



Citation for published version:

Lavi, R, Sattath, O & Zohar, A 2022, 'Redesigning Bitcoin's Fee Market', *ACM Transactions on Economics and Computation*, vol. 10, no. 1, 5, pp. 1-31. <https://doi.org/10.1145/3530799>

DOI:

[10.1145/3530799](https://doi.org/10.1145/3530799)

Publication date:

2022

Document Version

Peer reviewed version

[Link to publication](#)

Publisher Rights

Unspecified

© ACM, 2022. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *ACM Transactions on Economics and Computation*, volume 10, issue 1. <http://doi.acm.org/10.1145/3530799>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Redesigning Bitcoin’s fee market*

Ron Lavi¹, Or Sattath², and Aviv Zohar³

¹U. of Bath, UK, and Technion – Israel Institute of Technology

²Ben-Gurion University

³The Hebrew University

July 1, 2021

Abstract

The Bitcoin payment system involves two agent types: Users that transact with the currency and pay fees and miners in charge of authorizing transactions and securing the system in return for these fees. Two of Bitcoin’s challenges are (i) securing sufficient miner revenues as block rewards decrease, and (ii) alleviating the throughput limitation due to a small maximal block size cap. These issues are strongly related as increasing the maximal block size may decrease revenue due to Bitcoin’s pay-your-bid approach. To decouple them, we analyze the “monopolistic auction” [16], showing: (i) its revenue does not decrease as the maximal block size increases, (ii) it is resilient to an untrusted auctioneer (the miner), and (iii) simplicity for transaction issuers (bidders), as the average gain from strategic bid shading (relative to bidding one’s value) diminishes as the number of bids increases.

1 Introduction

Bitcoin’s security relies on its ability to attract honest miners and to incentivize them to invest large amounts of computational power [25]. This is done by rewarding miners for the creation of new blocks. The payment comprises a block reward that is a fixed amount per block, and a fee that each transaction pays for its inclusion in the block. Since the block reward is cut in half approximately every 4 years, transaction fees gradually become the miners’ main incentive to participate. The decision on which transactions to include in the block is made using a “pay-your-bid” auction: each transaction contains a fee for its inclusion in the block. A miner, in turn, includes the highest bidding transactions that fit in the block. In this auction, larger block sizes imply smaller transaction fees, since all

*A preliminary extended abstract of this paper appeared in [21].

transactions will aim to pay the lowest possible fee for their inclusion in the block. The Bitcoin block size therefore significantly influences miners’ revenue. Indeed, one of the main arguments against a block-size increase in Bitcoin has been that it may cause transaction fees to drop significantly (see, e.g., [17]). Thus, two of the Bitcoin system’s main problems, i.e., obtaining sufficient revenue for the miners as the block reward gradually decreases and the throughput limitation as a result of the maximal block size, are in fact strongly related. This connection, however, is a major obstacle to their resolution. Our motivation in this paper is to decouple the issues of revenue and maximal block size (which should be determined according to other technical considerations, like block propagation times) and to provide scalability from the economic perspective of the fee market.

While most of auction theory assumes a trusted auctioneer that honestly follows the protocol, a miner (the auctioneer of the block) cannot be trusted. If adding fake bids can increase profit, a miner can add transactions by moving funds from one of its accounts to another. Similarly, a miner can ignore bids. An untrusted miner can manipulate a 2nd price auction, for example, by adding a fake bid that is very close to the highest bid. Users will realize that and shade down their bids. Though cryptography can, in general, help resolve some of these issues, as far as the authors are aware, cryptographic techniques are not applicable to Bitcoin: they assume a trusted setup, that the miner is known in advance, or multiple communication rounds between all parties as in secure multi-party computation [26, 23, 1, 6].

We analyze here a potential solution to these issues, the *monopolistic auction* [16].¹ We observe that in the Bitcoin setting, the monopolistic auction is immune to untrusted miners, on the one hand, and it decouples the revenue issue from the maximal block size issue, on the other. In this auction, given bids b_1, \dots, b_n , the miner chooses which transactions to include in the block. All chosen transactions pay the same fee – the lowest bid in the block (as must be verified by the protocol). The maximal revenue resulting from this method cannot be increased by adding fake bids or by ignoring true bids. Other manipulations like side payments between miners and users are also not beneficial. In fact, we are not aware of any beneficial miner manipulation that can be exploited in the frame of this auction. Additionally and in contrast to a pay-your-bid auction, increasing the maximal block size does not decrease the revenue of the monopolistic auction. Choosing a threshold price is conceptually similar to dynamically choosing the block size – the miner chooses the size of the block it creates as a function of the bids received. In some bid instances, blocks will be small, while in others they will be large.

In a pay-your-bid auction, a user must strategically determine the lowest bid it can submit such that it will be included in the block (“bid shading”). Indeed, Bitcoin wallets use various fee estimation techniques. It is not clear how to do this optimally, and in any case, it requires some non-trivial computational effort (data gathering, statistical analysis, etc). Bid shading can also help in worst-case instances of the monopolistic auction, and as a result, [16] do not analyze its strategic properties but suggest it only as a benchmark for other truthful auctions

¹Termed the “optimal single price omniscient auction” in [16], this auction was suggested as a benchmark, and as far as we know, its game-theoretic properties have not been analyzed in the Bitcoin context.

they define. These other truthful auctions from [16] are less applicable vis-à-vis our needs (see Section 6). Although the approach is not truthful, in the monopolistic auction, a user, rather than paying her bid, pays some threshold value (a function of all bids) which w.h.p is only remotely related to one’s bid. The main technical novelty of this paper is to show that the expected gain from optimal bid shading relative to truthfully bidding the user’s value (which is her maximal willingness to pay) decreases to zero as the number of bids grows. On average, all users will benefit very little, if at all, from bid shading in the monopolistic auction. Since a near optimal strategy is to bid the user’s value, bidding in the monopolistic auction is simpler. Section 2 defines the formal framework used and proves this statement, and Section 4 verifies it via simulations using synthetic and actual Bitcoin bid distributions. The greater simplicity of the monopolistic auction is another one of its advantages over the pay-your-bid auction.

Our work analyzes a single block creation, and ignores “patient users”, who are willing to wait for future blocks. These users might lower their bids if they anticipate less competition in the future. Two important issues that remain to be investigated in future research are (1) how patient users shade down their bids and (2) how non-myopic users who create persistent transactions affect the bid distribution. Note that these issues affect the revenue not only of the monopolistic auction, but also of the pay-your-bid auction. We do not know, however, which of the two auctions is more strongly affected. Our analytical results suggest that the monopolistic auction collects at least as much revenue from impatient users as Bitcoin’s current mechanism. We discuss temporal considerations further in Section 6.1.

Subsequent Work. After the preprint version of this manuscript was published, Andrew Chi-Chih Yao [34] further analyzed the monopolistic price as a fee mechanism for Bitcoin, in the process resolving our main conjecture, Conjecture 1 – see the discussion there. In addition, Mark Friedenbach [14, 15] proposed an implementation of our proposed fee market as a Bitcoin soft-fork (i.e., an upgrade which is compatible with older version of the Bitcoin client, see [27] for details regarding a soft-fork).

Basu, Easley, O’Hara and Sirer [5] used a modified monopolistic price auction to determine inclusion in the block. They penalize miners if the blocks are not full, and additionally, they average the reward from blocks between different miners. Their goal is to maximize social welfare rather than the miners’ revenue. The main analytical difference between their work and ours is that theirs is based mostly on simulations while ours is proved for the most part formally. Perhaps surprisingly, even though their design goal and their mechanism is different than ours, in terms of abstract conclusions, our results agree with theirs: they also establish that the incentive of miners and users to manipulate diminishes as n grows.

Additional Related Work. There are relatively few works devoted to analyses of the Bitcoin fee market and fewer still that suggest modifications to the basic mechanism. [19] provided an early analysis of the economics of Bitcoin mining, including a game theoretic analysis of the incentives to mine on the longest chain. [4] considered the incentives of miners to distribute transactions to each other and designed ways to share the fees in exchange for distribution. [10] explored the security of Bitcoin and the incentive to mine blocks properly when block rewards diminish and fees dominate the revenue of miners. They showed that

variance in fees may undermine the security of the protocol and subject it to different forms of deviations and attacks. [7] explored related bribery attacks on the protocol that are paid for through promises of higher rewards for attackers that construct blocks off the longest chain. [29] considered the removal of the block limit altogether, arguing that delays in large block propagation times, which, in turn, imply a higher likelihood that the block is abandoned (often referred to as an orphaned block), will cause miners to restrict their own block sizes. The paper analyzed the resultant fee market that emerges. [18] modeled the transaction fee market and assumed that users benefit less if their transactions are delayed. They showed that such delays together with the congestion that may naturally occur in blocks due to queuing effects can lead to non-zero bids for transacting users even if blocks are not completely full. As far as we are aware, no previous work has explored a different mechanism for the fee market in crypto-currencies.

The paper [2] shows that a class of auctions that includes the monopolistic auction is “strategy-proof in the large”, when values are i.i.d. from some distribution with a finite support. Strategy-proofness in the large means that truthfulness is an epsilon BNE and epsilon goes to zero as the number of bidders goes to infinity. Our limit results on the monopolistic auction are similar, although we prove this for a different, stronger, notion – the discount ratio. See a more detailed discussion in Section 2.

Two recent papers [22, 11] take the axiomatic approach to characterize good block reward schemes, i.e., schemes that satisfy sets of axioms, and show that the proportional selection rule is the unique selection rule satisfying several reasonable axioms.

Several surveys that discuss the state-of-art of integrating game theoretic considerations with cryptocurrencies are given, e.g., in [9, 24, 3].

Terminology In general, we use Bitcoin’s terminology (e.g., users, transactions and miners). In some rare cases (especially when citing seminal auction theory results) we use the standard auction theory terminology instead (bidders, auctioneer, etc.).

Structure of this paper. Section 2 lays out the basic model, and further studies the properties of the monopolistic auction. Section 3 further generalizes the monopolistic auction to account for multiple strategic bids submitted by the same user. Section 4 conducts an extensive empirical evaluation of our theoretical claims and conjectures. This is especially useful since our theoretical bounds holds in the large n limit. Section 5 conducts a theoretical and experimental analysis of the suitability of the RSOP auction in our setting. Section 6 gives concluding remarks and suggestions for future research directions. Appendices A to F provide full proofs, and Appendix H lists the main notations and symbols we use.

2 The Monopolistic Auction

There are n transactions, each of which has a privately known value v_i for its inclusion in the block. This is the maximum fee that the transaction is willing to pay. The value can be deduced, e.g., by comparing to an alternative transaction cost via a bank, credit card,

etc. Transactions submit bids, and based on these, the miner decides which transactions to include in the block. The bid vector is sorted so that $\mathbf{b} = (b_1, \dots, b_n)$ satisfies $b_1 \geq \dots \geq b_n$. Let $k^*(\mathbf{b})$ be the index k that maximizes $k \cdot b_k$. In case of ties, k^* is chosen to be the maximal such index. Define:

$$R(\mathbf{b}) \equiv k^*(\mathbf{b}) \cdot b_{k^*(\mathbf{b})}, p^M(\mathbf{b}) \equiv b_{k^*(\mathbf{b})}, \quad (1)$$

where $R(\mathbf{b})$ is termed the “monopolistic revenue” and $p^M(\mathbf{b})$ is termed the monopolistic price. In the monopolistic auction, the miner includes in the block the $k^*(\mathbf{b})$ transactions who submitted the highest bids and charges a payment of $p^M(\mathbf{b}) = b_{k^*(\mathbf{b})}$ from each of these transactions. The miner’s revenue from this block is therefore $R(\mathbf{b})$.

An untrusted miner that can add fake bids does not have any incentive to do so in this auction. Specifically, the Bitcoin protocol can verify that all transactions in the same block pay the same price. However, it cannot prevent miners from adding fake bids by exploiting transactions that move funds between two accounts of the same miner. These fake bids can affect the payments of the “real” bids and increase the revenue of the miner, as is the case in a second-price auction. In the monopolistic auction, in contrast, the addition of fake bids cannot increase the revenue: given any vector of real bids \mathbf{b} , since all transactions in the block pay the same price, the maximal revenue is $R(\mathbf{b})$ if the miner does not create another block in the near future. However, a user might benefit from submitting $b_i < v_i$, as demonstrated by the following examples:

Example 1. *Suppose n users that have the same value $v_i = 1$. If all bids are $b_i = v_i$, then $R(\mathbf{b}) = n$, $k^*(\mathbf{b}) = n$, and $p^M(\mathbf{b}) = 1$, i.e., all transactions get accepted to the block and they all pay 1. A strategic user, say 1, can reduce her payment by reducing her bid to $b_1 = \frac{n-1}{n}$. With this bid, the monopolistic price decreases to $\frac{n-1}{n}$: if all transactions that bid 1 are accepted, the revenue is $n - 1$, and if all transactions that bid at least $b_1 = \frac{n-1}{n}$ are accepted, the revenue is still $n - 1$.*

The gain from bid shading in this example vanishes as the number of transactions (and the block size) increases. In other cases, gains do not vanish even when the number of transactions grows:

Example 2. *$v_1 = \dots = v_{\frac{n}{2}+1} = 2, v_{\frac{n}{2}+2} = \dots = v_n = 1$, given some even n . If user 1 is truthful and bids 2, her transaction will get accepted to the block and the price will be 2. A careful look reveals that she can reduce her price to 1 by bidding 1 instead.*

However, if the values for this example are taken from realistic distributions, then the probability that it will be realized will usually be very small.

