder only participates once.

Theorem 1. *If the revenue curve REV is L -Lipschitz and $\alpha \leq \frac{1}{L/m+1}$, the WDPP mechanism $(\vec{x}, \vec{p}, T_W^\alpha)$ is asymptotically stable with respect to F^n and has a unique equilibrium price q^* . Moreover, it is approximately welfare optimal: $\text{WELFARE}(\vec{x}(q^*)) \geq \frac{\text{OPT}}{2}$.*

We defer the proof to the appendix. As a proof sketch recall that a mechanism is asymptotically stable if the iteration $E_{T_W^\alpha}(q_0), E_{T_W^\alpha}(q_1), \dots$ converges to an equilibrium price z . Expanding $E_{T_W^\alpha}(q_0)$, we have

$$E_{T_W^\alpha}(q_0) = \alpha \frac{\text{WELFARE}(\vec{x}(q_0))}{m} + (1 - \alpha)q_0.$$

Since miners allocate block space to the highest bidders, we have that $\text{WELFARE}(\vec{x}(q_0))$ is a decreasing function of q_0 . Therefore, $E_{T_W^\alpha}$ is a monotone mixture as defined in Section 4. One can show that if $\text{REV}(q_0)$ is Lipschitz continuous, then so is the $\text{WELFARE}(\vec{x}(q_0))$. Thus we can use the ideas of Section 4 to prove that the WDPP mechanism is asymptotically stable.

To prove that the WDPP mechanism is approximately welfare optimal, recall that we need to show $\text{WELFARE}(\vec{x}(z)) \geq \Omega(\text{OPT})$ where z is an equilibrium price. Note that z is an equilibrium price if and only if $E_{T_W^\alpha}(z) = z$ and observe that $E_{T_W^\alpha}(z) = \alpha f(z) + (1 - \alpha)z$ is a convex combination of z with a function f of z . Observe that z is also a fixed point of f : $f(z) = z$. For the WDPP mechanism, $z = f(z) = \frac{\text{WELFARE}(\vec{x}(z))}{m}$, which we use when showing $\text{WELFARE}(\vec{x}(z)) \geq \frac{\text{OPT}}{2}$.

Remark. For all dynamic posted-price mechanisms $(\vec{x}, \vec{p}, T)$ we consider, we will always have that $E_T(z) = \alpha f(z) + (1 - \alpha)z$ (i.e. E_T is a convex combination of z with some function f of z). Thus if the price update rule T is asymptotically stable, it converges to an equilibrium price z where $f(z) = z$, which we use to show the mechanism is approximately welfare optimal.

6 Utilization-Based Dynamic Posted-Price Mechanism

To improve the robustness of the welfare-based dynamic posted-price mechanism to strategic behavior from miners, which in turn unravels the IC property for bidders, we can remove the influence of bid values on the price update when there is excess demand. That is, we can use a fixed multiplicative price change that is independent of bid values:

Definition 6.1 (Utilization-Based Update Rule). The (α, δ) -utilization-based update rule is the update rule

$$T_U^{\alpha, \delta}(q, B) := \alpha \frac{|B|}{m} (1 + \delta)q + (1 - \alpha)q. \quad (11)$$

where $\alpha \in (0, 1)$ is the convergence parameter and $\delta \in (0, \infty)$.

The utilization-based update rule with $\delta = 1$ and $\alpha = 1/8$ is the update rule used by EIP-1559 to dynamically set its posted price. While EIP-1559 uses a first-price auction to choose the allocation when there is excess demand, we adopt as the intended allocation rule the RM allocation rule which, in turn, removes incentives for non-truthful bidding by bidders:

Definition 6.2 (Utilization-Based Dynamic Posted-Price Mechanism). The *utilization-based dynamic posted-price* (UDPP) mechanism is an RDPP mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$ endowed with the utilization-based update rule $T_U^{\alpha, \delta}$.

Proposition 6.3. *The RDPP is ex post IC for myopic bidders and DSIC for myopic miners.*

Proof. First, we show the RDPP mechanism is DSIC for myopic miners. For a particular time step, let M denote the set of active bidders and let q denote the current posted price. The active miner maximizes utility by selecting a maximal allocation from $M(q)$, and any maximal allocation weakly dominates all other allocations. Thus choosing a uniformly random maximal allocation is a weakly dominant strategy. This proves the RDPP mechanism is DSIC for myopic miners.

Next, we show the RDPP mechanism is *ex post* IC for myopic bidders. Given that miners select a uniformly random maximal allocation from $M(q)$, we first consider the case where bidder i has value $v_i < q$. If bidder i bids $b_i \geq q$, there is a positive probability the active miner selects bidder i to receive a slot and pay q resulting in a negative utility. Thus bidding v_i strictly dominates bidding $b_i \geq q$ and weakly dominates bidding $b_i \in (v_i, q)$. Next, suppose bidder i has value $v_i \geq q$. For bidder i to receive a slot and pay q , she must bid $b_i \geq q$. Thus bidding v_i strictly dominates bidding $b_i < q$. For any bid $b_i \geq q$, bidder i has equal probability of winning a slot and paying q . Thus bidding v_i weakly dominates bidding $b_i > v_i$. This proves truthful bidding weakly dominates bidding $b_i \neq v_i$, and the mechanism is IC for myopic bidders. $\square$

Corollary 6.4. *The UDPP mechanism is ex post IC for myopic bidders and DSIC for myopic miners.*

Remark. Proposition 6.3 shows that a randomized dynamic posted-price mechanism shares the benefits of the tipless mechanism [36] (both are *ex post* IC for myopic bidders and DSIC for myopic miners). However, they are vulnerable to off-chain agreements [35] in that collusion between miners and bidders could improve the sum of their utilities: whenever the current posted price is smaller than the market-clearing price, the active miner can run an auction outside the blockchain to select the top m transactions instead of randomizing the allocation. This deviation can improve the joint utility of miners and bidders since the RM allocation might allocate block space to a low value bidder in place of a higher value bidder. On the positive side, if the posted price is above the market-clearing price, off-chain agreements do not improve the joint utility of bidders and miners. Any additional revenue miners receive results in a higher payment from a bidder that would still receive block space if there were no off-chain agreement.

Remark. A dynamic posted-price mechanism with the welfare-based update rule and the RM allocation rule (Definition 2.16) is also DSIC for myopic miners. However, the crucial difference is that there is a clear and sensible out-of-model deviation by non-myopic miners to that rule, while this deviation has gone away in the new design.

Interestingly, the fact the UDPP mechanism ignores bid values when there is excess demand, which is important for incentive alignment, can lead to instability in instances in

which the welfare-based dynamic-posted mechanism is asymptotically stable. This motivates the study of more stable update rules in Section 7.¹²

Proposition 6.5. *There exists a distribution F^2 where the WDPP mechanism $(\vec{x}, \vec{p}, T_W^\alpha)$ is asymptotically stable but the UDPP mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$ is unstable.*

Proof. For any positive v , let F be the distribution where each bidder has value v with probability 1. Suppose there is a single slot for sale and two bidders drawn i.i.d. from F in each time step. For the UDPP mechanism, if the current price $q_t \leq v$, the mechanism observes excess demand and the subsequent price increases to $q_t(1 + \alpha\delta)$. However, when the current price $q_t > v$, neither bidder will purchase, and the subsequent price decreases to $q_t(1 - \alpha)$. This repeats *ad infinitum*, and the utilization-based update rule is unstable with respect to F^2 .

For the welfare-based dynamic posted-price mechanism, whenever the current price $q_t > v$, neither bidder will purchase, and the subsequent price decreases to $q_t(1 - \alpha)$. Thus eventually, the mechanism reaches a price $q_t \leq v$. Whenever the current price $q_t \leq v$, the active miner allocates a single slot and the subsequent price increases to $q_t + \alpha(v - q_t)$ where $v - q_t \geq 0$ is the surplus of the mechanism. Thus the subsequent price is $q_t < q_{t+1} = \alpha v + (1 - \alpha)q_t \leq v$ whenever $q_t < v$ (and $q_{t+1} = v$ whenever $q_t = v$). This proves for all time steps $n \geq t$, $q_t \leq v$, and one bidder receives the slot. The sequence $q_t, q_{t+1}, \dots$ is increasing and bounded by v . From completeness of the Euclidean space, the sequence has a limit point $z \leq v$. This proves the welfare-based update rule is asymptotically stable with respect to F^2 . See Figure 4 for an empirical example of this behavior. $\square$

By assuming the revenue curve is strictly concave and Lipschitz continuous we can derive a condition for the stability of the UDPP mechanism (and EIP-1559).

Theorem 2. *Let F be a distribution over bidder valuations where values are bounded by $\bar{a}$ and F induces the revenue curve $\text{REV}(q)$ to be L -Lipschitz and strictly concave. Then if $\alpha \leq \frac{1}{L(1+\delta)/m+1}$, the UDPP mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$ is asymptotically stable with respect to F and has unique equilibrium price.*

The proof will require the same techniques of the proof for the stability of the TWDPP mechanism. We defer the proof to Section 7.1.

7 The Truncated Welfare-Based Dynamic Posted-Price Mechanism

In this section, we propose the *truncated welfare-based dynamic posted-price* (TWDPP) mechanism that combines desirable features of the WDPP and the UDPP mechanisms. In particular, whenever m or more bidders receive a slot, TWDPP behaves like the UDPP mechanism. Whenever fewer than m bidders receive a slot, we use bid values to compute the subsequent price just like the WDPP mechanism, but to reduce the variance of the update rule, we cap how much each bid contributes to the subsequent price.