Therefore, rather than allege that there are no profitable deviations from truth-telling, we will only claim that the expected gain from any non-truthful bid diminishes as n grows. For example, if the distribution of values has a finite support, bidding the next lower value in the support *weakly dominates* truth-telling. More pronounced bid shading may be beneficial; We will show, however, that the expected gain from such a deviation or from any other deviation becomes extremely small as n grows. Therefore, we believe that it makes sense to conclude that no deviations will likely happen.

Example 3. Consider the distribution F with finite support: $\Pr(v = 1) = 0.5, \Pr(v = 2) = 0.5$. For simplicity, assume that users can only submit bids of 1 or 2. Every user with $v = 2$ has a weakly dominant strategy to bid 1, since she derives 0 utility when the monopoly price is set to be 2, and by bidding 1 she increases the chance that the monopoly price will be 1. However, the probability that bidding 1 instead of 2 will actually make a difference is vanishingly small. Specifically, a user with $v = 2$ can improve the outcome (i.e., change the empirical monopoly price from 2 to 1) by bidding $b=1$, when the number of users is $2n$, only in the event that there are exactly $n + 1$ users with value equals to 2. The probability of this event is $\frac{\binom{2n}{n+1}}{2^{2n}} = \Theta\left(\frac{1}{\sqrt{n}}\right)$. We note that this is the most extreme example – as the support size increases, the bid shading suggested in this example becomes negligible.

We validate this intuition more generally via both theoretical and empirical analyses. First we define some terms. Fix a user i and a vector of bids $\mathbf{b}_{-i} \equiv (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ of the other users. Define the “strategic price”:

$$p^S(\mathbf{b}_{-i}) \equiv \inf\{b_i \in \mathbb{R} \mid p^M(b_i, \mathbf{b}_{-i}) \leq b_i\} = \min\{b_i \in \mathbb{R} \mid p^M(b_i, \mathbf{b}_{-i}) \leq b_i\}. \quad (2)$$

The equality is proven in Appendix A. In words, $p^S(\mathbf{b}_{-i})$ is the lowest possible bid for i to be included in the block, when the other bids are \mathbf{b}_{-i} . Figure 1 gives an example. This definition does not capture a case where a user can provide false-name bids (i.e., can split her transaction into several transactions), which is discussed in Section 3.

Figure 2 presents the strategic prices in an example with $n = 256$ users. In this example, p^S is monotonically decreasing among winning users (higher winning bids imply lower strategic prices) but not among all users. This property is formally proved in Claim 9 in Appendix B. Additionally, here and in most other simulations, almost all winning bids have the same strategic price, the exception being the very few lowest winning bids (i.e., those with prices slightly higher than the monopolistic price). The strategic price of these bids is higher, i.e., they gain less from being strategic.

While $p^S(\mathbf{b}_{-i})$ is the minimal price that user i must pay to be included in the block, she will pay $p^M(v_i, \mathbf{b}_{-i})$ if she bids her true value. This possible gain from bid shading is captured by the notion of the “discount ratio”:

$$\delta_i(v_i, \mathbf{b}_{-i}) \equiv \begin{cases} 1 - \frac{p^S(\mathbf{b}_{-i})}{p^M(v_i, \mathbf{b}_{-i})} & \text{if } v_i \geq p^S(\mathbf{b}_{-i}) \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, $0 \leq \delta_i(v_i, \mathbf{b}_{-i}) \leq 1$. If $v_i < p^S(\mathbf{b}_{-i})$, every bid of at most v_i loses. Hence, the gain from bid shading is 0. If $v_i \geq p^S(\mathbf{b}_{-i})$, user i can win and pay $p^S(\mathbf{b}_{-i})$. With an honest bid, she will pay $p^M(v_i, \mathbf{b}_{-i})$. She can thus save a percentage of $1 - \frac{p^S(\mathbf{b}_{-i})}{p^M(v_i, \mathbf{b}_{-i})}$ from her price by strategic bid shading.

Example 2 above shows that the discount factor does not vanish in the worst case, even as n grows. In this example, the discount factor is $\frac{1}{2}$ for all possible values of n . We therefore conduct an average-case analysis. Assume that all true values are drawn i.i.d. from some distribution F on $\mathbb{R}_{>0}$. The average discount ratio is then $\Delta_n^{average} = \mathbb{E}_{(v_1, \dots, v_n) \sim F}[\delta_1(v_1, \mathbf{v}_{-1})]$




Figure 1: The function $p^M(b_1, 1)$ as a function of b_1 , demonstrating that p^M is neither monotone nor continuous. As can be seen, the strategic price of the first user is 0.5.

(the choice of user 1 is arbitrary since all are symmetric a-priori). Note that this definition implicitly assumes that $\mathbf{b}_{-1} = \mathbf{v}_{-1}$. This is similar to the usual logic in equilibrium analysis – assume all others are truthful and show that user i does not have a reason to deviate from truthfulness as well. Indeed, we will show that all discount ratios go to zero as n goes to infinity. We therefore rely on this logic in all our definitions.

We also consider two stronger notions. In the first, for each realization v_1, \dots, v_n , we consider the maximal discount ratio among all users:

$$\delta_{max}(\mathbf{v}) \equiv \max_i \delta_i(v_i, \mathbf{v}_{-i}), \quad \Delta_n^{max} \equiv \mathbb{E}_{(v_1, \dots, v_n) \sim F} [\delta_{max}(\mathbf{v})]$$

Clearly, for all n , $\Delta_n^{max} \geq \Delta_n^{average}$ since for every \mathbf{v} and every i , $\delta_{max}(\mathbf{v}) \geq \delta_i(v_i, \mathbf{v}_{-i})$.

Theorem 1. *For any distribution F with a finite support size, $\lim_{n \rightarrow \infty} \Delta_n^{max} = 0$.*

Section 2.1 proves this theorem.

In the second extension, defined only for distributions with a finite bounded support, we fix player 1 and assume that she deterministically has a value that maximizes her discount ratio (the worst possible value for her, from our perspective) while all other values are probabilistic. This depicts a situation commonly assumed in game theory wherein a user knows




Figure 2: An example with $n = 256$ bids sampled i.i.d. from the uniform distribution on $[0, 1]$. For each user, the x-axis shows her values and the y-axis shows her strategic price. Winning bids, which are marked with “x”, must be at least as high as their respective strategic price, i.e., to the right of the black line.

her own value and (only) the distribution of the other values. Let $v_{max} \equiv \max \text{Support}(F)$.

$$\delta_n^{GT}(\mathbf{b}_{-1}) \equiv \max_{v_1 \in \text{Support}(F)} \delta_1(v_1, \mathbf{b}_{-1}) = \delta_1(v_{max}, \mathbf{b}_{-1}). \quad (3)$$

The equality is proven in Claim 10.

$$\Delta_n^{GT} \equiv \mathbb{E}_{(v_2, \dots, v_n) \sim F} [\delta_n^{GT}(\mathbf{v}_{-1})] \quad (4)$$

Theorem 2. *For any distribution F with a finite support size, $\lim_{n \rightarrow \infty} \Delta_n^{GT} = 0$.*

The proof of this theorem is very similar to that of Theorem 1, as explained in Section 2.1.

Our empirical analysis focuses on Δ_n^{max} , since it is well-defined even for distributions with unbounded support.

While $\lim_{n \rightarrow \infty} |\Delta_n^{GT} - \Delta_n^{max}| = 0$, the relation between Δ_n^{GT} and Δ_n^{max} depends on the distribution, i.e., one term is not always larger than the other. For example, consider the distribution $\Pr(v_i = 1) = p$ and $\Pr(v_i = \epsilon) = 1 - p$, where $0 < \epsilon < 1$. In this case, if $p < 0.5$,

then $\Delta_{n=2}^{GT} > \Delta_{n=2}^{max}$, and if $p > 0.5$, then $\Delta_{n=2}^{GT} < \Delta_{n=2}^{max}$, since:

$$\begin{aligned}\Delta_{n=2}^{GT} &= \mathbb{E}_{v_2 \sim F}[\delta_1(1, v_2)] = p\delta(1, 1) + (1-p)\delta(1, \epsilon) = p \cdot \frac{1}{2} + (1-p) \cdot \left(1 - \frac{\epsilon}{2}\right) \\ \Delta_{n=2}^{max} &= \mathbb{E}_{(v_1, v_2) \sim F}[\max_i \delta_i(v_1, v_2)] \\ &= p^2\delta_1(1, 1) + (1-p)^2\delta_1(\epsilon, \epsilon) + 2p(1-p)\max_i \delta_i(1, \epsilon) \\ &= p^2\delta_1(1, 1) + (1-p)^2\delta_1(\epsilon, \epsilon) + 2p(1-p)\delta_1(1, \epsilon) \\ &= (1-2p(1-p)) \cdot \frac{1}{2} + 2p(1-p) \cdot \left(1 - \frac{\epsilon}{2}\right)\end{aligned}$$

We believe that our theorems can be further generalized:

Conjecture 1 (Nearly Bayesian-Nash Incentive-Compatibility).

1. For any distribution F , $\lim_{n \rightarrow \infty} \Delta_n^{average} = 0$. In particular, $\Delta_n^{average} = O\left(\frac{1}{\sqrt{n}}\right)$, where the constant in the $O(\cdot)$ notation may depend on F .
2. If F has a bounded support (still possibly with an infinite cardinality) $\lim_{n \rightarrow \infty} \Delta_n^{max} = 0$. In particular, $\Delta_n^{max} = O\left(\frac{1}{\sqrt{n}}\right)$. The same holds for Δ_n^{GT} .
3. There exists a distribution F with unbounded support such that $\lim_{n \rightarrow \infty} \Delta_n^{max} > 0$. The same holds for Δ_n^{GT} .

The $O\left(\frac{1}{\sqrt{n}}\right)$ in the above conjecture can be achieved by Example 3.

Recently, Andrew Chi-Chih Yao resolved many issues that were left open in this work. Indeed, “The main purpose [of Yao’s paper] is to settle the Nearly IC Conjecture for the monopolistic price in the positive” [34]. Specifically, the first sentences in items 1 and 2 in the conjecture above were proved, and a slightly weaker claim was shown than the “in particular” claims in these items. This was proved also when accounting for multiple strategic bids – see Section 3. Item 3 was proved for the inverse distribution (see Eq. (8)), which was our primary candidate for this conjecture based on numerical data – see Section 4.

Game Theoretic Equilibrium Analysis. In a Bayesian-Nash equilibrium (BNE), user i ’s utility is $u_i(v_i, b_i, \mathbf{b}_{-i}) = v_i - p^M(\mathbf{b})$ if $b_i \geq p^M(\mathbf{b})$ and 0 otherwise. A strategy $s_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ gives i ’s bid $s_i(v_i)$. $(s_1(\cdot), \dots, s_n(\cdot))$ is an ϵ -BNE if $\forall i$ and $v_i, b_i \in \text{Support}(F)$: $\mathbb{E}_{\mathbf{v}_{-i}}[u_i(v_i, s_i(v_i), \mathbf{v}_{-i})] \geq \mathbb{E}_{\mathbf{v}_{-i}}[u_i(v_i, b_i, \mathbf{v}_{-i})] - \epsilon$. The “truthful” strategy $s_i(v_i) = v_i$ is an ϵ -BNE if for all i ,

$$\max_{v_i, b_i \in \text{Support}(F)} \mathbb{E}_{\mathbf{v}_{-i}}[u_i(v_i, b_i, \mathbf{v}_{-i}) - u_i(v_i, v_i, \mathbf{v}_{-i})] \leq \epsilon \quad (5)$$

Our discount-ratio analysis is similar but formally different. In fact, our approach is stronger.

Theorem 3. Assume that F is bounded, with v_{max} being the maximal point on the support of F . For any n , $s_i(v_i) = v_i$ is an ϵ -BNE with $\epsilon = v_{max} \cdot \Delta_n^{GT}$ in the monopolistic auction with n users.

The proof is given in Appendix E.

If F has a finite support, Theorem 2 states that $\lim_{n \rightarrow \infty} \Delta_n^{GT} = 0$. Therefore, there exists a function $N(\epsilon)$ such that, for every $\epsilon > 0$, $n \geq N(\epsilon)$, truthfulness is an ϵ -BNE in the game with n users. The opposite direction is not true; if truthfulness is an ϵ -BNE, this does not imply that the discount ratio is at most ϵ . This is because the deviation b_i in the BNE analysis is fixed before taking the expectation, while in our analysis b_i is a best response to a specific \mathbf{v}_{-i} and we take the expectation over these possibly different best responses. In this respect, the discount-ratio analysis is stronger.

2.1 Proof of Theorem 1

Recall that F is the distribution over true values. Given \mathbf{v} , define: $i^* \equiv \operatorname{argmax}_{i=1, \dots, n} v_i$, $k^* \equiv k^*(\mathbf{v}_{-i^*})$, and $p^S \equiv p^S(\mathbf{v}_{-i^*})$ ('S' stands for 'Strategic').