¹²Leonardos et al. [26] provides additional examples of unstable behavior for the UDPP mechanism.

Definition 7.1 (Truncated Welfare-Based Update Rule). The *truncated* (α, δ) -welfare-based update rule is the update rule

$$T_{\text{TW}}^{\alpha, \delta}(q, B) := \begin{cases} \alpha \frac{1}{m} \sum_{i \in B} \min\{b_i, (1 + \delta) \cdot q\} + (1 - \alpha) \cdot q & \text{for } |B| < m, \\ \alpha(1 + \delta) \cdot q + (1 - \alpha)q & \text{for } |B| = m, \end{cases} \quad (12)$$

where $\alpha \in (0, 1)$ is the convergence parameter and $\delta \in (0, \infty)$ is the truncation parameter.¹³

Definition 7.2 (Truncated Welfare-Based Dynamic Posted-Price Mechanism). The *truncated welfare-based dynamic posted-price* (TWDPP) mechanism is an RDPP mechanism $(\vec{x}, \vec{p}, T_{\text{TW}}^{\alpha, \delta})$ endowed with the truncated welfare-based update rule $T_{\text{TW}}^{\alpha, \delta}$.

Because the TWDPP mechanism is an RDPP mechanism, Proposition 6.3 implies:

Theorem 3. *The TWDPP mechanism is ex post IC for myopic bidders and DSIC for myopic miners.*

Recall from Proposition 6.5 that there is a distribution where the UDPP mechanism is unstable, but the WDPP mechanism is stable. Since the TWDPP mechanism behaves like the UDPP mechanism when blocks are full, the TWDPP mechanism can be unstable when the WDPP mechanism is stable (simply consider the same distribution constructed in Proposition 6.5). However, the following proposition shows the TWDPP mechanism can still provide advantages over the UDPP mechanism by constructing a distribution F^n where the UDPP mechanism is unstable even when the demand n is smaller than the supply m , but the TWDPP mechanism is stable.

Proposition 7.3. *There exists a distribution F^n with $n < m$ where the TWDPP mechanism $(\vec{x}, \vec{p}, T_{\text{TW}}^{\alpha, \delta})$ is asymptotically stable, but the UDPP $(\vec{x}, \vec{p}, T_{\text{U}}^{\alpha, \delta})$ mechanism is unstable.*

We defer the proof to the appendix. To study the stability of the TWDPP mechanism, we define the expected value of the random operator $T_{\text{TW}}^{\alpha, \delta}(q, \text{Unif}(2^{M(q)}))$:

$$E_{T_{\text{TW}}^{\alpha, \delta}}(q) := \alpha K_{\text{TW}}(q) + (1 - \alpha)q, \quad (13)$$

where K_{TW} is the kernel of the TWDPP mechanism:

$$K_{\text{TW}}(q) := \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^n \min\{v_i, (1 + \delta)q\} \cdot x_i(q) \cdot \mathbb{1}_{N(q) < m} \right. \quad (14)$$

$$\left. + (1 + \delta) \cdot q \cdot \mathbb{1}_{N(q) \geq m} \right], \quad (15)$$

where the expected value is taken over the values $v_1, v_2, \dots, v_n$ and random allocation $\text{Unif}(2^{M(q)})$. Recall $x_i(q) = 1$ if bidder i with $v_i \geq q$ receives a slot and $N(q) = \sum_{i=1}^n \mathbb{1}_{v_i \geq q}$.

¹³An alternative update rule could always round b_i to the interval $[(1 - \delta)q, (1 + \delta)q]$ so that $T(q, B) = \alpha \frac{1}{m} \sum_{i \in B} \max\{\min\{b_i, (1 + \delta)q\}, (1 - \delta)q\} + (1 - \alpha) \cdot q$. Here we study the update rule in Equation 12 to exclude out-of-model deviations where miners prefer to include higher value bids whenever $|B| = m$.

7.1 Stability

Because the kernel K_{TW} is neither increasing nor decreasing, we cannot directly apply the proof sketch in Section 4 to show the TWDPP mechanism is asymptotically stable. Instead, we will show that fixed-point iteration on $E_{\text{TW}}^{\alpha, \delta}$ starting from a positive q_0 converges to a fixed point under the assumption that K_{TW} is Lipschitz continuous and strictly concave.

Definition 7.4 (Concave function). A real-valued function $f : X \rightarrow \mathbb{R}$ is *concave* if for any $x \neq y \in X$ and $0 < \alpha < 1$, $f((1 - \alpha)x + \alpha y) \geq (1 - \alpha)f(x) + \alpha f(y)$, and *strictly concave* if we have a strict inequality.

Here, we show that for a function f satisfying Assumption 1, the fixed-point iteration on $g(x) = \alpha f(x) + (1 - \alpha)x$ (for some positive $\alpha < 1$) starting from a positive x_0 converges to a fixed point of f .¹⁴

Assumption 1. Suppose $X = [0, \bar{a}] \neq \emptyset$ is a finite interval and $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a real-valued function satisfying:

- (i) f is L -Lipschitz.
- (ii) The restriction of f to X is strictly concave.
- (iii) $f(x) = 0$ for all $x \geq \bar{a}$.

Assumption 2. Assume $f : X \rightarrow X$ with nonempty interval $X = [0, \bar{a}]$ is a real-valued function satisfying:

- (i) f is strictly concave and Lipschitz continuous.
- (ii) $f(x) \geq 0$ for all $x \in X$.
- (iii) Let $\bar{x} = \sup_{x \in X} f(x)$. Strict concavity implies f is strictly increasing on interval $I = [0, \bar{x})$ and strictly decreasing on interval $D = [\bar{x}, \sup X]$. Then assume the restriction of f to D is L -Lipschitz with $L < 1$.
- (iv) There is a positive $b \in X$ such that $f(b) < b$.

The main idea is to reduce the problem of computing a fixed point of f satisfying Assumption 1 to finding a fixed point for another function g satisfying Assumption 2 such that x is a fixed point of f if and only if x is a fixed point of g . The main observation is that whenever $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is nonincreasing on some interval $D \subseteq \mathbb{R}_{\geq 0}$, we can force the restriction of g to D to be Lipschitz continuous with Lipschitz constant $L < 1$ by using a similar argument from Lemma 4.5. We can leverage this fact to show that the restriction of g to some interval $[a, b] \subseteq D$ is a contraction— g maps $[a, b]$ to itself and the restriction of g to $[a, b]$ has Lipschitz constant $L < 1$. Then, we show that iterating g starting from any positive

¹⁴It is known that iterating an increasing function $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ will either converge to a fixed point or diverge [22]. In Section 4, we showed that if f is decreasing and Lipschitz continuous, then the fixed-point iteration of $g(x) = \alpha f(x) + (1 - \alpha)x$ will converge to a unique fixed point by setting a sufficiently small positive α . Here, we replace the monotonicity assumption by the assumption f is strictly concave.

x_0 will either converge to a point outside $[a, b]$ or eventually visit a point $x_t \in [a, b]$. The moment $x_t \in [a, b]$ for some t , from Banach contraction principle and the fact the restriction of g to $[a, b]$ is a contractive mapping implies the fixed-point iteration of g starting from x_t will converge to a unique fixed point contained in $[a, b]$.

Lemma 7.5. *If $f : X \rightarrow X$ satisfies Assumption 2, then the fixed-point iteration of f starting from any $x_0 \in X$ converges to a fixed point of f .*

Lemma 7.6. *Let $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a function satisfying Assumption 1. By setting $0 \leq \alpha \leq \frac{1}{L+1}$, the function $g(x) = \alpha f(x) + (1 - \alpha)x$ satisfies Assumption 2. Then, the fixed-point iteration of g starting from any positive x_0 converges to a fixed point of f .*

We defer the proof of the Lemmas to Appendix D.1. In Example D.13, we run simulations demonstrating the fixed-point iteration described in Lemma 7.6.

Lemma 7.6 immediately implies that if the kernel K_{TW} is Lipschitz continuous and strictly concave on a bounded interval $[0, \bar{a}]$, the TWDPP mechanism is asymptotically stable. We can now state our main theorem.

Theorem 4. *Let F be a distribution over bidder valuations where all buyers have value at most $\bar{a}$. Assume F induces the kernel K_{TW} to be L -Lipschitz and strictly concave on interval $[0, \bar{a}]$. Then for any $\alpha \leq \frac{1}{L+1}$, the TWDPP mechanism is asymptotically stable with respect to F .*

Proof. We first check that K_{TW} satisfies all properties in Assumption 1. The first and second properties are clear. For the third property, note that $K_{\text{TW}}(q) \geq 0$, and observe that no bidder purchases at price $> \bar{a}$. Thus $1 - F(q) = 0$ and $K_{\text{TW}}(q) = 0$ for all $q \geq \bar{a}$. This proves K_{TW} satisfies Assumption 1. From Lemma 7.6, for any $\alpha \leq \frac{1}{L+1}$, the fixed-point iteration of

$$E_{T_{\text{TW}}^{\alpha, \delta}}(q_0) = \alpha K_{\text{TW}}(q_0) + (1 - \alpha)q_0$$

starting from positive q_0 converges to a fixed point q^* . This proves the intended TWDPP mechanism is asymptotically stable. $\square$

The proof of the stability of the UDPP mechanism is identical to the proof of the stability of the TWDPP mechanism.