Lemma 1. For any F with finite support size, there exists a constant $c > 0$ (which may depend on F but not on n) s.t. $\lim_{n \rightarrow \infty} \Pr(k^* < c \cdot n) = 0$.

Proof. Let $v_{max} = \max \operatorname{Support}(F)$, k_{max} be the number of users in \mathbf{v} whose value is v_{max} , and $p_{max} = \Pr_{v_i \sim F}[v_i = v_{max}]$. By linearity of expectation, $\mathbb{E}[k_{max}] = p_{max} \cdot n$. A Chernoff bound implies that $\lim_{n \rightarrow \infty} \Pr(k_{max} < \frac{9np_{max}}{10}) = 0$. Users who bid v_{max} win; therefore, $k^* + 1 \geq k_{max}$. Thus, $\lim_{n \rightarrow \infty} \Pr(k^* + 1 < \frac{9np_{max}}{10}) = 0$. Choosing $c = \frac{8p_{max}}{10}$ completes the proof of the lemma. \square

Appendix B shows three useful properties:

Property 1. $\forall \mathbf{v}, i, p^M(p^S(\mathbf{v}_{-i}), \mathbf{v}_{-i}) = p^S(\mathbf{v}_{-i})$.

Property 2. Let y be the smallest element in the support of F that is at least p^S . Then, $y = p^M(\mathbf{v})$ implies $p^S \geq \frac{k^*}{k^*+1} \cdot p^M(\mathbf{v})$.

Property 3. $\forall \mathbf{v}, i, j: v_i \geq v_j$ implies $\delta_i(v_i, \mathbf{v}_{-i}) \geq \delta_j(v_j, \mathbf{v}_{-j})$.

With these, we prove a second lemma:

Lemma 2. $\lim_{n \rightarrow \infty} \Pr(p^S < \frac{k^*}{k^*+1} p^M(\mathbf{v})) = 0$.

Proof. Define A as the event where $p^S < \frac{k^*}{k^*+1} p^M(\mathbf{v})$. The proof defines an event B s.t. (i) $\lim_{n \rightarrow \infty} \Pr(B) = 0$, and (ii) $A \subseteq B$. This implies $\lim_{n \rightarrow \infty} \Pr(A) = 0$ as claimed.

Define $\operatorname{num}(\mathbf{v}, z) \equiv |\{v_i | v_i \geq z\}|$ and the random variables $n_z = \operatorname{num}(\mathbf{v}, z)$ for any real number z . Let B be the event in which $\exists x, y \in \operatorname{Support}(F)$ s.t. $x > y$ and $n_x \cdot x > n_y \cdot y \geq (n_x - 1)x$. Since the support is of finite size, there is a fixed number of such pairs

$x > y \in \text{Support}(F)$. By Claim 16 in Appendix D and the union bound over these finite set of pairs, we conclude that $\lim_{n \rightarrow \infty} \Pr(B) = 0$. We show that $A \subseteq B$, i.e., $p^S < \frac{k^*}{k^*+1} p^M(\mathbf{v})$ implies $\exists x > y \in \text{Support}(F)$ s.t. $n_x \cdot x > n_y \cdot y \geq (n_x - 1)x$. Let $x = p^M(\mathbf{v})$ and let y be the smallest element in the support of F which is at least p^S . By Property 2, the event $p^S < \frac{k^*}{k^*+1} p^M(\mathbf{v})$ implies $x > y$. Since $p^M(v_{i^*}, \mathbf{v}_{-i^*}) = x$, it follows that $n_x \cdot x > n_y \cdot y$. Furthermore,

$$n_y \cdot y \geq n_{p^S} \cdot p^S = \text{num}(\mathbf{v}, p^S) \cdot p^S = \text{num}((p^S, \mathbf{v}_{-i^*}), p^S) \cdot p^S \geq \text{num}((p^S, \mathbf{v}_{-i^*}), x) \cdot x = (n_x - 1)x$$

where the first step follows since $n_y = n_{p^S}$ and $p^S \leq y$, the third step follows since $v_{i^*} \geq x > y \geq p^S$, the fourth step follows from Property 1 that shows that $p^M(p^S, \mathbf{v}_{-i^*}) = p^S$, and the fifth step follows since $v_{i^*} \geq x > y \geq p^S$. \square

Let the “bad” event E_1 be the case where $k^* < c \cdot n$ (c is taken from Lemma 1) and the “bad” event E_2 be the case where $p^S < \frac{k^*}{k^*+1} p^M(\mathbf{v})$. By Property 3, $\delta_{\max}(\mathbf{v}) = \delta_{i^*}(\mathbf{v})$. Therefore, $\Delta_n^{\max} = \mathbb{E}_{\mathbf{v}}[\delta_{i^*}(\mathbf{v})]$. If E_2 does not hold then $p^S \geq \frac{k^*}{k^*+1} p^M(\mathbf{v})$ and therefore $\delta_{i^*}(\mathbf{v}) \leq \frac{1}{k^*+1}$. To conclude:

$$\begin{aligned} \lim_{n \rightarrow \infty} \Delta_n^{\max} &= \lim_{n \rightarrow \infty} \left[\Pr(E_1 \cup E_2) \mathbb{E}_{\mathbf{v}}[\delta_{i^*}(v_{i^*}, \mathbf{v}_{-i^*}) | E_1 \cup E_2] \right. \\ &\quad \left. + \Pr(E_1^c \cap E_2^c) \mathbb{E}_{\mathbf{v}}[\delta_{i^*}(v_{i^*}, \mathbf{v}_{-i^*}) | E_1^c \cap E_2^c] \right] \\ &\leq \lim_{n \rightarrow \infty} \left[\Pr(E_1) + \Pr(E_2) + \Pr(E_1^c \cap E_2^c) \cdot \mathbb{E}_{\mathbf{v}}\left[\frac{1}{k^*+1} | E_1^c \cap E_2^c\right] \right] \\ &\leq \lim_{n \rightarrow \infty} \left[\Pr(E_1) + \Pr(E_2) + \Pr(E_1^c \cap E_2^c) \cdot \frac{1}{c \cdot n} \right] \\ &\leq \lim_{n \rightarrow \infty} \left[\Pr(E_1) + \Pr(E_2) + \frac{1}{c \cdot n} \right] = 0, \end{aligned}$$

implying Theorem 1. The proof of Theorem 2 is very similar. The main difference between the two theorems is that in the latter, we need to set the value of the first player to the maximal value in the support of F . All statements of the above proof then hold with respect to the first player, instead of player i^* .

3 Multiple strategic bids

In some cases it is beneficial to split one’s bid to several separate transactions with several separate bids. In fact, such a strategy sometimes enables a losing transaction to be included in the block, as the following example demonstrates:

Example 4. Let $\mathbf{v} = (5, 2, 1, 1)$. With bids $\mathbf{b} = \mathbf{v}$, the monopolistic price is 5 and the second user loses. However, if she submits two separate transactions with a bid of 1 each, the monopolistic price will be 1 and all transactions will be included.

Our empirical evaluation below accounts for such situations, showing that the benefit from using multiple bids also goes to zero as n increases. For this purpose, we generalize p^S to capture the possible benefit of splitting a single bid to u different bids. Formally, define $p^{\text{multibid}}(\mathbf{b}_{-i})$ as the minimum of $u \cdot b_i^{(u)}$ over all $u \in \mathbb{N}^+$ and all $b_i^{(1)} \geq \dots \geq b_i^{(u)} \in \mathbb{R}$ such that $b_i^{(u)} \geq p^M(b_i^{(1)}, \dots, b_i^{(u)}, \mathbf{b}_{-i})$. Choosing $u = 1$ gives us p^S , hence $p^S(\mathbf{b}_{-i}) \geq p^{\text{multibid}}(\mathbf{b}_{-i})$. Example 4 shows that there are cases in which the inequality is strict. It is easy to show that $p^{\text{multibid}}(\mathbf{b}_{-i})$ can be written as:

$$p^{\text{multibid}}(\mathbf{b}_{-i}) \equiv \min\{u \cdot b_i \mid u \in \mathbb{N}^+, b_i \in \mathbb{R}, b_i \geq p^M(\overbrace{b_i, \dots, b_i}^{u \text{ times}}, \mathbf{b}_{-i})\}. \quad (6)$$

The minimum is well defined using an argument similar to that used for p^S . In particular, for any positive integer u , use $f_u(b_i) = p^M(b_i, \dots, b_i, \mathbf{b}_{-i})$ (where b_i appears u times). By Claim 11, $u \in \{1, \dots, n\}$. Since for every u , the infimum is contained in the set, then so is the infimum over the union over $u \in \{1, \dots, n\}$.

We believe that Theorem 1 and Conjecture 1 hold in the case of multiple bids, where in the definition of δ , p^{multibid} replaces p^S . In our empirical evaluation we use p^{multibid} instead of p^S . In all the distributions we examined, except the discrete distribution, we *never* encountered a case in which the strategic player has an advantage placing multiple bids. Even in the discrete case, the effect of such multiple bids was negligible.

The definition of p^{multibid} gives little information on its algorithmic computation, since the optimization is done on an infinite set. We provide an alternative definition which yields a polynomial-time algorithm. We use this algorithm in our simulations. Fix a player i and $\mathbf{w} = \mathbf{v}_{-i}$. Assume w.l.o.g. that \mathbf{w} is sorted, $w_1 \geq w_2, \dots, \geq w_{n-1}$. For every $k^*(\mathbf{w}) \leq j \leq n-1$, define

$$f(j) \equiv \max \left\{ \left\lceil \frac{R(\mathbf{w})}{w_j} \right\rceil, j+1 \right\}.$$

Intuitively, if user i (the strategic user) adds $f(j) - j$ additional bids to \mathbf{w} , all equal to $\frac{R(\mathbf{w})}{f(j)}$, then this will be the new monopolistic price, and user i will win. It is thus clear that these manipulations should be considered in order to determine the optimal solution of Eq. (6). While a-priori it may be possible to consider other manipulations as well, the next theorem shows that considering this type of manipulations is sufficient.

Theorem 4.

$$p^{\text{multibid}}(\mathbf{w}) = \min_{k^*(\mathbf{w}) \leq j \leq n-1} \frac{R(\mathbf{w})}{f(j)} (f(j) - j).$$

Furthermore, if j^* is the index that minimizes this term, then $b^* = \frac{R(\mathbf{w})}{f(j^*)}$ and $u^* = f(j^*) - j^*$ minimizes the r.h.s. of Eq. (6).

For intuition, we refer the reader to Example 4. Indeed, in that example, $\mathbf{w} = (5, 1, 1)$, $k^*(\mathbf{w}) = 1$, $R(\mathbf{w}) = 5$. We can calculate $f(1) = 2$, $f(2) = 5$, $f(3) = 5$. Therefore, $j^* = 3$, resulting in $b^* = 1$, $u^* = 2$ as described in the example above.

Proof. To prove this theorem, we begin with a useful Lemma:

Lemma 3. *Let b, u be the arguments which minimize Eq. (6), j be the integer which satisfies $w_j \geq b > w_{j+1}$ and $u^* = f(j) - j$. Then, $u^* \geq 1$ and $u \geq u^*$.*

Proof. Since $f(j) \geq j + 1$, $u^* \geq 1$. If $f(j) = j + 1$, $u^* = 1$ and the claim immediately follows. Thus assume that $f(j) = \lceil \frac{R(\mathbf{w})}{w_j} \rceil$. Suppose towards a contradiction that $u \leq u^* - 1$. Then,

$$R(\overbrace{b, \dots, b}^{u \text{ times}}, \mathbf{w}) = b \cdot (u + j) \leq w_j \cdot (u^* - 1 + j) = w_j \cdot (f(j) - 1) < R(\mathbf{w}) \quad (7)$$

where the first step follows since by Claim 12

$$b = p^M(\overbrace{b, \dots, b}^{u \text{ times}}, \mathbf{w})$$

and since there are exactly j bids in \mathbf{w} that are at least b ; and the last inequality follows from $f(j) < \frac{R(\mathbf{w})}{w_j} + 1$.

However Eq. (7) shows a contradiction since adding bids can only increase the monopolistic revenue. This completes the proof of Lemma 3. \square

We now prove Theorem 4. Let b, u be the arguments that determine $p^{\text{multibid}}(\mathbf{w})$ (i.e., (b, u) minimizes Eq. (6)). Let j be the integer that satisfies $w_j \geq b > w_{j+1}$ (where $j = n - 1$ if $w_{n-1} \geq b$). By Claim 13, $j \geq k^*(\mathbf{w})$. Let $u^* = f(j) - j$ and $b^* = \frac{R(\mathbf{w})}{f(j)}$. Then,

$$b \cdot (u + j) = R(\overbrace{b, \dots, b}^{u \text{ times}}, \mathbf{w}) \geq R(\mathbf{w}) = b^* \cdot (u^* + j)$$

where the first step is the same as the first step in Eq. (7) above. Thus, $b \cdot u \geq b^* \cdot u^* + j \cdot (b^* - b)$. By Lemma 3, $u \geq u^*$, which implies $b \cdot u \geq b^* \cdot u^*$ (since $b \cdot u < b^* \cdot u^*$ and $u \geq u^*$ implies $b < b^*$ and $b \cdot u < b^* \cdot u^* + j \cdot (b^* - b)$, a contradiction). Thus,

$$p^{\text{multibid}}(\mathbf{w}) = b \cdot u \geq b^* \cdot u^* \geq \min_{k^*(\mathbf{w}) \leq j \leq n-1} \frac{R(\mathbf{w})}{f(j)} (f(j) - j),$$

and the first part of the theorem follows. The second part follows in a straight-forward way from the first part. \square

4 Empirical Evaluation

We provide empirical evidence to support the above conjectures and to supply greater detail on the rate at which the discount ratio converges to zero. We use both synthetic data as well as transaction data taken from the Bitcoin blockchain. The synthetic data are generated using four distributions:

1. A uniform discrete distribution over the integers $1, \dots, 100$. This distribution has a finite support size and therefore satisfies the requirements of Theorem 1.