Proof of Theorem 2. Note the expected value of the price update rule $T_{\text{U}}^{\alpha, \delta}(q, B)$ is

$$E_{T_{\text{U}}^{\alpha, \delta}}(q) = \alpha \frac{\text{REV}(q)(1 + \delta)}{m} + (1 - \alpha)q.$$

Observe $\text{REV}(q)(1 + \delta)/m$ is $\frac{L(1+\delta)}{m}$ -Lipschitz and satisfies all conditions in Assumption 1. From Lemma 7.6, $E_{T_{\text{U}}^{\alpha, \delta}}(q)$ is asymptotically stable whenever $\alpha \leq \frac{1}{L(1+\delta)/m+1}$. $\square$

7.2 Welfare guarantee at equilibrium

Because the active miner selects a random maximal allocation whenever $N(q) > m$, it is more challenging to show the TWDPP mechanism is approximately welfare optimal. That is because the social welfare is not a monotone function of the posted price. As the posted price decreases, the welfare might decrease because the active miner has a higher chance to allocate block space to a lower bid instead of a higher bid.

The main idea is to show that if q is an equilibrium price, the probability that $N(q) \geq m$ is at most $1/(1 + \delta)$. Whenever $N(q) \leq m$, miners deterministically allocate slots to the top $N(q)$ bidders. Thus we will show the welfare loss from the cases where $N(q) > m$ is a constant fraction of the optimal welfare.

Theorem 5. *If q is an equilibrium price of the TWDPP mechanism $(\vec{x}, \vec{p}, T_{TW}^{\alpha, \delta})$, then*

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{OPT}{2(1 + \delta)} \min\{1, \delta\}.$$

We defer the proof to the appendix. By setting $\delta = 1$, we get that the TWDPP mechanism obtains $1/4$ of the optimal welfare at equilibrium.

8 Experimental Results

In this section, we describe some experimental results supporting our theoretical findings. We defer additional experiments to Appendix A. We implement an agent-based simulator (included in the supplemental materials) with which we simulate the following dynamic posted-price mechanisms:

- *Welfare-Based Dynamic Posted-Price* (WDPP) mechanism $(\vec{x}, \vec{p}, T_W^\alpha)$: the deterministic dynamic posted-price mechanism with the maximum value allocation rule and the welfare-based update rule.
- *Utilization-Based Dynamic Posted-Price* (UDPP) mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$: the randomized dynamic posted-price mechanism with the random maximal allocation rule and the utilization-based update rule.
- *Truncated Welfare-based Dynamic Posted-Price* (TWDPP) mechanism $(\vec{x}, \vec{p}, T_{TW}^{\alpha, \delta})$: the randomized dynamic posted-price mechanism with the random maximal allocation rule and the truncated welfare-based update rule.

While UDPP and TWDPP mechanisms are *ex post* IC, the WDPP mechanism is not. However, for the simulations, we always assume truthful bidding (in this sense, the results are overly optimistic for WDPP). Recall UDPP mechanism is equivalent to the EIP-1559 protocol in the case that a bidder i bids $b_i = 0$ and submits a bid cap $c_i = v_i$.

For the experimental setup, at each time step t , n_t bidders are drawn i.i.d. from distribution F_t . We use the same random seed across all mechanisms. We set the convergence

parameter $\alpha = 1/16$. We set the truncation parameter $\delta = 1$. Decreasing δ would reduce how much each bid contributes to increasing the subsequent price.

Throughout the simulations, we track three quantities for each time step t : the demand n_t , the posted price q_t , and the ratio between the achieved welfare $\sum_{i \in B_t} v_i$ and the optimal welfare $\sum_{i=1}^{n_t} v_i \cdot x_i^*(\vec{v}_t)$, where $\vec{v}_t$ is drawn from $F_t^{n_t}$, the n_t -dimensional product distribution on F_t .

8.1 Excess demand

As an illustration, consider the case with constant demand $n_t = 200$ above the supply of $m = 100$ and three different value distributions: a uniform distribution over 0 to 200, an exponential distribution with a mean of 100, and a scaled and shifted Pareto distribution with a median around 100. See Figure 1.

The WDPP mechanism has the fastest convergence for distributions with low variance, but is highly volatile when the distribution has high variance (e.g., a Pareto distribution). This kind of high price volatility in WDPP would also raise doubts about the assumption of bidder myopia since they could benefit by timing their bids. Because the UDPP mechanism ignores bid values, it converges to a higher average price than the TWDPP mechanism, with the effect of providing lower welfare. This phenomenon can be observed clearly for the exponential distribution. The TWDPP mechanism more robustly handles both the low- and high-variance cases, consistently achieving the best welfare of around 4/5 of the optimal welfare (outperforming the worst-case 1/4 bound).

9 Conclusion

First-price auctions have been the standard transaction fee mechanism for blockchains since Bitcoin’s inception. While incentive-aligned for miners, first-price auctions provide a bad user experience for users under high competition when it is difficult to estimate how much to bid to get a transaction accepted. Moreover, typical alternatives from mechanism design theory do not work in the blockchain setting because a miner may deviate from the intended auction allocation rules (with the effect of also hiding information from other participants).

The question, then, is how to design a transaction fee mechanism that provides a good user experience and is at the same time robust to manipulation by miners. The Ethereum community has proposed and recently adopted EIP-1559, which is a posted-price and first-price hybrid. EIP-1559 reduces to the incentive-compatible UDPP mechanism that we study when the demand at the current posted price is lower than the supply, but it is not incentive compatible otherwise.

We have proposed the *Truncated Welfare-based Dynamic Posted-Price (TWDDP) mechanism*. This mechanism is incentive compatible for myopic bidders and aligns incentives for myopic miners to follow the rules, even during periods of market instability. It converges to a stable price under strict concavity and Lipschitz continuity assumptions on a quantity derived from the distribution of bidder valuations. At this stable price, the mechanism is also approximately welfare optimal. Our experimental results, conducted in simulation, confirm our theoretical findings.




Figure 1: Simulations with constant demand $n_t = 200$ for all t , constant supply $m = 100$, and uniform, exponential, and Pareto value distributions.

9.1 Future work

This line of work leaves many interesting open questions. As with previous work, we study transaction fee mechanisms under the assumption that miners are myopic. However, as blockchain mining becomes more centralized through mining pools, miners may become motivated to take on non-myopic strategies. Thus it is interesting to understand the strategic behavior of non-myopic miners, seeking bounds on gains available to miners for deviating from the intended allocation rule as well as rules that are robust even to forward-looking miners.

Another interesting direction is to consider the effect of new transaction fee mechanisms on miner incentives to follow the consensus protocol. Block withholding, for example, where a miner hides one or more blocks in an attempt to fork honest blocks from the blockchain, is a forward-looking strategy that can be effective for miners who are motivated by block rewards [13, 38, 16] or transaction fees [5]. The redesign of the transaction fee mechanism modifies the incentives of miners, and leaves the interesting open question of understanding the impact of this change on block withholding strategies.

Lastly, while we model single-dimensional auctions in which bidders have distinct private values but identical transaction sizes, the transaction fee auctions that take place within Ethereum more closely resemble *knapsack auctions*, where different transactions can occupy different amounts of block space. Thus, a relevant question for transaction fee markets is that of the design of credible and communication-efficient mechanisms for multi-dimensional settings such as knapsack auctions [1, 15, 9].

References

- [1] Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020.
- [2] Soumya Basu, David Easley, Maureen O’Hara, and Emin Sirer. Towards a functional fee market for cryptocurrencies. *Available at SSRN 3318327*, 2019.
- [3] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78:771–789, 2010.
- [4] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, and Ian Norden. Ethereum improvement proposal 1559: Fee market change for eth 1.0 chain. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md>.
- [5] Miles Carlsten, Harry Kalodner, S Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167, 2016.
- [6] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Efficient mechanisms with dynamic populations and dynamic types. Technical report, Harvard University, 2010.

- [7] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.
- [8] Lothar Collatz. *Functional analysis and numerical mathematics*. Academic Press, 2014.
- [9] Constantinos Daskalakis, Maxwell Fishelson, Brendan Lucier, Vasilis Syrgkanis, and Santhoshini Velusamy. Simple, credible, and approximately-optimal auctions. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 713–713, 2020.
- [10] Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43(1):192–198, 2007.
- [11] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.
- [12] Meryem Essaidi, Matheus V.X. Ferreira, and S. Matthew Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. In *Forthcoming in Innovations in Theoretical Computer Science*, ITCS ’22, New York, NY, USA, 2022. Association for Computing Machinery.
- [13] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security*, pages 436–454. Springer, 2014.
- [14] Joan Feigenbaum and Scott Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Current Trends in Theoretical Computer Science: The Challenge of the New Century Vol 1: Algorithms and Complexity Vol 2: Formal Models and Semantics*, pages 403–434. World Scientific, 2004.
- [15] Matheus VX Ferreira and S Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 683–712, 2020.
- [16] Matheus VX Ferreira and S Matthew Weinberg. Proof-of-stake mining games with perfect randomness. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 433–453, 2021.
- [17] Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [18] Emily Graffeo. Ether mining revenue and transaction volume soared to record highs in may, report says. URL <https://markets.businessinsider.com/news/currencies/ether-mining-revenue-eth-transaction-uncle-block-subsidy-fee-2021-6>.
- [19] Mohammad T Hajiaghayi. Online auctions with re-usable goods. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 165–174, 2005.