Figure 3: A histogram of the sum-of-outputs on the x-axis (in log scale) and its frequency in the collected data (on the y-axis).

2. The uniform distribution over $[0, 1]$. Notice that here the support size is infinite.
3. The half normal distribution.² Here, the probability for value x decreases exponentially with x , hence, even though arbitrarily high transaction values will be seen, they are highly unlikely.
4. The inverse distribution

$$F(x) = 1 - \frac{1}{x}, \quad x \in [1, \infty). \quad (8)$$

Here, the probability for a value x decreases polynomially with x .

The uniform distribution has no tail, the half normal distribution has a light tail (it decreases exponentially fast), and the inverse distribution has a heavy tail.

Bitcoin blockchain data was collected from 1000 consecutive blocks, which constitute roughly one week of activity ending on October 28th, 2016. Figure 3 shows a histogram of the collected data. The data obviously do not contain the values, since this is not how Bitcoin operates. The simulations estimate the value v , as a function $v = v(x)$ of the transaction size x (x is the sum of all outputs of the Bitcoin transaction). We use three alternative functions: $v(x) = \log x$, $v(x) = \sqrt{x}$, $v(x) = x$. Note that multiplying all bids by a scalar does

²Recall that the half normal distribution arises from taking the absolute value of a normal random variable. This is needed in our setting since valuations are non-negative.




Figure 4: The average discount ratio of a player from selfish bidding as a function of the number of players that participate.

not change the discount ratio, which is the case when, e.g., the value is some percentage of x .

We empirically evaluated $\Delta_n^{average}$ and Δ_n^{max} (where $p^{multibid}$ replaces p^S) as a function of the number of bids, n . For each $n \in \{2^3, 2^4, \dots, 2^{17}\}$, we conducted 100 simulation runs. Each run samples n bids and calculates the n empirical discount ratios. $\Delta_n^{average}$ is the average of the n individual discount ratios and Δ_n^{max} is the maximum over the n individual discount ratios. Each point in the graph is the average of 100 runs.

Figure 4 shows that $\Delta_n^{average}$ behaves almost identically for all tested distributions. In particular, as stated in the first part of Conjecture 1, $\Delta_n^{average}$ decreases linearly with the number of bids n and for all distributions, even those with an infinite and unbounded support size.

Figure 5 shows that Δ_n^{max} behaves differently for some of the tested distributions. For the uniform distribution, Δ_n^{max} decreases linearly with the number of bids, supporting the second part of Conjecture 1. The half normal distribution behaves similarly (even though it is not bounded). For the inverse distribution, Δ_n^{max} does not seem to decrease with the number of users, and we believe that this is in fact an example that supports the third part of Conjecture 1. The three Bitcoin distributions behave as follows: the log of the sum decreases the fastest, the square root of the sum decreases more slowly, and the sum decreases the slowest. For example, for $n = 2048$ (roughly the current Bitcoin block-size), Δ_n^{max} is about 0.2%, 5%, and 19% for the three distributions, respectively. When $n = 2^{17} \approx 130,000$, Δ_n^{max} is about 0.0007%, 0.05%, and 1%.




Figure 5: The maximal discount ratio of a player from selfish bidding as a function of the number of players.

These findings show that there is a qualitative difference between $\Delta_n^{average}$ and Δ_n^{max} , mainly for the inverse distribution and for the Bitcoin distribution with $v(x) = x$. In these distributions, although the average transaction will benefit very little from bid shading, some transactions can nonetheless obtain significant benefit. This outcome can, in principle, result from two different reasons. The first is that $\Delta_n^{average}$ also accounts for the losing transactions, whose discount ratio is zero. The second is that Δ_n^{max} only accounts for the single winning user with the maximal discount ratio. However, typical user behavior (as demonstrated in Figure 2) shows that almost all winning bids have the lowest strategic price and, therefore, the highest discount ratio. Thus, the explanation for the difference between the two discount ratios seems to be the first.

Figure 6 shows all simulation points for the Bitcoin log-of-sum distribution. The x-axis is the resulting *block size* of the simulation run (in the previous figures it was n) and the y-axis is the maximal discount ratio of the simulation run. Three main conclusions can be drawn from Figure 6. First, block sizes (number of winners) range from 1 to about 25,000 when $v(x) = \log x$ (recall that the total number of bids ranges from 2^3 to 2^{17}).³ Second, this range in block sizes is the main reason why Δ_n^{max} is smaller in the Bitcoin log-of-sum distribution, as larger block sizes imply smaller discount factors. Third, the distribution of

³The throughput when $v(x) = \log x$ will therefore be about $0.1n$ (n is the number of all transactions). For $v(x) = x$, block sizes range from 1 to about 1,500, implying a throughput of about $0.01n$. This can be compared to Bitcoin's actual throughput in October 2016 which varied between 1300 and 1900 transactions per-block (the variation is mostly due to variations in transaction sizes in bytes).




Figure 6: A scatter plot of all simulation points for the Bitcoin distribution with $v(x) = \log x$.

the simulation points, which include outliers, is not normal. We do not have an explanation for this finding. The last two remarks apply to all tested distributions.

5 The Suitability of the RSOP Auction in Our Context

The monopolistic-price auction is miner-honest in the sense that a myopic miner cannot increase her revenue by diverging from the suggested protocol (for example, add false bids or delete existing bids). However, as discussed above, even impatient users can sometimes benefit from submitting a strategic bid which is different than their true value.

In this section we discuss an alternative approach in which we use a well-known auction which is truthful for impatient users. That is, an impatient user will maximize her gain by simply submitting her value. On the other hand, in this second mechanism the challenge is to show that miners will be honest and will not manipulate the protocol, as we further discuss below. In particular, we will show empirical evidence suggesting that the profit for the miners from such manipulations becomes negligible as the number of users increases.

More specifically, our second suggested mechanism uses the following RSOP (Random Sampling Optimal Price) auction, defined by [16]. In the definition, the following notation is used: For a subset $A \subseteq [n]$ of users, and a bid vector \mathbf{b} , let $\mathbf{b}_A = (b_i)_{i \in A}$. We define $p_A^M(\mathbf{b}) \equiv p^M(\mathbf{b}_A)$. When \mathbf{b} is clear from the context, we simply write p_A^M .

Definition 1 (The RSOP auction [16]). *Upon receiving n bids, $\mathbf{b} = b_1, \dots, b_n$, the miner performs the following:*

1. Randomly partition the bids to two disjoint sets A and B (each bid is placed in A with probability $\frac{1}{2}$, otherwise it is placed in B).
2. Compute the monopolistic price for each set: p_A^M, p_B^M . The monopolistic price of an empty set is defined to be zero.
3. The set of winning bids is $A' \cup B'$, where:

$$A' = \{i \in A : b_i \geq p_B^M\} \quad ; \quad B' = \{i \in B : b_i \geq p_A^M\}.$$

The transactions in A' each pay p_B^M , and the transactions in B' pay p_A^M each. The revenue obtained in the auction is therefore

$$RSOP(\mathbf{b}) = |A'| \cdot p_B^M + |B'| \cdot p_A^M.$$

For the RSOP mechanism, Goldberg *et al.* provide the following results that hold under “the usual” auction theory assumptions, most notably that: (1) the auctioneer is honest; (2) bidders do not collude; (3) each bidder can submit exactly one bid; (4) bidders are impatient, i.e. they do not obtain utility from winning in future auctions (we discuss these assumptions further below).

Theorem 5 ([16], Observation 6.2 and Theorem 6.4).

1. Truthfulness. *The RSOP auction is truthful, i.e., a bidder maximizes her utility by reporting her value, even when the other bidders do not reveal their value.*
2. Maximal Revenue. *Fix any parameter h . Let \mathbf{b} be any bid vector of n bids with $b_i \in [1, h]$ for all i . Then*

$$\lim_{n \rightarrow \infty} \max_{\mathbf{b}} \frac{R(\mathbf{b})}{RSOP(\mathbf{b})} = 1.$$

To instantiate the RSOP-based mechanism, users must specify a bid that ideally represents their values for each transaction, and the miners are asked to create a block with all transactions they wish to *potentially* include. Unlike the current Bitcoin protocol, here not all transactions in a block are valid. After the block is mined, and propagated to the Bitcoin nodes in the network, they determine which transactions are valid, by running Algorithm 1.

Algorithm 1 RSOP block verification

- 1: A node receives a new block B .
 - 2: Check validity of the block and of all the included transactions according to Bitcoin's current rules. When referring to transactions in previous blocks, consider only valid transactions (as explained below).
 - 3: Compute the sets A and B using the block hash as a seed to a Cryptographically Secure Pseudo-Random Number Generator (CSPRNG).
 - 4: The transactions that are considered valid are the ones in A' and B' as in Def. 1.
 - 5: Transactions in A' pay p_B^M and transactions in B' pay p_A^M . A fraction $1 - \alpha$ of that revenue goes to the miner who mined the current block, and the rest goes to the future miner who would mine the next valid block. Invalid transactions do not pay anything. The parameter $0 \leq \alpha \leq 1$ needs to be specified as part of the protocol.
-

The algorithm has the following advantages. Since the source of the randomness (block hash) used in line 3 is the same for all the Bitcoin nodes, all nodes will reach consensus on the set of valid transactions. Notice that the block hash is determined by its contents including the solution to a proof-of-work puzzle. A miner cannot easily manipulate the choice of the sets A and B to maximize revenue⁴ (if she could, it could destroy the truthfulness of the protocol). She *can* choose to forgo a block in which A and B were not set to her liking but will need to generate the proof-of-work from scratch, which is extremely costly. There are however several concerns with the proposed algorithm:

User truthfulness. The assumption that users cannot send multiple bids is unrealistic in the Bitcoin setting. Therefore, we do not know whether the suggested protocol is truthful even for impatient users and honest miners. Under the assumption that the user cannot control which bids will be associated with subset A and which with subset B , we do not know how to construct even tailored worst-case examples or if there exists an efficient algorithm to find beneficial deviations.⁵

Miner honesty: adding false bids. Miners are able to manipulate the protocol by adding false bids, especially if at step 5 of Algorithm 1, $\alpha = 0$. The purpose of choosing $\alpha > 0$ is to mitigate this problem. The following example demonstrates this issue:

Example 5. *Suppose $\alpha = 0$. There are two users with $b_1 = h$, $b_2 = \ell$ where $h > 2\ell$. With probability $\frac{1}{2}$ both users will fall to the same set, and the revenue for the miner would be 0. If they fall into different sets, only the h user will be included, and will pay ℓ . Therefore, the expected revenue is $\frac{\ell}{2}$. A strategic miner can create many false transactions that bid h , by this receiving a revenue of approximately h . The false bids clearly cannot create a profit for the miner but also do not harm her because she pays them to herself. As α increases, the profitability of this strategy decreases.*

⁴A more careful cryptographic analysis is required to establish this statement, but is outside the scope of this work. This statement is fairly simple to show in the random oracle model.

⁵For the monopolistic auction, we discuss this issue in Sections 3 and 4.

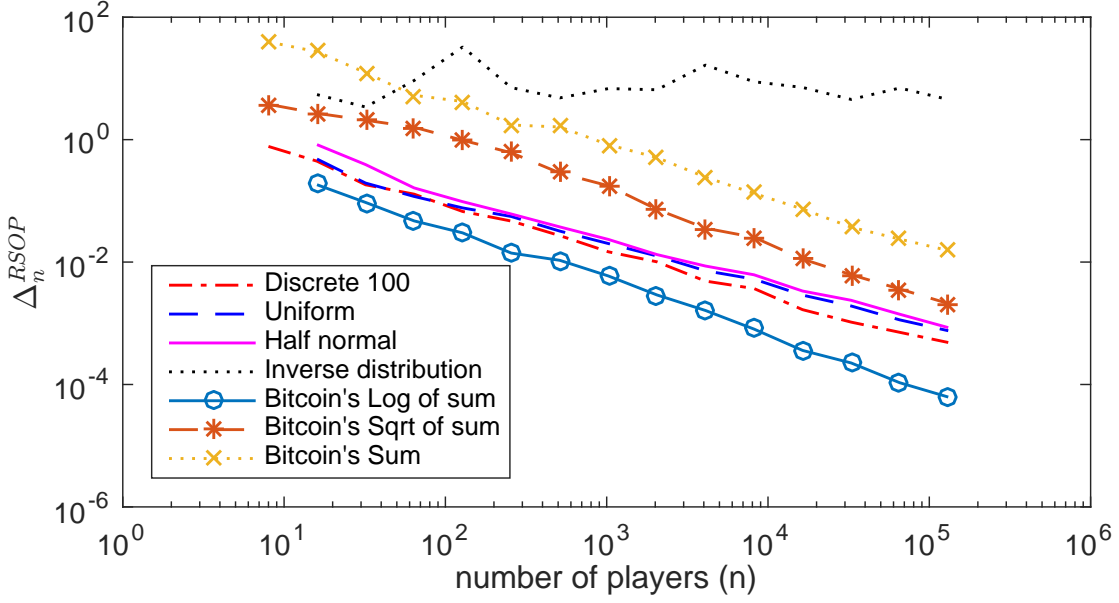


Figure 7: Δ_n^{RSOP} as a function of the number of players, for various distributions.