- [20] Gerd Herzog and Peer Chr Kunstmann. A fixed point theorem for decreasing functions. *Numerical Functional Analysis and Optimization*, 34(5):530–538, 2013.
- [21] Gur Huberman, Jacob Leshno, and Ciamac C Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *Bank of Finland Research Discussion Paper*, (27), 2017.
- [22] John Kennan. Uniqueness of positive fixed points for increasing concave functions on $\mathbb{R}^n$: An elementary result. *Review of Economic Dynamics*, 4(4):893–899, 2001.
- [23] Robert Kleinberg and S Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.
- [24] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, 310(1-3):159–180, 2004.
- [25] Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin’s fee market. In *The World Wide Web Conference*, pages 2950–2956, 2019.
- [26] Stefanos Leonardos, Barnabé Monnot, Daniël Reijsbergen, Stratis Skoulakis, and Georgios Piliouras. Dynamical analysis of the eip-1559 ethereum fee market. *arXiv preprint arXiv:2102.10567*, 2021.
- [27] Dov Monderer and Moshe Tennenholtz. Distributed games: From mechanisms to protocols. *AAAI/IAAI*, 99:32–37, 1999.
- [28] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [29] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009. URL <http://www.bitcoin.org/bitcoin.pdf>.
- [30] Stephen O’Neal. Ethereum miners vote to increase gas limit, causing community debate, Jun 2020. URL <https://cointelegraph.com/news/ethereum-miners-vote-to-increase-gas-limit-causing-community-debate>.
- [31] David C Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, pages 411–439. Cambridge University Press, 2007.
- [32] David C. Parkes and Jeffrey Shneidman. Distributed implementations of vickrey-clarke-groves mechanisms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS ’04, page 261–268, USA, 2004. IEEE Computer Society. ISBN 1581138644.
- [33] David C. Parkes and Satinder Singh. An mdp-based approach to online mechanism design. In *Proceedings of the 16th International Conference on Neural Information Processing Systems*, NIPS’03, page 791–798, Cambridge, MA, USA, 2003. MIT Press.

- [34] Pooja Ranjan and Tim Beiko. Eip-1559 community outreach report. <https://medium.com/ethereum-cat-herders/eip-1559-community-outreach-report-aa18be0666b5/>, 2020. Accessed: 2021-02-11.
- [35] Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7): 78–86, 2010.
- [36] Tim Roughgarden. Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. *arXiv preprint arXiv:2012.00854*, 2020.
- [37] Halsey Lawrence Royden and Patrick Fitzpatrick. *Real analysis*, volume 32. Macmillan New York, 1988.
- [38] Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 515–532. Springer, 2016.
- [39] Jeffrey Shneidman and David C Parkes. Specification faithfulness in networks with rational nodes. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 88–97, 2004.
- [40] Andrew Chi-Chih Yao. An incentive analysis of some bitcoin fee designs. In *47th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 1:1–1:12, July 2020.

A Other Experimental Results

A.1 Under demand

Here we replicate the experiments from Section 8.1 but with demand $n_t = 67$ smaller than the supply $m = 100$. When the demand is lower than the supply, the market-clearing price is 0. In general, we cannot expect our mechanisms to converge to a posted price of 0, but we can hope the mechanisms converge to a price that sells to all almost all bidders. We give the results in Figure 2.

Unfortunately, the WDPP mechanism continues to have high price volatility, which contributes to low welfare on average. Both the UDPP and the TWDPP mechanisms sell to almost all bidders, with the TWDPP having a slightly higher welfare because it converges to a lower average price.




Figure 2: Simulations with constant demand $n_t = 67$ for all t , constant supply $m = 100$, and uniform, exponential, and Pareto value distributions.

A.2 Demand shocks

We repeat the experiments from Section 8.1 with a demand shock: the demand starts at 200 transactions per time step, jumps to 600, and later returns to 200. The results are given in Figure 3.

We observe that all mechanisms converge quickly to an equilibrium price under demand shock. During demand shock, the WDPP mechanism’s instability increases for distributions with high variance (e.g., the Pareto distribution). The TWDPP mechanism has strictly higher welfare for all value distributions because it converges to an average price closer to the market-clearing price.

We observe that the TWDPP mechanism can have price volatility when values are drawn from the uniform distribution over $[0, 200]$. That happens because whenever a block is full, the TWDPP mechanism uses the utilization-based update rule. Thus when the TWDPP mechanism reaches a price q close to 175, there is a high probability that there are at least m bidders with values above q causing the subsequent increase to $q(1 + \alpha)$. However, this instability is unpredictable (differently from the instability we will observe in Section A.3) and even under instability, the TWDPP mechanism extracts higher welfare than the UDPP mechanism.




Figure 3: The demand follows a step function that starts at 200, jumps to 600, and finally returns to 200. As with the previous simulations, we use a fixed supply of $m = 100$ slots per time step, and we provide results for uniform, exponential, and Pareto value distributions.

A.3 Instability of EIP-1559 utilization-based update rule

This section considers the case where all bidders have a deterministic value of 100. The experiment in Figure 4 serves as empirical support for Proposition 6.5. It provides an example of a distribution where the WDPP mechanism is asymptotically stable, but the utilization-based update rule from EIP-1559 (and hence the UDPP mechanism) is unstable.

From the experiment, we observe the WDPP mechanism converges to a stable price of 100 while the UDPP mechanism oscillates around 100. The price oscillation weakens the bidder's myopic assumption. In particular, a patient bidder would wait to enter a bid only when the price is below 100. For this example, one can check the TWDPP mechanism is also unstable because the truncated welfare-based update rule reduces to the welfare-based

update rule.




Figure 4: Simulation with constant demand $n = 200$, constant supply $m = 100$, and deterministic bidder value $v = 100$.

Next, we modify the example in Figure 4 and consider the case where the demand is smaller than the supply. The experiment in Figure 5 serves as empirical support for Proposition 7.3. It provides an example of a distribution where the WDPP and TWDP mechanisms are asymptotically stable, but the UDPP mechanism is unstable even when demand is lower than the supply.




Figure 5: Simulation with constant demand $m = 67$, constant supply $m = 100$, and deterministic bidder value $v = 100$.

B Existing Proposals of Blockchain Transaction Fee Mechanisms

B.1 EIP-1559

First-price auctions (Definition 2.12) have been the standard choice for transaction fee mechanisms since Nakamoto [29] first proposed using them for Bitcoin. Buterin et al. [4]’s EIP-1559 proposal for Ethereum is closest to ours since their mechanism is endowed with an update rule for their posted price. However, their proposal contains a first-price auction component and a posted-price component, and here we study purely posted-price mechanisms.

Definition B.1 (EIP-1559 [4]). The mechanism proposed in EIP-1559 is a dynamic mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$ endowed with the utilization-based update rule $T_U^{\alpha, \delta}$ (Definition 6.1) and a initial posted price q_0 . Each bidder submits a tuple (b_i, c_i) where b_i is the bid and c_i is a bid cap. Given history $B_1, B_2, \dots, B_{t-1}$ and active bidders M_t , the mechanism computes the posted price $q_t = T(q_{t-1}, B_{t-1})$ for time step t . The intended allocation rule allocates slots to the top m bids b_i with $b_i + q_t \leq c_i$. For all bidders that receive a slot, the intended payment rule charges $b_i + q_t$, but miners are paid only b_i . Thus the utility of the active miner is

$$u_0^{\text{EIP-1559}}(q_t, B_t) := \sum_{i \in B_t} (p_i(B_1, \dots, B_t) - q_t).$$

For EIP-1559, Buterin et al. [4] propose setting $\delta = 1$ in the hope the mechanism converges to an equilibrium that sells $m/2$ of the slots. Formally studying the stability of EIP-1559 is more challenging since truthful bidding is not an *ex post* Nash equilibrium, and any formal analysis would likely require analyzing bidders participating in a Bayesian Nash equilibrium. An alternative approach is to study EIP-1559 without the first-price auction component. The resulting mechanism is equivalent to the UDPP mechanism we study in Section 6.

B.2 Monopolistic price auction

We study mechanisms that sell a fixed supply of block slots. Nevertheless, since blockchain space is a digital good [17], it is not obvious why the supply must be fixed a priori. With that in mind, Lavi et al. [25] propose using the *monopolistic price auction* as a transaction fee mechanism, allowing miners to sell an unlimited supply of block slots. They showed the monopolistic price auction is DSIC for myopic miners and approximately IC for myopic bidders when values are drawn i.i.d. from a bounded distribution $F = [\underline{a}, \bar{a}]$ and demand n gets large – i.e., $n \rightarrow \infty$.

Definition B.2 (Monopolistic price auction [25]). The *monopolistic price auction* is a static mechanism where the allocation rule allows an unlimited supply of slots to be allocated. First, the mechanism orders the active bids in decreasing order $b_1 \geq b_2 \geq \dots \geq b_n$ and lets

$$k = \arg \max_i i \cdot b_i.$$

Then, it allocates slots to bidders 1 through k . All bidders that receive a slot pay b_k to the active miner.

The monopolistic price is DSIC for myopic miners, and Yao [40] showed it is also approximately IC for myopic bidders for any distribution F when the demand $n \rightarrow \infty$. Unfortunately, there is an inherent trade-off between having larger blocks and loss in social welfare.¹⁵ Large block sizes result in higher network propagation delay, which increases the chance of forks. As a result, an unlimited supply of slots brings with it the possibility of new denial-of-service attacks.

¹⁵Ethereum allows the blockchain protocol to converge to a fixed supply m that is agreed upon by the majority of miners by allowing each active miner to upvote or downvote the supply by $\approx 0.1\%$. Thus one might expect m to converge to a block size that balances the trade-offs between latency and supply.

B.3 Random sampling optimal price

Goldberg et al. [17] proposed the *random sampling optimal price* (RSOP) mechanism as a randomized truthful mechanism for selling an unlimited supply of items. The RSOP mechanism randomly partitions all bids into two groups and runs the monopolistic price auction in each. Lavi et al. [25] propose using RSOP as a transaction fee mechanism but recognize that it has undesirable properties. Crucially, RSOP has communication complexity $\Omega(n)$ since it requires all bids to be stored on the blockchain.