More generally, a dishonest miner can always use the following simple strategy to obtain the revenue $R(b)$: (1) Compute the monopolistic price over all bids; (2) Add sufficiently many false transactions whose bids are equal to the monopolistic price. Adding sufficiently many such false bids will change the RSOP revenue to that of the monopolistic price mechanism.

We conjecture that the strategy in the preceding paragraph is always beneficial to the miner, when $\alpha = 0$:

Conjecture 2. *For every \mathbf{b} and all choices of A and B , the RSOP revenue is at most the monopolistic revenue. In particular, $RSOP(\mathbf{b}) \leq R(\mathbf{b})$.*

The right hand side ($R(\mathbf{b})$) is the revenue a manipulating miner can obtain with the strategy above, and the left hand side is the RSOP revenue, for *any* allocation of the sets A and B (even one chosen adversarially). When $\alpha = 0$, this is the revenue of the miner which uses Algorithm 1 honestly, i.e., without false bids.

In the empirical evaluation reported below, we have never encountered a counterexample to Conjecture 2. To measure the expected “gain ratio” from performing this strategy, we define

$$\Delta_n^{RSOP} \equiv \mathbb{E}_{(v_1, \dots, v_n) \sim F} \left[\frac{R(\mathbf{v})}{RSOP(\mathbf{v})} - 1 \right].$$

For distributions with bounded support sizes Theorem 5 proves that Δ_n^{RSOP} goes to zero as n goes to infinity. Furthermore, Figure 7 demonstrates that this indeed happens in a reasonably fast manner for almost all the distributions we considered in Section 4, with the only exception being the inverse distribution. To summarize this issue:

- **Gaining from adding false bids when $\alpha = 0$.** In the RSOP mechanism the miner can gain from adding false transactions with a bid equal to the monopolistic price. This is an easy strategy to perform, and as far as we can tell, it never harms the miner when $\alpha = 0$. If the miner does so, the mechanism changes and becomes similar to the monopolistic-price mechanism. This, in turn, may harm the bidder-truthfulness property. However, the gain ratio from performing this strategy decreases as the number of users increases (with the exception of the Inverse distribution) – see Figure 7.
- **Choosing $\alpha > 0$ to eliminate the gain from adding false bids.** We conjecture that, for many bid distributions and when the number of users is fairly large, it is possible to set $\alpha > 0$ in a way that will eliminate the profitability of this strategy. Note that if $\alpha > 0$, the miner needs to pay some of her false bids to other miners. On the other hand, setting $\alpha > 0$, and especially high values (say, $\alpha = 0.9$), introduces risks of side-payments. A miner may provide the following service. She asks users to submit transactions with 0 bid, and pay directly to the miner a fixed per transaction fee. The miner will guarantee that she will eventually include these transactions in a block which contain only other 0 bid transactions (which will then all be valid). In such a case the miner gets all the revenue for herself, whereas if she follows the protocol she will only receive a fraction $1 - \alpha$ of the payments. If α is low, this is less likely to happen, as this requires non-trivial coordination between the miner and the user (e.g., trust regarding the transfer of payments, regarding inclusion in the block, etc).

Miner honesty: removing true bids. The miner might also be able to increase her profit by removing true bids. The following example demonstrates this:

Example 6. *There are n bids with value 2 and n bids with value 1. The miner can obtain a revenue of $2n$ w.h.p. by removing all bids of value 1. The expected revenue for the honest miner is always at most $2n$, and in the limit of large n , the revenue is $2n - \sqrt{\frac{n}{\pi}} + o(\sqrt{n})$, see Appendix F Cor. 2.*

Note that the maximal revenue that can be obtained in this example by removing bids is exactly the monopolistic revenue. We do not know if this is true in general, however we strongly suspect that this is the case. More specifically, we conjecture:

Conjecture 3. *For any vector of bids \mathbf{b} , $\max_{\mathbf{b}' \subseteq \mathbf{b}} RSOP(\mathbf{b}') \leq R(\mathbf{b})$.*

In fact, Conjectures 2 and 3 are equivalent: by choosing $\mathbf{b}' = \mathbf{b}$, Conjecture 3 implies 2. To show the other direction, Conjecture 2 implies that $RSOP(\mathbf{b}') \leq R(\mathbf{b}')$ and by the monotonicity of $R(\cdot)$, we reach $RSOP(\mathbf{b}') \leq R(\mathbf{b})$, as required.

Note that under this conjecture, the gain ratio Δ_n^{RSOP} goes to zero for any distribution with finite support even if the miner can remove true bids, using Theorem 5. Note that this form of the conjecture is harder to verify empirically as it requires considering different subsets for removal. But the equivalence to Conjecture 2, and the empirical evaluation performed above provides some supporting evidence for it.

We do not know an efficient algorithm to find the best bids to remove, and therefore the strategy for the miner is left open. We do not have any suggestions regarding how to mitigate this case. Unlike the previous case (adding false bids) α is irrelevant to this strategy by the miner. Perhaps it can be shown that removing true bids is rarely beneficial, under some distributional assumptions.

Block size. Notice that in the RSOP mechanisms blocks may contain many transactions that do not get accepted (their bid may be below the price that is eventually determined). One possible way to prevent this is to commit to all transactions in the Merkle tree in some canonical order of bid size and then eventually only reveal transactions that are needed to establish the monopolistic price and transactions that win this bid. While we do not have a fully fleshed out scheme, we do believe that clever use of data structures and cryptographic schemes may help in reducing the amount of wasted space in blocks in this way.

6 Discussion

The simple pay-your-bid fee mechanism used in Bitcoin has several disadvantages: (i) A fixed (hard-coded) maximal block size implies non-optimal revenue extraction when the prefixed block-size is too small or too large for the current bid instance, (ii) Bitcoin wallets must invest computational effort in bid shading, and (iii) deciding on the maximal block size involves economic rather than purely technological considerations (e.g., block propagation time [12, 31]). For single block creation, the proposed monopolistic auction effectively solves all these issues while maintaining miner non-manipulability.

The monopolistic auction is ϵ -truthful (ϵ goes to zero as the number of users grows). Alternatively, we can take a truthful auction where a user cannot influence its winning price, e.g., the RSOP auction is truthful and its revenue $RSOP(\mathbf{v})$ satisfies $\lim_{n \rightarrow \infty} \max_{\mathbf{v}} \frac{R(\mathbf{v})}{RSOP(\mathbf{v})} = 1$ [16]. Section 5 describes an implementation in Bitcoin (e.g., the required randomness needs to be verified by other miners) and discusses major drawbacks of it: First, to verify correctness, the block must contain all rejected transactions, creating significantly larger blocks. Second, truthfulness does not rule out beneficial bid splits. We do not know whether RSOP is resilient to such a manipulation. We discussed the effect of bid splits in the monopolistic auction in Section 3. Third, although $\frac{R(\mathbf{v})}{RSOP(\mathbf{v})} \approx 1$, nonetheless $R(\mathbf{v}) - RSOP(\mathbf{v})$ still usually grows to infinity (e.g., if $R(\mathbf{v}) = O(n + \sqrt{n})$ and $RSOP(\mathbf{v}) = O(n)$). Thus, a miner can significantly increase its revenue by engaging in manipulative behavior, for example, (i) deletion of true bids (however, we do not know if there exists a polynomial-time algorithm to find which bids to delete), (ii) adding many false bids, all of which are equal to the monopolistic price of the true bids (a possible solution to this attack is to split the block fees between the miner that created the block and the miner that will create the next block). In light of these shortcomings, we believe that the monopolistic auction is better suited to implement improvements in Bitcoin's fee market.

6.1 Temporal considerations

An important issue we did not cover is how the current mechanism would behave in a realistic setting where the interaction with miners is repeated. Just as in other mechanisms (for example, the second price auction) moving from a single shot interaction to a repeated one does not necessarily maintain the mechanism’s truthfulness. In our case, one could consider a repeated setting, in which users who did not manage to have their transactions included in a block still persist and request a transfer from the miners again (indeed, by default, Bitcoin transactions persist in the miner’s queue if they were not already included in a block). In this setting, the monopolistic mechanism, which clears the queue of its highest value transactions will slowly cause the distribution of transactions in the queue to be more skewed towards lower-bid transactions compared to the distribution of new incoming transactions. We then expect that occasionally the mechanism will cause a price fluctuation, when the weight of lower-bid transactions shifts the monopolistic price down. In this case, the occasional drop in price may entice patient users (who we did not consider in our model) to bid lower values and have their transactions accepted whenever such a drop occurs. The mechanism is therefore not truthful if applied in this manner.

Another important aspect to consider if patient users are present is the effect on the strategic behavior of miners. We have argued that in the one-shot monopolistic auction, miners have no incentive to add false bids or delete true ones. We believe this will be the case for small miners even when patient users exist, since small miners cannot anticipate to mine blocks often. However, with large miners (or mining pools), a different analysis will be required. Consider an example where at every round, 51% of new transactions are of value 1, and 49% of transactions are of value 2. In this case, a myopic miner may include all transactions in his block with a monopolistic price of 1. A more far-sighted miner that has almost 100% of the hashing power⁶ may instead alternate between blocks that contain only transactions with value 1 and blocks that only contain transactions with value 2, relying on the ability to defer transactions by one block, and gaining in the process.

On the face of it, a simple posted-price technique could handle such temporal considerations. Specifically, gather statistics about the underlying distribution and calculate an optimal fixed price for transactions to be included in the block. However, since the true underlying distribution is really unknown, such a method is not straight-forward and contains several significant challenges. First, it is hard to learn the full distribution, since blocks only contain information about accepted transactions, and therefore there is no consensus regarding the tail of the bidding distribution. Second, the distribution itself may change as a result of the mechanism’s own behavior if transactions that are not included in blocks persist and attempt to enter future blocks. Third, such a statistical method is highly sensitive to miner manipulations as they can easily alter sampled transactions and their transaction fees. One recent quite sophisticated attempt to handle these issues is contained in Ethereum’s recent proposal EIP1559 [8] which is examined thoroughly in [30] (which also assumes that users are myopic with regards to their transactions, as we do) and in a related analysis in [13].

⁶Of course, such a miner could double-spend transactions, censor users, etc., so, in some sense, the manipulation which we discuss is the least of our worries.

A key idea used within EIP1559 proposes burning money to avoid the miners manipulating the posted amount that transactions are charged. The analysis in [13] further shows that EIP1559 is not always stable and a different posted price mechanism is proposed. All this demonstrates that successfully handling such issues even in a myopic setting and we believe that this is an open problem which deserves and requires significant further effort.

6.2 Future Work

The monopolistic auction can also be explored further via a careful equilibrium analysis. The astute reader may notice that the current model allows for an equilibrium in which all users bid 0. This is an unrealistic setting since in practice there exists a maximum limit on the block size, and not all transactions can be included in the block. Still, an elaborate game-theoretic analysis of all possible equilibria in our model is interesting.

One of the modeling uncertainties in this work regards the users' values. We chose to analyze various distributions (see, e.g., Figure 5) since we have little information regarding the maximal willingness to pay of the users. This issue has been studied extensively in several economic contexts [28, 20, 33] and we believe that studying it in Bitcoin will be beneficial.

The work in [32] shows that, given Bitcoin's current reward scheme, it is beneficial to miners to introduce a time gap between the time a miner observes the last block and the time the miner starts mining the new block. It could be interesting to examine whether our proposed new reward scheme eliminates this issue as well.

7 Acknowledgments

We thank Alaa Mozalbat, who designed and implemented the simulation code, and Ittay Eyal and Reshef Meir for valuable discussions. R.L. was partly supported by a Marie-Curie fellowship "Advance-AGT", by an ARCHES award from the MINERVA foundation, and by the ISF-NSFC joint research program (grant No. 2560/17). O.S. was supported by ERC Grant 280157, by the Israel Science Foundation (personal grant 682/18 and Quantum Technologies and Science Program grant 2137/19), and by the Cyber Security Research Center at Ben-Gurion University. A.Z. was supported by the Israel Science Foundation (grant 616/13) and by a grant from the HUJI Cyber Security Research Center in conjunction with the Israel National Cyber Bureau (grant 039-9230).