B.4 Generalized second-price auction

The *generalized second-price* (GSP) auction [11] orders the bids in decreasing order $b_1 \geq b_2 \geq \dots \geq b_n$ and allocates slots to the top m bidders with each bidder $i \in [m]$ paying b_{i+1} (assuming $n \geq m + 1$). In the context of transaction fee mechanisms, Basu et al. [2] proposed a modified generalized second-price auction where bidder $i \in [m]$ pays b_m . Note their payment rule is identical to the monopolistic price auction but with a fixed supply. To incentivize miners to not include fake bids and inflate the m -th highest bid, the mechanism splits the revenue for block B_t among the creators of future blocks $B_{t+1}, B_{t+2}, \dots, B_{t+k}$. As the demand and the number of miners increase, the mechanism is approximately incentive compatible for myopic bidders and myopic miners.

An out-of-model consideration pointed out by Lavi et al. [25] is collusion between an active miner with bidders where both agree to execute payments outside the blockchain. For example, a bidder bids $b_i = 0$ but agrees to pay the intended bid outside the blockchain.¹⁶ To motivate bidders to accept the collusion, the miner might divide the revenue with bidders. Moreover, even without colluding with bidders, a miner could create blocks with m fake bids equal to 0 and share the revenue of blocks $B_{t-k}, B_{t-k+1}, \dots, B_{t-1}$ for free. The active miner loses $1/k$ of the revenue they could obtain by honestly creating block B_t . Thus including only fake transactions in blocks is an approximate equilibrium.

C Omitted Proofs from Section 5

Proof of Theorem 1. Recall $x_i(q)$ is the indicator random variable for bidder $i \in [n] = \{1, \dots, n\}$ with value v_i receiving a slot when the posted price is q . For fixed $\vec{v}$, $x_i(q)$ is a nonincreasing function of q because the mechanism always allocates slots to the highest bidders. Assuming truthful bidding, the expected value of $T_W^\alpha(q, B)$ is

$$E_{T_W^\alpha}(q) := \alpha \underbrace{\frac{1}{m} \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \right]}_{\text{Kernel } K_W(q)} + (1 - \alpha)q, \quad (16)$$

where the expectation is taken over $\vec{v}$ drawn from F^n . From the proof outline in Section 4, we first check $E_{T_W^\alpha}$ is a monotone mixture. Note that $K_W(q) = \text{WELFARE}(\vec{x}(q))/m$ and that $\text{WELFARE}(\vec{x}(q))$ is a nonincreasing function of q since $x_i(q) = 1$ for the top m bids

¹⁶For simplicity, we consider the case where the auction has no reserve price.

with $v_i \geq q$. Thus $E_{T_W^\alpha}$ is a monotone mixture. Next, we prove the kernel K_W is Lipschitz continuous under the assumption REV is Lipschitz continuous.

Claim C.1. *If REV is L -Lipschitz, then kernel K_W is $(L/m + 1)$ -Lipschitz.*

Proof. We decompose the welfare as the revenue plus the surplus. Since the revenue is L -Lipschitz, it suffices to show the surplus is m -Lipschitz. Fix any positive $q < q'$ and because the welfare is decreasing, we have $K_W(q) \geq K_W(q')$. Thus,

$$\begin{aligned} 0 &\leq m \cdot (K_W(q) - K_W(q')) \\ &= \underbrace{\mathbb{E} \left[\sum_{i=1}^n (v_i - q) \cdot x_i(q) \right]}_{\text{Expected surplus at price } q} - \underbrace{\mathbb{E} \left[\sum_{i=1}^n (v_i - q') \cdot x_i(q') \right]}_{\text{Expected surplus at price } q'} + \text{REV}(q) - \text{REV}(q') \\ &= \mathbb{E} \left[\sum_{i=1}^n (v_i - q) \cdot (x_i(q) - x_i(q')) - (v_i - q' - v_i + q) \cdot x_i(q') \right] + \text{REV}(q) - \text{REV}(q'). \end{aligned}$$

Observe $x_i(q') \leq x_i(q)$, and whenever $x_i(q) - x_i(q') = 1$, bidder i has value $v_i \leq q'$. Thus the difference between the expected surplus at price q and q' is at most

$$(q' - q) \mathbb{E} \left[\sum_{i=1}^n x_i(q) - x_i(q') \right] + (q' - q) \mathbb{E} \left[\sum_{i=1}^n x_i(q') \right] = (q' - q) \mathbb{E} \left[\sum_{i=1}^n x_i(q) \right] \leq (q' - q) \cdot m.$$

From the assumption REV is L -Lipschitz, we have

$$0 \leq K_W(q) - K_W(q') \leq (q' - q)(L/m + 1).$$

This proves K_W is $(L/m + 1)$ -Lipschitz. $\square$

From Lemma 4.5 with $\alpha \leq 1/(L/m + 1)$, we get that $E_{T_W^\alpha}(q)$ is a contractive mapping with Lipschitz constant $1 - \alpha < 1$. From the Banach Contraction Principle (Lemma 4.6), for any initial positive price q_0 , the iteration $q_0, E_{T_W^\alpha}(q_0), E_{T_W^\alpha}^2(q_0), \dots$ converges to a unique fixed point q^* . From Observation 4.4, q^* is also a unique fixed point for the kernel K_W . This proves the intended mechanism is asymptotically stable with respect to distribution F^n .

Claim C.2. *If q is a fixed point for K_W , then $\text{WELFARE}(\vec{x}(q)) \geq \frac{1}{2} \text{OPT}$.*

Proof. Observe $q \cdot m$ is the expected welfare of the WDPP mechanism at equilibrium price q since q is a fixed point for K_W . That is,

$$\text{WELFARE}(\vec{x}(q)) = \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \right] = m \cdot K_W(q) = q \cdot m,$$

where the second equality is simply the definition of K_W and the third equality follows from the fact $q = K_W(q)$ is a fixed point of K_W . Assume for contradiction $q \cdot m \leq \frac{1}{2} \text{OPT}$, then the optimal welfare

$$\text{OPT} = \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \right] = \mathbb{E} \left[\sum_{i=1}^n v_i (x_i^*(\vec{v}) - x_i(q)) \right] + \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \right]$$

$$\begin{aligned}
&< q\mathbb{E}\left[\sum_{i=1}^n(x_i^*(\vec{v}) - x_i(q))\right] + \text{WELFARE}(\vec{x}(q)) \quad \text{Because } x_i^*(\vec{v}) - x_i(q) = 1 \text{ only if } v_i < q. \\
&\leq \frac{OPT}{2} + \text{WELFARE}(\vec{x}(q)) \quad \text{Because } q \leq \frac{OPT}{2m} \text{ and } \sum_{i=1}^n x_i^*(\vec{v}) - x_i(q) \leq m. \\
&= \frac{OPT}{2} + \text{WELFARE}(\vec{x}(q)).
\end{aligned}$$

Rearranging the inequality gives $q \cdot m = \text{WELFARE}(\vec{x}(q)) > \frac{OPT}{2}$, a contradiction. Thus $\text{WELFARE}(\vec{x}(q)) = q \cdot m \geq \frac{1}{2}OPT$. $\square$

This proves Theorem 1. $\square$

D Omitted Proofs from Section 7

Proof of Proposition 7.3. For any δ, α , positive value v , and supply $m \geq 4(1 + \delta)/\delta$, assume that at each time step there are $n = \lfloor m \frac{1+\delta/2}{1+\delta} \rfloor \leq m - 1$ bidders with value v . For the UDPP mechanism $(\vec{x}, \vec{p}, T_U^{\alpha, \delta})$, if the current price $q_t \leq v$, the subsequent price increases to

$$q_{t+1} = q_t \left(1 + \alpha \left(\frac{n}{m}(1 + \delta) - 1\right)\right) \geq q_t \left(1 + \alpha \left(1 + \frac{\delta}{2} - \frac{1 + \delta}{m} - 1\right)\right) \geq q_t \left(1 + \frac{\alpha\delta}{4}\right).$$

If the current price $q_t > v$, no bidder purchases, and the subsequent price decreases to

$$q_{t+1} = q_t(1 - \alpha).$$

This repeats *ad infinitum*, and the allocation-based update rule is unstable with respect to distribution F^n .