References

- [1] Abe, M., Suzuki, K.: M+1-st price auction using homomorphic encryption. In: Naccache, D., Paillier, P. (eds.) *Public Key Cryptography, 5th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2002, Paris, France, February 12-14, 2002, Proceedings. Lecture Notes in Computer Science*, vol. 2274, pp. 115–124. Springer (2002). https://doi.org/10.1007/3-540-45664-3_8

- [2] Azevedo, E.M., Budish, E.: Strategy-proofness in the Large. *The Review of Economic Studies* **86**(1), 81–116 (08 2018). <https://doi.org/10.1093/restud/rdy042>
- [3] Azouvi, S., Hicks, A.: Sok: Tools for game theoretic models of security for cryptocurrencies (2019), <http://arxiv.org/abs/1905.08595>
- [4] Babaioff, M., Dobzinski, S., Oren, S., Zohar, A.: On Bitcoin and red balloons. In: *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*. pp. 56–73 (2012). <https://doi.org/10.1145/2229012.2229022>
- [5] Basu, S., Easley, D., O'Hara, M., Siner, E.G.: Towards a functional fee market for cryptocurrencies (2019), <http://arxiv.org/abs/1901.06830>
- [6] Bogetoft, P., Damgård, I., Jakobsen, T.P., Nielsen, K., Pagter, J., Toft, T.: A practical implementation of secure auctions based on multiparty integer computation. In: Crescenzo, G.D., Rubin, A.D. (eds.) *Financial Cryptography and Data Security, 10th International Conference, FC 2006, Anguilla, British West Indies, February 27-March 2, 2006, Revised Selected Papers. Lecture Notes in Computer Science, vol. 4107*, pp. 142–147. Springer (2006). https://doi.org/10.1007/11889663_10
- [7] Bonneau, J.: Why buy when you can rent? - bribery attacks on bitcoin-style consensus. In: *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*. pp. 19–26 (2016). https://doi.org/10.1007/978-3-662-53357-4_2
- [8] Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., Bakhta, A.: EIP-1559: Fee market change for ETH 1.0 chain (2019), <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md>
- [9] Böhme, R., Christin, N., Edelman, B., Moore, T.: Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives* **29**(2), 213–38 (May 2015). <https://doi.org/10.1257/jep.29.2.213>
- [10] Carlsten, M., Kalodner, H.A., Weinberg, S.M., Narayanan, A.: On the instability of Bitcoin without the block reward. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*. pp. 154–167 (2016). <https://doi.org/10.1145/2976749.2978408>
- [11] Chen, X., Papadimitriou, C.H., Roughgarden, T.: An axiomatic approach to block rewards. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies, AFT 2019, Zurich, Switzerland, October 21-23, 2019*. pp. 124–131. ACM (2019). <https://doi.org/10.1145/3318041.3355470>
- [12] Decker, C., Wattenhofer, R.: Information propagation in the bitcoin network. In: *13th IEEE International Conference on Peer-to-Peer Computing, IEEE P2P 2013, Trento, Italy, September 9-11, 2013, Proceedings*. pp. 1–10 (2013). <https://doi.org/10.1109/P2P.2013.6688704>

- [13] Ferreira, M.V., Moroz, D.J., Parkes, D.C., Stern, M.: Dynamic posted-price mechanisms for the blockchain transaction-fee market (2021)
- [14] Friedenbach, M.: Rebatable fees & incentive-safe fee markets. <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2017-September/015093.html> (2017)
- [15] Friedenbach, M.: Forward blocks: On-chain/settlement capacity increases without the hard-fork. <http://freico.in/forward-blocks-scalingbitcoin-paper.pdf> (2018)
- [16] Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M.E., Wright, A.: Competitive auctions. *Games and Economic Behavior* **55**(2), 242–269 (2006). <https://doi.org/10.1016/j.geb.2006.02.003>
- [17] Houy, N.: The economics of bitcoin transaction fees (2014). <https://doi.org/10.2139/ssrn.2400519>
- [18] Huberman, G., Leshno, J.D., Moallemi, C.C.: Monopoly without a monopolist: An economic analysis of the bitcoin payment system. <https://ssrn.com/abstract=3025604> (2017)
- [19] Kroll, J.A., Davey, I.C., Felten, E.W.: The economics of bitcoin mining, or bitcoin in the presence of adversaries. In: *Proceedings of WEIS*. vol. 2013 (2013)
- [20] Lange, C., Martin, C., Chabanet, C., Combris, P., Issanchou, S.: Impact of the information provided to consumers on their willingness to pay for champagne: comparison with hedonic scores. *Food quality and preference* **13**(7-8), 597–608 (2002)
- [21] Lavi, R., Sattath, O., Zohar, A.: Redesigning bitcoin’s fee market. In: Liu, L., White, R.W., Mantrach, A., Silvestri, F., McAuley, J.J., Baeza-Yates, R., Zia, L. (eds.) *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*. pp. 2950–2956. ACM (2019). <https://doi.org/10.1145/3308558.3313454>
- [22] Leshno, J., Strack, P.: Bitcoin: An impossibility theorem for proof-of-work based protocols. *Cowles Foundation Discussion Paper* (2019)
- [23] Lipmaa, H., Asokan, N., Niemi, V.: Secure vickrey auctions without threshold trust. In: Blaze, M. (ed.) *Financial Cryptography, 6th International Conference, FC 2002, Southampton, Bermuda, March 11-14, 2002, Revised Papers. Lecture Notes in Computer Science*, vol. 2357, pp. 87–101. Springer (2002). https://doi.org/10.1007/3-540-36504-4_7
- [24] Liu, Z., Luong, N.C., Wang, W., Niyato, D., Wang, P., Liang, Y.C., Kim, D.I.: A survey on applications of game theory in blockchain (2019), <http://arxiv.org/abs/1902.10865>

- [25] Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008), <http://www.bitcoin.org/bitcoin.pdf>
- [26] Naor, M., Pinkas, B., Sumner, R.: Privacy preserving auctions and mechanism design. In: Feldman, S.I., Wellman, M.P. (eds.) Proceedings of the First ACM Conference on Electronic Commerce (EC-99), Denver, CO, USA, November 3-5, 1999. pp. 129–139. ACM (1999). <https://doi.org/10.1145/336992.337028>
- [27] Narayanan, A., Bonneau, J., Felten, E.W., Miller, A., Goldfeder, S.: Bitcoin and Cryptocurrency Technologies - A Comprehensive Introduction. Princeton University Press (2016), <http://press.princeton.edu/titles/10908.html>
- [28] Noussair, C., Robin, S., Ruffieux, B.: Revealing consumers’ willingness-to-pay: A comparison of the bdm mechanism and the vickrey auction. *Journal of Economic Psychology* **25**(6), 725 – 741 (2004). <https://doi.org/https://doi.org/10.1016/j.joep.2003.06.004>, <http://www.sciencedirect.com/science/article/pii/S0167487003000837>
- [29] Rizun, P.R.: A transaction fee market exists without a block size limit. Block Size Limit Debate Working Paper (2015), <https://www.bitcoinunlimited.info/resources/feemarket.pdf>
- [30] Roughgarden, T.: Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559 (2020)
- [31] Sompolinsky, Y., Zohar, A.: Secure high-rate transaction processing in bitcoin. In: Financial Cryptography and Data Security - 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers. pp. 507–527 (2015). https://doi.org/10.1007/978-3-662-47854-7_32
- [32] Tsabary, I., Eyal, I.: The gap game. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018. pp. 713–728. ACM (2018). <https://doi.org/10.1145/3243734.3243737>
- [33] Wicker, P.: Willingness-to-pay in non-profit sports clubs. *International Journal of Sport Finance* **6**(2), 155 (2011)
- [34] Yao, A.C.: An incentive analysis of some bitcoin fee designs (2018), <http://arxiv.org/abs/1811.02351>

A The Definition of Strategic Price is Well Formed

The goal of this section is to show that the infimum in Eq. (2) is equal to the minimum. This is shown in Claim 1 below, using $f(b_i) = p^M(b_i, b_{-i})$. Claim 2 shows that $f(b_i)$ satisfies the property required by Claim 1, and clearly, there exists b_i such that $b_i \geq f(b_i)$, e.g., taking any $b_i > \max_{j \neq i} b_j$.

Claim 1. Let $f : \mathfrak{R}_{\geq 0} \rightarrow \mathfrak{R}_{\geq 0}$ be any function that satisfies the following property: If $x < f(x)$ there exists $\delta > 0$ such that for any $0 < \epsilon < \delta$, $f(x + \epsilon) = f(x)$. Let $A = \{x \in \mathfrak{R}_{\geq 0} \mid x \geq f(x)\}$. Then, if A is not empty, $\inf A \in A$.

Proof. Note that if A is not empty than the infimum of A is well defined since A is bounded from below by 0. Let $x^* = \inf A$. Assume towards a contradiction that $x^* \notin A$, i.e., $x^* < f(x^*)$. Thus, by the property of f , there exists $\epsilon > 0$ such that $x^* + \epsilon < f(x^*)$ and $f(x) = f(x^*)$ for every $x^* < x < x^* + \epsilon$. Since $x^* = \inf A$ there exists $x \in A$, $x^* < x < x^* + \epsilon$. For this x , $f(x) \leq x < x^* + \epsilon < f(x^*) = f(x)$, a contradiction. \square

Claim 2. Fix any \mathbf{b}_{-i} and any $b_i < p^M(b_i, \mathbf{b}_{-i})$. Then there exists $\delta > 0$ such that for any $0 < \epsilon < \delta$, $p^M(b_i + \epsilon, \mathbf{b}_{-i}) = p^M(b_i, \mathbf{b}_{-i})$.

Proof. Let $\mathbf{v} = (b_i, \mathbf{b}_{-i})$ and assume w.l.o.g. that \mathbf{v} is ordered, i.e., $v_1 \geq v_2 \geq \dots \geq v_n$. Let k be such that $v_k = p^M(\mathbf{v})$. Let j be the minimal index of bid b_i in \mathbf{v} , i.e., $v_{j-1} > v_j = b_i$. Note that $j > k$ since $v_j < v_k$ by assumption. Also note that $v_j \cdot j < v_k \cdot k$ since the monopolistic price is v_k . Choose $\delta > 0$ such that $v_j + \delta < v_{j-1}$ and $(v_j + \delta) \cdot j < v_k \cdot k$. By definition, for any $0 < \epsilon < \delta$, $p^M(v_j + \epsilon, \mathbf{v}_{-j}) = p^M(\mathbf{v})$, and the claim follows. \square

B Auxiliary Technical Claims

This section proves multiple useful properties of the monopolistic auction. We recommend that this section be read linearly. Recall that $\text{num}(\mathbf{v}, z) \equiv |\{v_i \mid v_i \geq z\}|$.

Claim 3. If $p^M(\mathbf{v}_{-i}) \leq v_i$, then $p^M(\mathbf{v}_{-i}) \leq p^M(\mathbf{v})$. Furthermore, if $p^M(\mathbf{v}_{-i}) = v_i$, then $p^M(\mathbf{v}_{-i}) = p^M(\mathbf{v})$.

Proof. Let $x = p^M(\mathbf{v}_{-i})$ and $y = p^M(\mathbf{v})$. Assume by contradiction that $x > y$. By definition of the monopolistic price for \mathbf{v}_{-i} , $x \cdot \text{num}(\mathbf{v}_{-i}, x) > y \cdot \text{num}(\mathbf{v}_{-i}, y)$. Since $v_i \geq x > y$, $\text{num}(\mathbf{v}, x) = \text{num}(\mathbf{v}_{-i}, x) + 1$ and $\text{num}(\mathbf{v}, y) = \text{num}(\mathbf{v}_{-i}, y) + 1$. Thus, $x \cdot \text{num}(\mathbf{v}, x) > y \cdot \text{num}(\mathbf{v}, y) - y + x > y \cdot \text{num}(\mathbf{v}, y)$, contradicting $y = p^M(\mathbf{v})$.

For the second part of the claim, note that the monopolistic price cannot increase when adding v_i , otherwise v_i will not be included in the optimal block and therefore it cannot make a difference. The first part of this claim shows that the monopolistic price cannot decrease, hence the equality. \square

Claim 4. For any \mathbf{v} and any i , $p^S(\mathbf{v}_{-i}) < p^M(\mathbf{v}_{-i})$.

Proof. Let $k^* = k^*(\mathbf{v}_{-i})$, and $b^* = \frac{k^*}{k^*+1} p^M(\mathbf{v}_{-i})$. Suppose by contradiction that $p^M(b^*, \mathbf{v}_{-i}) > b^*$. This would imply that $p^M(b^*, \mathbf{v}_{-i}) = p^M(\mathbf{v}_{-i})$. But this is a contradiction since $\text{num}((b^*, \mathbf{v}_{-i}), b^*) \cdot b^* \geq (k^*+1) \cdot b^* = k^* \cdot p^M(\mathbf{v}_{-i})$. Thus, $p^M(b^*, \mathbf{v}_{-i}) \leq b^*$. Therefore, $p^S(\mathbf{v}_{-i}) = \min_b p^M(b, \mathbf{v}_{-i}) \leq p^M(b^*, \mathbf{v}_{-i}) \leq b^* < p^M(\mathbf{v}_{-i})$, which concludes the proof. \square

Corollary 1. Fix v_1, \dots, v_n and let $i^* \equiv \text{argmax}_{i=1, \dots, n} v_i$. Then,

$$p^S(\mathbf{v}_{-i^*}) < p^M(\mathbf{v}_{-i^*}) \leq p^M(\mathbf{v}).$$

This is an immediate consequence of Claims 4 and 3.

Claim 5. For any \mathbf{v} and any i , let $x = p^M(\mathbf{v})$. Then, $p^M(x, \mathbf{v}_{-i}) \leq x$

Proof. Suppose towards a contradiction that $y = p^M(x, \mathbf{v}_{-i}) > x$. Therefore $y \cdot \text{num}((x, \mathbf{v}_{-i}), y) > x \cdot \text{num}((x, \mathbf{v}_{-i}), x)$. Note that $\text{num}(\mathbf{v}, y) \geq \text{num}((x, \mathbf{v}_{-i}), y)$ and $\text{num}(\mathbf{v}, x) \leq \text{num}((x, \mathbf{v}_{-i}), x)$. This implies $y \cdot \text{num}(\mathbf{v}, y) \geq y \cdot \text{num}((x, \mathbf{v}_{-i}), y) > x \cdot \text{num}((x, \mathbf{v}_{-i}), x) \geq x \cdot \text{num}(\mathbf{v}, x)$. Therefore $y \cdot \text{num}(\mathbf{v}, y) > x \cdot \text{num}(\mathbf{v}, x)$ which contradicts the fact that $x = p^M(\mathbf{v})$. \square

Property 1. $\forall \mathbf{v}, i, p^M(p^S(\mathbf{v}_{-i}), \mathbf{v}_{-i}) = p^S(\mathbf{v}_{-i})$.

Proof. Recall the definition, $p^S(\mathbf{v}_{-i}) \equiv \min\{b \mid p^M(b, \mathbf{v}_{-i}) \leq b\}$. It immediately follows that $p^M(p^S(\mathbf{v}_{-i}), \mathbf{v}_{-i}) \leq p^S(\mathbf{v}_{-i})$. Assume towards a contradiction that the inequality is strict, and let $x = p^M(p^S(\mathbf{v}_{-i}), \mathbf{v}_{-i})$. But then Claim 5 implies that $p^M(x, \mathbf{v}_{-i}) \leq x$, which is a contradiction since $p^S(\mathbf{v}_{-i})$ is supposed to be the minimal such x . \square

Fixing any \mathbf{v} , recall that $i^* \equiv \text{argmax}_{i=1, \dots, n} v_i$, $k^* \equiv k^*(\mathbf{v}_{-i^*})$, $p^S \equiv p^S(\mathbf{v}_{-i^*})$.

Property 2. Let y be the smallest element in the support of F that is at least p^S . Then, $y = p^M(\mathbf{v})$ implies $p^S \geq \frac{k^*}{k^*+1} \cdot p^M(\mathbf{v})$.

Proof. By Corollary 1,

$$p^S < p^M(\mathbf{v}_{-i^*}) \leq p^M(v_{i^*}, \mathbf{v}_{-i^*}) = p^M(\mathbf{v}). \quad (9)$$

I.e., $y \leq p^M(\mathbf{v}_{-i^*}) \leq p^M(\mathbf{v})$ (because $p^M(\mathbf{v}_{-i^*})$ is in the support of F). Thus, since $p^M(\mathbf{v}) = y$,

$$p^M(\mathbf{v}_{-i^*}) = p^M(\mathbf{v}). \quad (10)$$

By definition of k^* , there exist k^* values in \mathbf{v}_{-i^*} that are at least $p^M(\mathbf{v}_{-i^*})$. Overall,

$$\text{num}(\mathbf{v}_{-i^*}, p^S) = \text{num}(\mathbf{v}_{-i^*}, y) = \text{num}(\mathbf{v}_{-i^*}, p^M(\mathbf{v})) = k^*. \quad (11)$$

Therefore, $R(\mathbf{v}_{-i^*}, p^S) = (k^* + 1)p^S$: this follows from Property 1, and Eq. (11). Thus, for any number z ,

$$(k^* + 1)p^S \geq z \cdot \text{num}((\mathbf{v}_{-i^*}, p^S), z) \geq z \cdot \text{num}(\mathbf{v}_{-i^*}, z). \quad (12)$$

Take $z \equiv p^M(\mathbf{v})$. Note that $\text{num}(\mathbf{v}_{-i^*}, p^M(\mathbf{v})) = k^*$ using Eq. (10). Using Eq. (12), $(k^* + 1)p^S \geq p^M(\mathbf{v}) \cdot k^*$, and the claim follows. \square

Claim 6. For any $\mathbf{v} = (v_1, \dots, v_n)$, and any v'_i such that $v_i \geq v'_i \geq p^M(\mathbf{v})$, $p^M(\mathbf{v}) \geq p^M(v'_i, \mathbf{v}_{-i})$.

We remark that $p^M(\mathbf{v})$ is not necessarily monotonically increasing as a function of v_i , for example, $p^M(2, 0) = 2 > 1 = p^M(2, 1)$. Hence, the importance of the condition $v'_i \geq p^M(\mathbf{v})$.

Proof. Denote $x = p^M(\mathbf{v})$ and $y = p^M(v'_i, \mathbf{v}_{-i})$. Assume towards a contradiction that $x < y$. Therefore

$$\begin{aligned} y \cdot \text{num}(\mathbf{v}, y) &\geq y \cdot \text{num}((v'_i, \mathbf{v}_{-i}), y) \\ &> x \cdot \text{num}((v'_i, \mathbf{v}_{-i}), x) = x \cdot \text{num}(\mathbf{v}, x) \end{aligned}$$

where the first step follows since $v_i \geq v'_i$, the second step follows since y is the monopolistic price for (v'_i, \mathbf{v}_{-i}) (the inequality is strict because $y > x$), and the third step follows since $v_i \geq v'_i \geq x$. However $y \cdot \text{num}(\mathbf{v}, y) > x \cdot \text{num}(\mathbf{v}, x)$ contradicts the fact that x is the monopolistic price for v . \square

Claim 7. For any $\mathbf{v} = (v_1, \dots, v_n)$, and any v'_i such that $v_i \geq v'_i \geq p^M(v'_i, \mathbf{v}_{-i})$, $p^M(\mathbf{v}) \geq p^M(v'_i, \mathbf{v}_{-i})$.

Proof. Assume towards a contradiction that $p^M(\mathbf{v}) < p^M(v'_i, \mathbf{v}_{-i})$. But then $v_i \geq v'_i \geq p^M(\mathbf{v})$ and Claim 6 implies that $p^M(\mathbf{v}) \geq p^M(v'_i, \mathbf{v}_{-i})$ which is a contradiction. \square

Claim 8. For any $\mathbf{v} = (v_1, \dots, v_n)$, and any i , if $v_i < p^M(\mathbf{v})$ then $v_i < p^S(\mathbf{v}_{-i})$.

Proof. Denote $x \equiv p^M(\mathbf{v})$ and $y \equiv p^S(\mathbf{v}_{-i}) = p^M(y, \mathbf{v}_{-i})$ (according to Property 1). Assume towards a contradiction that $x > v_i \geq y$. Since $x \equiv p^M(\mathbf{v})$ and our assumption $y < x$,

$$x \cdot \text{num}(\mathbf{v}, x) > y \cdot \text{num}(\mathbf{v}, y) \quad (13)$$

Furthermore, since $y = p^M(y, \mathbf{v}_{-i})$,

$$y \cdot \text{num}((y, \mathbf{v}_{-i}), y) \geq x \cdot \text{num}((y, \mathbf{v}_{-i}), x). \quad (14)$$

We also have

$$\text{num}(\mathbf{v}, y) = \text{num}((y, \mathbf{v}_{-i}), y), \quad (15)$$

since $v_i \geq y$, and similarly,

$$\text{num}(\mathbf{v}, x) = \text{num}((y, \mathbf{v}_{-i}), x), \quad (16)$$

since $v_i < x$ and $y < x$.

By plugging in Eqs. (15) and (16) in Eq. (14),

$$y \cdot \text{num}(\mathbf{v}, y) \geq x \cdot \text{num}(\mathbf{v}, x). \quad (17)$$

Combining Eqs. (13) and (17), we finally show that $x \cdot \text{num}(\mathbf{v}, x) > x \cdot \text{num}(\mathbf{v}, x)$, which proves the contradiction. \square

Claim 9. For any v_1, \dots, v_n and i, j , $v_i > v_j \geq p^S(\mathbf{v}_{-j})$ implies that $p^S(\mathbf{v}_{-j}) \geq p^S(\mathbf{v}_{-i})$.

We remark that the condition $v_j \geq p^S(\mathbf{v}_{-j})$ is important as without it the claim is not true. For example, take $\mathbf{v} = (2, 1, 0)$:

$$p^S(\mathbf{v}_{-2}) = p^S(2, 0) = 1 > \frac{2}{3} = p^S(2, 1) = p^S(\mathbf{v}_{-3}).$$

Proof. Let $b = p^S(\mathbf{v}_{-j})$. We will prove that $b \geq p^M(b, \mathbf{v}_{-i})$, which implies the claim since $p^S(\mathbf{v}_{-i}) \equiv \min\{b \mid b \geq p^M(b, \mathbf{v}_{-i})\}$.

By applying Claim 6 with respect to $\mathbf{u} = (v_1, \dots, v_{j-1}, b, v_{j+1}, \dots, v_n)$ and

$$\mathbf{u}' = (v_1, \dots, v_{i-1}, v_j, v_{i+1}, \dots, v_{j-1}, b, v_{j+1}, \dots, v_n)$$

implies that $p^M(\mathbf{u}) \geq p^M(\mathbf{u}')$ (note that $v_i \geq v_j \geq b \geq p^M(b, \mathbf{v}_{-j}) = p^M(\mathbf{u})$ where the last inequality follows from the definition of $p^S(\mathbf{v}_{-j})$; therefore $u_i \geq u'_i \geq p^M(\mathbf{u})$). Thus

$$b \geq p^M(\mathbf{u}) \geq p^M(\mathbf{u}') = p^M(b, \mathbf{v}_{-i}).$$

The last equality holds since the monopolistic price is oblivious to the identities of the users. The claim thus follows. \square

Property 3. $\forall \mathbf{v}, i, j$: $v_i \geq v_j$ implies $\delta_i(v_i, \mathbf{v}_{-i}) \geq \delta_j(v_j, \mathbf{v}_{-j})$.

Proof. If $v_j < p^S(\mathbf{v}_{-j})$, $\delta_j(v_j, \mathbf{v}_{-j}) = 0$ and the claim follows. Thus assume $v_j \geq p^S(\mathbf{v}_{-j})$, and we have $v_i > v_j \geq p^S(\mathbf{v}_{-j}) \geq p^S(\mathbf{v}_{-i})$ where the last inequality follows from Claim 9. Since p^M is the same in both terms, and $p^S(\mathbf{v}_{-j}) \geq p^S(\mathbf{v}_{-i})$, we have $\delta_i(v_i, \mathbf{v}_{-i}) \geq \delta_j(v_j, \mathbf{v}_{-j})$. \square

Claim 10. For any \mathbf{v}_{-i} and $v_i > v'_i$, $\delta_i(v_i, \mathbf{v}_{-i}) \geq \delta_i(v'_i, \mathbf{v}_{-i})$.

Proof. If $v'_i < p^S(\mathbf{v}_{-i})$, $\delta_i(v'_i, \mathbf{v}_{-i}) = 0$ and the claim follows. Thus assume $v'_i \geq p^S(\mathbf{v}_{-i})$. Assume towards a contradiction that $v'_i < p^M(v'_i, \mathbf{v}_{-i})$. Claim 8 then implies that $v'_i < p^S(\mathbf{v}_{-i})$ which is a contradiction to our assumption above. Thus, $v_i > v'_i \geq p^M(v'_i, \mathbf{v}_{-i})$. Claim 7 then implies that $p^M(v_i, \mathbf{v}_{-i}) \geq p^M(v'_i, \mathbf{v}_{-i})$. By the definition of the discount ratio we have $\delta_i(v_i, \mathbf{v}_{-i}) \geq \delta_i(v'_i, \mathbf{v}_{-i})$ as claimed. \square

C Additional Auxiliary Claims for the Multi-Bid Case

Claim 11. For any \mathbf{b}_{-i} and $u \in \mathbb{N}^+$, define⁷

$$p_u^{\text{multibid}}(\mathbf{b}_{-i}) = \min\{u \cdot b_i \mid b_i \in \mathbb{R}, b_i \geq p^M(\overbrace{b_i, \dots, b_i}^{u \text{ times}}, \mathbf{b}_{-i})\}.$$

Then, for any $u \geq n + 1$, $p_u^{\text{multibid}}(\mathbf{b}_{-i}) > p_1^{\text{multibid}}(\mathbf{b}_{-i})$.