For the TWDPP mechanism $(\vec{x}, \vec{p}, T_{TW}^{\alpha, \delta})$, if the current price $q_t > v$, no bidder purchases, and the subsequent price decreases to

$$q_{t+1} = q_t(1 - \alpha).$$

If the current price $q_t < \frac{v}{1+\delta}$, all n bidders purchase, and the subsequent price increases to

$$q_{t+1} = \alpha \frac{n}{m} \min\{v, (1 + \delta)q_t\} + (1 - \alpha)q_t \leq q_t(\alpha(1 + \delta) + 1 - \alpha) = q_t(1 + \alpha\delta) < v.$$

Thus eventually, the mechanism reaches a price $q_t \in [\frac{v}{1+\delta}, v] = X$. If the current price $q_t \in X$, all n bidders purchase, and the subsequent price is

$$q_{t+1} = T_{TW}^{\alpha, \delta}(q_t) = \alpha \frac{n}{m} \min\{v, (1 + \delta)q_t\} + (1 - \alpha)q_t = \alpha \frac{n}{m} v + (1 - \alpha)q_t \leq v.$$

Additionally,

$$T_{TW}^{\alpha, \delta}(q_t) = \alpha \frac{n}{m} v + (1 - \alpha)q_t$$

$$\begin{aligned}
&\geq \alpha v \left(\frac{1 + \delta/2}{1 + \delta} - \frac{1}{m} \right) + (1 - \alpha) \frac{v}{1 + \delta} && \text{Because } n \geq m \frac{1 + \delta/2}{1 + \delta} - 1, \\
&= \alpha v \left(\frac{1}{1 + \delta} + \frac{\delta/2}{1 + \delta} - \frac{1}{m} \right) + (1 - \alpha) \frac{v}{1 + \delta} \\
&\geq \alpha v \left(\frac{1}{1 + \delta} + \frac{\delta}{4(1 + \delta)} \right) + (1 - \alpha) \frac{v}{1 + \delta} && \text{Because } m \geq 4(1 + \delta)/\delta, \\
&\geq \frac{v}{1 + \delta}.
\end{aligned}$$

Thus $T_{\text{TW}}^{\alpha, \delta}$ maps X to itself. For all $q, q' \in X$, we have

$$|T_{\text{TW}}^{\alpha, \delta}(q) - T_{\text{TW}}^{\alpha, \delta}(q')| = (1 - \alpha)|q - q'|.$$

Thus the restriction of $T_{\text{TW}}^{\alpha, \delta}$ to X is a contractive mapping (Definition 4.2). From the Banach Contraction Principle (Lemma 4.6), for all $q_t \in X$, the fixed point iteration

$$q_t, T_{\text{TW}}^{\alpha, \delta}(q_t), T_{\text{TW}}^{\alpha, \delta}(T_{\text{TW}}^{\alpha, \delta}(q_t)), \dots$$

converges to a unique fixed point $z \in X$. This proves the TWDPP mechanism is asymptotically stable with respect to distribution F^n . See Figure 5 for an empirical example of this behavior. $\square$

D.1 Omitted proofs from Section 7.1

The following is a well-known fact about concave functions.

Proposition D.1. *If f and g are (strictly) concave functions, then:*

(i) $f + g$ is (strictly) concave.

We will require the following fact:

Lemma D.2 (Theorem 3.1 in [22]). *If f is strictly concave and $f(0) \geq 0$, then f has at most one positive fixed point.*

Lemma D.2 guarantees a concave function f satisfying Assumption 2 has at most one positive fixed point, but it does not guarantee existence. For the one-dimensional case, uniqueness will follow directly from the intermediate value theorem.

Proof of Lemma 7.5. First, we show $f : X \rightarrow X$ with $X = [0, \bar{a}]$ has at least one nonnegative fixed point.

Claim D.3. *There is a positive $a \in X$ such that $f(a) > a$ if and only if f has a unique positive fixed point x^* .*

Proof. We first prove that if there is a positive $a \in X$ such that $f(a) > a$, then f has a unique positive fixed point.

$\Rightarrow$ From (iv), there is a positive $b \in X$ such that $f(b) < b$. Let

$$g(x) = f(x) - x, \tag{17}$$

and observe g is continuous (because f is continuous) and x is a fixed point for f if and only if $g(x) = 0$. Since g is continuous, $g(a) > 0$, and $g(b) < 0$, we can use the intermediate value theorem to conclude there is an $x^* \in [\min\{a, b\}, \max\{a, b\}] \subseteq X$ such that $g(x^*) = 0$. Thus $x^* > 0$ is a positive fixed point for f . From Lemma D.2, x^* is the unique positive fixed point of f .

Next, we prove that if f has a unique positive fixed point, then there is a positive $a \in X$ such that $f(a) > a$.

$\Leftarrow$ Let $z \in X$ be a positive fixed point of f . For some $\alpha \in (0, 1)$, let $a = \alpha \cdot z > 0$ and note $a \in X$. From the strict concavity of f ,

$$\begin{aligned} f(\alpha z) &= f(\alpha z + (1 - \alpha) \cdot 0) \\ &> \alpha f(z) + (1 - \alpha)f(0) && \text{From strict concavity of } f, \\ &\geq \alpha f(z) && \text{From (ii), } f(0) \geq 0, \\ &= \alpha z && \text{Because } z = f(z) \text{ is a fixed point of } f. \end{aligned}$$

The chain of inequalities witnesses $f(a) = f(\alpha z) > \alpha \cdot z = a > 0$. Thus there is a positive $a \in X$ such that $f(a) > a$. $\square$

Claim D.4. *f has at least one nonnegative fixed point and at most one positive fixed point.*

Proof. Let's first consider the case where f has no positive fixed point. From Claim D.3, if f has no positive fixed point, then $0 \leq f(x) \leq x$ for all positive $x \in X$. Then the sequence

$$f(\bar{a}), f\left(\frac{\bar{a}}{2}\right), f\left(\frac{\bar{a}}{3}\right), \dots$$

converges to 0. By continuity of f , $0 = f(0)$ is a fixed point of f and 0 is a nonnegative fixed point of f whenever f has no positive fixed point. Second, consider the case where f has a positive fixed point. From Lemma D.2, this fixed point is unique. $\square$

Let x^* be the largest nonnegative fixed point of f . Define the partition of X into sets

$$A := \{x \in X : f(x) \geq x\}, \tag{18}$$

$$L := \{x \in X : f(x) < x\}. \tag{19}$$

Claim D.5. *$A = [\inf X, x^*]$ and $L = (x^*, \sup X]$.*

Proof. Observe $x^* \in A$ because $f(x^*) = x^*$. Suppose for contradiction there are $x \in A$ and $y \in L$ such that $x > y$. Because $0 < y < x$, there is an $\alpha \in (0, 1)$ such that $y = (1 - \alpha)x$. From the strict concavity of f ,

$$\begin{aligned} f(y) &= f(0 \cdot \alpha + x(1 - \alpha)) \\ &> \alpha f(0) + (1 - \alpha)f(x) && \text{From strict concavity of } f, \\ &\geq (1 - \alpha)x = y && \text{Because } f(0) \geq 0 \text{ and } f(x) \geq x \text{ since } 0, x \in A. \end{aligned}$$

The chain of inequalities witnesses $y \in A$, a contradiction to $y \in L = X \setminus A$. $\square$

We are ready to show that iterating f starting from a positive x_0 converges to x^* . Define $x_t = f(x_{t-1})$ for $t = 1, 2, \dots$. It suffices to show that the sequence $x_0, x_1, \dots$ converges to a limit point z since the continuity of f implies z is a fixed point of f . When z is positive and f has a unique positive fixed point, we will have $z = x^*$. First, let us consider the case where f has no positive fixed point.

Claim D.6. *If f has no positive fixed point, the sequence $x_0, x_1, \dots$ converges to 0.*

Proof. From Claim D.4, 0 is the unique nonnegative fixed point of f and $L = X \setminus \{0\}$. If $x_0 = 0$, then $f(x_0) = 0$. If $x_0 \in L$, from (19) and (ii), $0 \leq x_1 = f(x_0) < x_0$. By induction, the sequence

$$x_0 > x_1 > x_2 > \dots \geq 0.$$

is decreasing and lower bounded by 0.

Claim D.7. *If the sequence $x_0, x_1, \dots$ is monotone and bounded, the sequence $x_0, x_1, \dots$ converges to a limit point z . Moreover, z is a fixed point of f .*

Proof. Since the Euclidean space is a complete metric space, the sequence $x_0, x_1, \dots$ converges to a limit point z . By continuity of f , z is a fixed point of f . $\square$

Thus when f has no positive fixed point, the nonincreasing sequence $x_0, x_1, \dots$ converges to 0, the unique nonnegative fixed point of f . $\square$

For the case where f has a positive fixed point, from Lemma D.2, f has a unique positive fixed point x^* . Next, we divide the proof into the case where $x^* < \bar{x} = \sup_{x \in X} f(x)$ and the case where $x^* \geq \bar{x}$.

Claim D.8. *If f has a positive fixed point and $x^* < \bar{x}$, the sequence $x_t, x_{t+1}, \dots$ converges to x^* .*

Proof. Observe $[\inf X, x^*] = A \subseteq I = [\inf X, \bar{x})$. We consider separately the case where $x_t \in A$ and the case $x_t \in L$.

Case 1: For the case $x_t \in A$,

$$\begin{aligned} x_t &\leq x_{t+1} = f(x_t) && \text{Because } x_t \in A, \\ &\leq f(x^*) && \text{Because } f \text{ is increasing on } I \ni x_t, x^* \text{ and } x_t \leq x^*, \\ &= x^* && \text{Because } x^* = f(x^*) \text{ is a fixed point.} \end{aligned}$$

Thus $x_{t+1} \in A$. By induction, $x_{t+i} \in A$ for all $i \geq 0$. Thus the sequence

$$x_t \leq x_{t+1} \leq \dots \leq x^*$$

is nondecreasing and upper bounded by x^* . From Claim D.7, the sequence $x_t, x_{t+1}, \dots$ converges to x^* .

Case 2: For the case where $x_t \in L$, we divide the proof into two subcases. First, for the case where $x_{t+i} \in L = (x^*, \sup X]$ for all $i \geq 0$, the sequence

$$x_t > x_{t+1} > x_{t+2} > \dots > x^*$$

is decreasing and lower bounded by x^* . From Claim D.7, the sequence $x_t, x_{t+1}, \dots$ converges to x^* .

Second, for the case where $x_{t+i} \notin L$ for some $i \geq 0$, $x_{t+i} \in A = X \setminus L$. Thus invoke Case 1 to conclude the sequence $x_{t+i}, x_{t+i+1}, x_{t+i+2}, \dots$ converges to x^* . This proves the sequence $x_t, x_{t+1}, \dots$ converges to x^* whenever f has a positive fixed point and $x^* < \bar{x}$. $\square$

Claim D.9. *If f has a positive fixed point and $\bar{x} \leq x^*$, the sequence $x_t, x_{t+1}, \dots$ converges to x^* .*

Proof. We first show the restriction of f to $[\bar{x}, f(\bar{x})]$ is a contractive mapping — i.e., f maps $[\bar{x}, f(\bar{x})]$ to itself and the restriction of f to $[\bar{x}, f(\bar{x})]$ is L -Lipschitz with $L < 1$.

Claim D.10. *The restriction of f to $[\bar{x}, f(\bar{x})]$ is a contractive mapping.*

Proof. From assumption (iii), the restriction of f to the interval $[\bar{x}, f(\bar{x})] \subseteq D$ is L -Lipschitz with $L < 1$. Thus it suffices to show that f maps $[\bar{x}, f(\bar{x})]$ to itself. Let $x \in [\bar{x}, f(\bar{x})]$, then $f(x) \leq f(\bar{x})$ since $\bar{x} \leq x$ and the restriction of f to D is decreasing. Suppose for contradiction $f(x) < \bar{x}$, then

$$\begin{aligned} |f(x) - f(\bar{x})| &= f(\bar{x}) - f(x) \\ &> f(\bar{x}) - \bar{x} && \text{From the assumption } f(x) < \bar{x}, \\ &\geq x - \bar{x} \Rightarrow \Leftarrow && \text{Observe } x \leq f(\bar{x}) \text{ since } x \in [\bar{x}, f(\bar{x})]. \end{aligned}$$

We reach a contradiction to the fact that the restriction of f to D is L -Lipschitz with $L < 1$. Thus $\bar{x} \leq f(x) \leq f(\bar{x})$, proving f maps $[\bar{x}, f(\bar{x})]$ to itself. This proves the restriction of f to $[\bar{x}, f(\bar{x})]$ is a contractive mapping. $\square$

It follows that if $x_t \in [\bar{x}, f(\bar{x})]$, the sequence $x_t, x_{t+1}, \dots$ will converge to a limit point.

Claim D.11. *For any $x_t \in [\bar{x}, f(\bar{x})]$, the sequence $x_t, x_{t+1}, x_{t+1}, \dots$ converges to x^* .*

Proof. The proof follows from the Banach Contraction Principle (Lemma 4.6) and the fact the restriction of f to $[\bar{x}, f(\bar{x})]$ is a contractive map (Claim D.10). $\square$

Recall from Assumption 2, $I = [\inf X, \bar{x})$ is the interval where f is nondecreasing and $D = [\bar{x}, \sup X]$ is the interval f is nonincreasing. Next, we consider separately the case where $x_t \in I$ and the case $x_t \in D$.

Case 1: For the case $x_t \in I$, observe

$$I = [\inf X, \bar{x}) \subset [\inf X, x^*] = A,$$

from the assumption $\bar{x} \leq x^*$. Thus $x_t \in A$. Let's consider two distinct cases depending on whether $x_{t+i} \in A$ for all $i \geq 0$ or $x_{t+i} \notin A$ for some $i \geq 0$. For the case where $x_{t+i} \in A$ for all $i \geq 0$, the sequence

$$x_t \leq x_{t+1} \leq x_{t+2} \leq \dots \leq \sup_{x \in X} f(x) = f(\bar{x})$$

is nondecreasing and upper bounded by $f(\bar{x}) < \infty$ (because f is Lipschitz continuous):

Claim D.12. $\sup_{x \in X} f(x) < \infty$.

Proof. From (iv), there is a $b \in X$ such that $f(b) < b$. Because f is L -Lipschitz,

$$\sup_{x \in X} f(x) - f(b) = |f(\bar{x}) - f(b)| \leq L|\bar{x} - b| \leq L \sup X.$$

Thus $\sup_{x \in X} f(x) \leq L \sup X + f(b) < L \sup X + b < (L + 1) \sup X < \infty$. $\square$

From Claim D.7, the nondecreasing sequence $x_t, x_{t+1}, \dots$ converges to x^* . For the case where $x_{t+i} \notin A$ for some i , let $i \geq 1$ be the first element that satisfies this condition — that is, $x_{t+i} \notin A$ while $x_{t+i-j} \in A$ for all $1 \leq j \leq i$. Observe

$$x_{t+i} = f(x_{t+i-1}) \leq \sup_{x \in X} f(x) = f(\bar{x}),$$

and the fact $x_{t+i} \notin A$ implies $x_{t+i} \in X \setminus A = L = (x^*, \sup X]$. Thus $x_{t+i} > x^* \geq \bar{x}$. Thus

$$\bar{x} \leq x^* < x_{t+i} = f(x_{t+i-1}) \leq f(\bar{x}).$$

The chain of inequalities witnesses $x_{t+i} \in [\bar{x}, f(\bar{x})]$ and from Claim D.11, the sequence $x_{t+i}, x_{t+i+1}, x_{t+i+2}, \dots$ converges to x^* .

Case 2: For the case where $x_t \in D = [\bar{x}, \sup X]$, observe $x_t \geq \bar{x} > 0$. We consider separately the case where $x_{t+i} \in D$ for all $i \geq 0$ and the case where $x_{t+i} \notin D$ for some $i \geq 0$.

For the case where $x_{t+i} \in D$ for all $i \geq 0$, from Banach contraction principle, the sequence $x_t, x_{t+1}, \dots$ converges to $x^* \geq \bar{x}$ since the restriction of f to D is L -Lipschitz with $L < 1$ (iv).

For the case where $x_{t+i} \notin D$ for some $i \geq 0$, $x_{t+i} \in X \setminus D = I$. From Case 1, the sequence $x_{t+i}, x_{t+i+1}, \dots$ converges to x^* . Cases 1 and 2 prove the sequence $x_t, x_{t+1}, \dots$ converges to x^* whenever $\bar{x} \leq x^*$ and f has a positive fixed point. $\square$

Claims D.6, D.8, and D.9 together prove that the sequence $x_0, x_1, \dots$ converges to a fixed point of f . $\square$

Proof of Lemma 7.6. Recall that $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is L -Lipschitz, that the restriction of f to $X = [0, \bar{x}]$ is strictly concave, and that $f(x) = 0$ for all $x \geq \bar{a}$. For positive $\alpha \leq \frac{1}{L+1}$, define

$$g(x) = \alpha f(x) + (1 - \alpha)x.$$

First, we show the restriction of $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ to X maps X to itself and satisfies Assumption 2.

For (i), observe that g is a convex combination of x and $f(x)$. From Proposition D.1, g restricted to X is strictly concave (because f restricted to X is strictly concave).

For (ii), from the assumption f is nonnegative, g is also nonnegative.

For (iii), for all $x \in X$,

$$\begin{aligned} 0 \leq g(x) &= \alpha f(x) + (1 - \alpha)x \\ &= \alpha(f(x) - f(\bar{a})) + (1 - \alpha)x && \text{From (iii), } f(\bar{a}) = 0, \\ &\leq \alpha(\bar{a} - x)L + (1 - \alpha)x && \text{From the fact } f \text{ is } L\text{-Lipschitz,} \\ &\leq \bar{a} - x + x - \alpha x && \text{From the fact } \alpha \leq 1/L, \end{aligned}$$

$$\leq \bar{a}.$$

Thus g maps X to itself. Let $I \subseteq X$ be the interval where f is increasing and let $D \subseteq X$ be the interval where f is decreasing. Then the restriction of g to D is a monotone mixture (Definition 4.3) with kernel f . From Lemma 4.5, the restriction of g to D is $(1 - \alpha)$ -Lipschitz for $\alpha \leq \frac{1}{L+1}$. This proves the restriction of g to X satisfies (iii).

For (iv), let $b = \bar{a} > 0$ and from the assumption $f(b) = 0 < b$,

$$g(b) = \alpha f(b) + (1 - \alpha)b < b.$$

The work above proves the restriction of g to X satisfies Assumption 2. Consider the fixed-point iteration of g starting from a positive x_0 . Let $x_{t+1} = g(x_t)$ for $t \geq 0$. If $x_t \geq \bar{a}$, then

$$g(x_t) = (1 - \alpha)x_t.$$

Let x_t be the first element in the sequence $x_0, x_1, \dots$ such that $x_t \leq \bar{a}$. From Lemma 7.5, the fixed-point iteration of g restricted to X starting from $x_t \in X$ converges to a nonnegative fixed point of g . Finally, observe that a fixed point of g is also a fixed point of f . $\square$

Example D.13. Consider the strictly concave functions $f_1(x) = 4 - (x - 2)^2$ and $f_2(x) = \frac{f(x)}{2}$, and for $\alpha = 0.4$, define functions $g_1(x) = \alpha f_1(x) + (1 - \alpha)x$ and $g_2(x) = \alpha f_2(x) + (1 - \alpha)x$. In Figure 6, we plot the fixed-point iteration of g_1 (respectively g_2), demonstrating that for a function satisfying Assumption 1 (e.g., f_1 (resp. f_2)), we can construct a function satisfying Assumption 2 (e.g., g_1 (resp. g_2)). Then the fixed-point iteration of g_1 (resp. g_2) converges to a fixed point for f_1 (resp. f_2).

D.2 Omitted proofs from Section 7.2

Proposition D.14. Consider the RDPP mechanism $(\vec{x}, \vec{p}, T)$ with equilibrium price q . If $\text{WELFARE}(\vec{x}(q)) \geq \frac{q \cdot m}{1 + \delta}$ and $\Pr[N(q) < m] \geq \frac{\delta}{1 + \delta}$, then

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{OPT}{2(1 + \delta)} \min\{1, \delta\}.$$

Proof. Let c be a positive constant. First, we consider two cases:

Case 1. For the case $q \geq \frac{c}{m} OPT$, from the assumption $\text{WELFARE}(\vec{x}(q)) \geq \frac{q \cdot m}{1 + \delta}$, we have

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{q \cdot m}{1 + \delta} \geq \frac{c}{1 + \delta} OPT. \quad (20)$$

Case 2. Next, consider the case $q < \frac{c}{m} OPT$. From the assumption $\Pr[N(q) < m] \geq \frac{\delta}{1 + \delta}$, we might hope to not lose too much welfare in the event $\{N(q) \geq m\}$ since this event is correlated with bidders having smaller values. Moreover, it is easy to estimate the welfare whenever $N(q) \leq m$ — it is simply the sum of all $v_i \geq q$. Since $q < \frac{c}{m} OPT$, we can upper bound the welfare loss from the top $N(q) - m$ bidders with values $v_i < q$. Let $N_{-i}(q) := \sum_{j \neq i} \mathbb{1}_{v_j \geq q}$ and let $\vec{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ be the vector $\vec{v}$ excluding the i -th entry. Then,

$$\text{WELFARE}(\vec{x}(q)) = \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \cdot \mathbb{1}_{N(q) \leq m} \right] + \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \cdot \mathbb{1}_{N(q) > m} \right]$$




Figure 6: (Top-left) g_1 with initial point $q_0 = 0.1$; (Top-right) g_2 with initial point $q_0 = 0.1$; (Bottom-left) g_1 with initial point $q_0 = 4$; (Bottom-right) g_2 with initial point $q_0 = 4$. The solid line (blue) denotes function f_1 or f_2 . The dashed-dot line (red) denotes function g_1 or g_2 . The dashed line (black) is the function $y(x) = x$. Dots (green) denote the following sequence of points in $\mathbb{R}^2$: $(q_0, f(q_0)), (f(q_0), f^2(q_0)), (f^2(q_0), f^3(q_0)), \dots$

$$\begin{aligned}
&\geq \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \cdot \mathbb{1}_{N(q) \leq m} \right] \\
&= \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \cdot \mathbb{1}_{v_i \geq q} \cdot \mathbb{1}_{N(q) \leq m} \right] \\
&= \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \cdot \mathbb{1}_{v_i \geq q} \cdot \mathbb{1}_{N_{-i}(q) \leq m-1} \right] \\
&= \sum_{i=1}^n \mathbb{E}_{\vec{v}_{-i} \sim F^{n-1}} \left[\mathbb{1}_{N_{-i}(q) \leq m-1} \cdot \mathbb{E}_{v_i \sim F} \left[v_i \cdot x_i^*(\vec{v}) \cdot \mathbb{1}_{v_i \geq q} \cdot \mathbb{1}_{N_{-i}(q) \leq m-1} | \vec{v}_{-i} \right] \right].
\end{aligned}$$

The third line observes that given $N(q) \leq m$, we have that $x_i(q) = 1$ if and only if bidder i is one of the top m bidders with value $v_i \geq q$. The fourth line observes that the following events are equivalent:

- Bidder i with value $v_i \geq q$ is one of the top m bidders when there are at most m bidders with value at least q .
- Bidder i with value $v_i \geq q$ is one of the top m bidders when among bidders $[n] \setminus \{i\}$ there are at most $m - 1$ bidders with value at least q .

The fifth line follows from linearity of expectation and the law of total expectation by observing that $\vec{v}$ is drawn from the product distribution $F^n = \underbrace{F \times \dots \times F}_{n \text{ times}}$. Next, observe

that if among bidders $[n] \setminus \{i\}$ there are at most $m - 1$ bidders with value at least q , then the event that bidder i has value at least q implies bidder i is one of the top m bidders. Thus $\text{WELFARE}(\vec{x}(q))$ is at least

$$\begin{aligned}
&\geq \sum_{i=1}^n \mathbb{E}_{\vec{v}_{-i} \sim F^{n-1}} \left[\mathbb{1}_{N_{-i}(q) \leq m-1} \cdot \mathbb{E}_{v_i \sim F} [v_i \cdot \mathbb{1}_{v_i \geq q} | \vec{v}_{-i}] \right] \\
&= \sum_{i=1}^n \Pr[N_{-i}(q) \leq m-1] \mathbb{E}[v_i \cdot \mathbb{1}_{v_i \geq q}] \quad \text{From independence between } v_i \text{ and } \vec{v}_{-i}, \\
&\geq \Pr[N(q) < m] \mathbb{E} \left[\sum_{i=1}^n v_i \cdot \mathbb{1}_{v_i \geq q} \right] \quad N(q) < m \text{ implies } N_{-i}(q) \leq m-1, \\
&\geq \frac{\delta}{1+\delta} \mathbb{E} \left[\sum_{i=1}^n v_i \cdot \mathbb{1}_{v_i \geq q} \right] \quad \text{From the assumption } N(q) < m \text{ with probability at least } \delta/(1+\delta), \\
&= \frac{\delta}{1+\delta} \left(\mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \right] + \mathbb{E} \left[\sum_{i=1}^n v_i (\mathbb{1}_{v_i \geq q} - x_i^*(\vec{v})) \right] \right) \\
&\geq \frac{\delta}{1+\delta} \left(\mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \right] - \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i^*(\vec{v}) \cdot \mathbb{1}_{v_i < q} \right] \right) \\
&> \frac{\delta}{1+\delta} (OPT - q \cdot m) \quad \text{From the fact } v_i \cdot \mathbb{1}_{v_i < q} < q \text{ and } \sum_{i=1}^n x_i^*(\vec{v}) \leq m,
\end{aligned}$$

$$\geq \frac{\delta}{1+\delta}(1-c)OPT \quad \text{From the assumption } q \cdot m < c \cdot OPT.$$

The second line observes that for independent random variables X and Y , $\mathbb{E}[X \cdot Y] = \mathbb{E}[X] \mathbb{E}[Y]$. The chain of inequalities proves that for Case 2,

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{\delta}{1+\delta}(1-c)OPT.$$

Combining Case 1 and 2, we get that

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{OPT}{1+\delta} \min\{c, (1-c)\delta\}.$$

Setting $c = 1/2$ proves Proposition D.14. $\square$

Proof of Theorem 5. Let $N_{-i}(q) = \sum_{j \neq i} \mathbb{1}_{v_j \geq q}$ be the number of bidders that would purchase a slot at price q excluding bidder i . Recall that the value profile $\vec{v}$ is drawn from the product distribution F^n and whenever the demand $N(q) > m$, the active miner selects m random bidders with values at least q . Recall that the expected value of the random operator $T_{\text{TW}}^{\alpha, \delta}(q, B)$ is

$$E_{T_{\text{TW}}^{\alpha, \delta}}(q) = \alpha K_{\text{TW}}(q) + (1-\alpha)q,$$

and if q is an equilibrium price, $q = E_{T_{\text{TW}}^{\alpha, \delta}}(q)$ is a fixed point of $E_{T_{\text{TW}}^{\alpha, \delta}}$. From the expression for $E_{T_{\text{TW}}^{\alpha, \delta}}$, we get that $q = K_{\text{TW}}(q)$ is also a fixed point of K_{TW} . By the definition of K_{TW} (Equation 14), if q is a fixed point of K_{TW} , we have

$$q = K_{\text{TW}}(q) \geq (1+\delta)q \Pr[N(q) \geq m]. \quad (21)$$

Thus, the probability that at price q , the demand is at most $m-1$ is

$$\Pr[N(q) < m] = 1 - \Pr[N(q) \geq m] \geq \frac{\delta}{1+\delta}. \quad (22)$$

Next, we lower bound $\text{WELFARE}(\vec{x}(q))$ in terms of q and δ . Because q is a fixed point of K_{TW} , we have

$$\begin{aligned} q \cdot m &= K_{\text{TW}}(q) \cdot m = \mathbb{E} \left[\sum_{i=1}^n \min\{v_i, (1+\delta)q\} \cdot x_i(q) \cdot \mathbb{1}_{N(q) < m} + (1+\delta) \cdot q \cdot \mathbb{1}_{N(q) \geq m} \right] \\ &\leq \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(q) \cdot \mathbb{1}_{N(q) < m} + (1+\delta) \cdot q \cdot \mathbb{1}_{N(q) \geq m} \right] \\ &= \text{WELFARE}(\vec{x}(q)) + \mathbb{E} \left[\sum_{i=1}^n ((1+\delta) \cdot q - v_i) \cdot x_i(q) \cdot \mathbb{1}_{N(q) \geq m} \right] \\ &\leq \text{WELFARE}(\vec{x}(q)) + \delta \cdot q \cdot \mathbb{E} \left[\sum_{i=1}^n x_i(q) \cdot \mathbb{1}_{N(q) \geq m} \right] \quad \text{Observing } x_i(q) = 1 \text{ only if } v_i \geq q. \\ &= \text{WELFARE}(\vec{x}(q)) + \delta \cdot q \cdot m \cdot \Pr[N(q) \geq m] \end{aligned}$$

$$\leq \text{WELFARE}(\vec{x}(q)) + \frac{\delta \cdot q \cdot m}{1 + \delta} \quad \text{From (21).}$$

Rearranging the inequality, we get

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{q \cdot m}{1 + \delta}. \quad (23)$$

From (22), (23), and Proposition D.14, we establish

$$\text{WELFARE}(\vec{x}(q)) \geq \frac{OPT}{2(1 + \delta)} \min\{1, \delta\}.$$

□