Proof. Let $b = \frac{p_u^{\text{multibid}}(\mathbf{b}_{-i})}{u}$. Then, $b \cdot (u + \text{num}(\mathbf{b}_{-i}, b)) \geq k^*(\mathbf{b}_{-i}) \cdot p^M(\mathbf{b}_{-i})$. Therefore,

$$p_u^{\text{multibid}}(\mathbf{b}_{-i}) = b \cdot u \geq \frac{k^*(\mathbf{b}_{-i}) \cdot p^M(\mathbf{b}_{-i})}{u + \text{num}(\mathbf{b}_{-i}, b)} \cdot u \quad (18)$$

If $k^*(\mathbf{b}_{-i}) \geq 2$ then $\frac{u}{u + \text{num}(\mathbf{b}_{-i}, b)} \cdot k^*(\mathbf{b}_{-i}) > 1$ since $u > n > \text{num}(\mathbf{b}_{-i}, b)$. Therefore Eq. (18) implies that $p_u^{\text{multibid}}(\mathbf{b}_{-i}) > p^M(\mathbf{b}_{-i})$. Using Claim 4, $p^M(\mathbf{b}_{-i}) > p^S(\mathbf{b}_{-i})$, and since $p^S(\mathbf{b}_{-i}) = p_1^{\text{multibid}}(\mathbf{b}_{-i})$, we complete the proof in the case $k^*(\mathbf{b}_{-i}) \geq 2$.

If $k^*(\mathbf{b}_{-i}) = 1$ then $\frac{u}{u + \text{num}(\mathbf{b}_{-i}, b)} \cdot k^*(\mathbf{b}_{-i}) > \frac{1}{2}$. Therefore Eq. (18) implies that $p_u^{\text{multibid}}(\mathbf{b}_{-i}) > \frac{p^M(\mathbf{b}_{-i})}{2} \geq p^S(\mathbf{b}_{-i}) = p_1^{\text{multibid}}(\mathbf{b}_{-i})$, implying the claim. (Here, in the second inequality we used the fact that $k^*(\mathbf{b}_{-i}) = 1$, and bidding half of the highest bid will win in this case.) \square

⁷The minimum is well defined – see the argument next to Eq. (6).

Claim 12. Fix any \mathbf{b}_{-i} and let b, u be the arguments which minimize Eq. (6). Then,

$$b = p^M(\overbrace{b, \dots, b}^{u \text{ times}}, \mathbf{b}_{-i}).$$

Proof. By definition $b \geq p^M(\overbrace{b, \dots, b}^{u \text{ times}}, \mathbf{b}_{-i})$. Suppose towards a contradiction that the inequality is strict, and take some b' such that

$$b > b' \geq p^M(\overbrace{b', \dots, b'}^{u \text{ times}}, \mathbf{b}_{-i}).$$

Claim 6 implies that

$$b' \geq p^M(\overbrace{b', \dots, b'}^{u \text{ times}}, \mathbf{b}_{-i}),$$

contradicting the minimality of (b, u) since $b \cdot u > b' \cdot u$. \square

Claim 13. Fix any \mathbf{b}_{-i} and let (b, u) be the arguments which minimize Eq. (6). Then,

$$b \leq p^M(\mathbf{b}_{-i}).$$

Proof. Let $b' = p^M(\mathbf{b}_{-i})$. Then, the second part of Claim 3 implies that $b' = p^M(b', \mathbf{b}_{-i})$. Thus $(b', u' = 1)$ satisfies the requirement of Eq. (6). Hence $b \cdot u \leq b'$ and the claim follows. \square

D Properties of the binomial and trinomial distributions

Claim 14. Let $X \sim B(n, p)$ be a binomial random variable where $0 < p < 1$. For every i , $\Pr(X = i) = O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. If np is an integer, the mode (most likely value of the distribution) is its mean np .⁸ The case where np is not an integer can be handled in a similar manner. Therefore, our goal

⁸This can be shown directly by computing $\frac{\Pr(X=i)}{\Pr(X=i+1)}$.

is to show that $\Pr(X = np) = O\left(\frac{1}{\sqrt{n}}\right)$. By using Stirling's approximation⁹ $n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$,

$$\begin{aligned} \Pr(X = np) &= \binom{n}{np} p^{np} (1-p)^{n(1-p)} \\ &\approx \frac{\sqrt{2\pi n} \left(\frac{n}{e}\right)^n}{\sqrt{2\pi np} \left(\frac{np}{e}\right)^{np} \sqrt{2\pi n(1-p)} \left(\frac{n(1-p)}{e}\right)^{n(1-p)}} p^{np} (1-p)^{n(1-p)} \\ &= \frac{1}{\sqrt{2\pi np(1-p)}} = O\left(\frac{1}{\sqrt{n}}\right). \end{aligned}$$

Of course, this is only valid when p is treated as a constant. □

Claim 15. Let $(X_1, X_2, X_3) \sim \text{Trinomial}(n, p_1, p_2, p_3)$ be a triple of trinomial random variables where $p_1, p_2 > 0$. For every function f , $\Pr(X_1 = f(X_1 + X_2)) = O\left(\frac{1}{\sqrt{n}}\right)$.

Recall that the trinomial distribution is the multinomial distribution restricted to the case $k = 3$.

Proof. Our first goal is to prove a very weak bound $\Pr(X_3 \geq \alpha n) = O\left(\frac{1}{\sqrt{n}}\right)$ for $\alpha = \frac{p_3+1}{2}$. Since this bound is so weak it could be proved using the Chernoff bound, or even Chebyshev's inequality.

Lemma 4 (Chebyshev's inequality). For every $k \in \mathbb{R}$, and every random variable X with expectation μ and standard deviation σ ,

$$\Pr(|X - \mu| \geq k\sigma) \leq \frac{1}{k^2} \tag{19}$$

Recall that $X_3 \sim \text{Binomial}(n, p_3)$, and therefore, $E[X] = np_3$ and has a standard deviation $\sqrt{np_3(1-p_3)}$. By observing that $\Pr(X - \mu \geq k\sigma) \leq \Pr(|X - \mu| \geq k\sigma)$ and plugging in $k = \frac{\sqrt{n(1-p_3)}}{2\sqrt{p_3}}$ in Lemma 4,

$$\Pr(X_3 \geq \alpha n) \leq \frac{4p_3}{n(1-p_3)} = O\left(\frac{1}{\sqrt{n}}\right), \tag{20}$$

where in the equality we use the fact that $p_3 < 1$ (since $p_1, p_2 > 0$ and $p_1 + p_2 + p_3 = 1$).

⁹One could take into account the error term, i.e., $n! = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n (1 + O(\frac{1}{n}))$, and reach the same conclusion.

Conditioned that $X_3 = z$, X_1 has a Binomial distribution: $X_1 \sim B(n - z, \frac{p_1}{1-p_3})$ (here we use the condition $p_1, p_2 > 0$ to conclude that $0 < \frac{p_1}{1-p_3} < 1$). Therefore,

$$\begin{aligned}
\Pr(X_1 = f(X_1 + X_2)) &= \sum_{z=0}^n \Pr(X_1 = f(X_1 + X_2) | X_3 = z) \Pr(X_3 = z) \\
&= \sum_{z=0}^n \Pr(X_1 = f(n - z) | X_3 = z) \Pr(X_3 = z) \\
&\leq \sum_{z=0}^{\lfloor \alpha n \rfloor} \Pr(X_1 = f(n - z) | X_3 = z) \Pr(X_3 = z) + \sum_{z=\lfloor \alpha n \rfloor + 1}^n \Pr(X_3 = z) \\
&\leq \sum_{z=0}^{\lfloor \alpha n \rfloor} \Pr(X_1 = f(n - z) | X_3 = z) \Pr(X_3 = z) + O\left(\frac{1}{\sqrt{n}}\right) \\
&\leq \sum_{z=0}^{\lfloor \alpha n \rfloor} O\left(\frac{1}{\sqrt{n - z}}\right) \Pr(X_3 = z) + O\left(\frac{1}{\sqrt{n}}\right) \\
&\leq \sum_{z=0}^{\lfloor \alpha n \rfloor} O\left(\frac{1}{\sqrt{(1 - \alpha)n}}\right) \Pr(X_3 = z) + O\left(\frac{1}{\sqrt{n}}\right) = O\left(\frac{1}{\sqrt{n}}\right),
\end{aligned}$$

where in the second inequality we use Eq. (20), and Claim 14 in the third inequality. \square

Claim 16. *Let F be any distribution with a finite support size and $\mathbf{v} = (v_1, \dots, v_n)$ be n i.i.d. draws from F . Define $\text{num}(\mathbf{v}, z) \equiv |\{v_i | v_i \geq z\}|$ and random variables $n_z = \text{num}(\mathbf{v}, z)$ for any real number z . Then, for any arbitrary $x, y \in \text{Support}(F)$ such that $x > y$,*

$$\lim_{n \rightarrow \infty} \Pr(n_x \cdot x > n_y \cdot y \geq (n_x - 1)x) = 0.$$

Proof. The term $n_x \cdot x > n_y \cdot y \geq (n_x - 1)x$ is the same as $n_x > \frac{y}{x}n_y \geq n_x - 1$, which is true if and only if $n_x = \lfloor \frac{y}{x}n_y + 1 \rfloor$. The triple $(n_x, n_y - n_x, n - n_y)$ is a trinomial distribution with probabilities $p_1 = \Pr_{v_i \sim F}(v_i \geq x)$, $p_2 = \Pr_{v_i \sim F}(x > v_i \geq y)$, $p_3 = \Pr_{v_i \sim F}(v_i < y)$ (p_i depends on F but not on n). Denoting $f(n_y) = \lfloor \frac{y}{x}n_y + 1 \rfloor$, we conclude that $\Pr(n_x = f(n_y)) = \Pr(n_x \cdot x > n_y \cdot y \geq (n_x - 1)x)$. Claim 15 therefore implies the current claim. \square

E Proof of Theorem 3

We first show a lemma:

Lemma 5. *For every $v_1, b_1, \mathbf{v}_{-1}$,*

$$p^M(v_{\max}, \mathbf{v}_{-1}) - p^S(\mathbf{v}_{-1}) \geq u_1(v_1, b_1, \mathbf{v}_{-1}) - u_1(v_1, v_1, \mathbf{v}_{-1})$$

Proof. Assume first that $v_1 \geq p^M(v_1, \mathbf{v}_{-1})$. Then:

$$\begin{aligned}
p^M(v_{max}, \mathbf{v}_{-1}) - p^S(\mathbf{v}_{-1}) &\geq p^M(v_1, \mathbf{v}_{-1}) - p^S(\mathbf{v}_{-1}) \\
&= [v_1 - p^S(\mathbf{v}_{-1})] - [v_1 - p^M(v_1, \mathbf{v}_{-1})] \\
&= u_1(v_1, p^S(\mathbf{v}_{-1}), \mathbf{v}_{-1}) - u_1(v_1, v_1, \mathbf{v}_{-1}) \\
&\geq u_1(v_1, b_1, \mathbf{v}_{-1}) - u_1(v_1, v_1, \mathbf{v}_{-1})
\end{aligned}$$

where the first transition follows by Claim 7 and the third transition follows by Property 1. Therefore if $v_1 \geq p^M(v_1, \mathbf{v}_{-1})$ the claim holds. Now assume $v_1 < p^M(v_1, \mathbf{v}_{-1})$. In this case, $u_1(v_1, v_1, \mathbf{v}_{-1}) = 0$ and $u_1(v_1, b_1, \mathbf{v}_{-1}) \leq 0$ (since by Claim 8 in this case $v_1 < p^S(\mathbf{v}_{-1})$). Thus, $p^M(v_{max}, \mathbf{v}_{-1}) - p^S(\mathbf{v}_{-1}) \geq 0 \geq u_1(v_1, b_1, \mathbf{v}_{-1}) - u_1(v_1, v_1, \mathbf{v}_{-1})$ hence the claim holds in this case as well. \square

We now prove the theorem:

$$\begin{aligned}
\Delta_n^{GT} &= \mathbb{E}_{\mathbf{v}_{-1} \sim F} \left[\max_{v_1 \in \text{Support}(F)} \delta_1(v_1, \mathbf{v}_{-1}) \right] \\
&= \mathbb{E}_{\mathbf{v}_{-1} \sim F} [\delta_1(v_{max}, \mathbf{v}_{-1})] \\
&= \mathbb{E}_{\mathbf{v}_{-1} \sim F} \left[1 - \frac{p^S(\mathbf{v}_{-1})}{p^M(v_{max}, \mathbf{v}_{-1})} \right] \\
&\geq \frac{1}{v_{max}} \mathbb{E}_{\mathbf{v}_{-1} \sim F} [p^M(v_{max}, \mathbf{v}_{-1}) - p^S(\mathbf{v}_{-1})] \\
&\geq \max_{v_1, b_1 \in \text{Support}(F)} \left\{ \frac{1}{v_{max}} \mathbb{E}_{\mathbf{v}_{-1} \sim F} [u_1(v_1, b_1, \mathbf{v}_{-1}) - u_1(v_1, v_1, \mathbf{v}_{-1})] \right\}.
\end{aligned}$$

Here the second step follows from Claim 10, the third step follows since a user with the highest possible valuation v_{max} will necessarily win, the fourth step follows from $p^M(v_{max}, \mathbf{v}_{-1}) \leq v_{max}$, and the fifth step follows from Lemma 5. Eq. (5) directly follows for $\epsilon = v^{max} \cdot \Delta_n^{GT}$, which completes the proof of the theorem.

F Analysis of Example 6

Claim 17.

$$\lim_{n \rightarrow \infty} \frac{RSOP(\overbrace{2, \dots, 2}^{n \text{ times}}, \overbrace{1, \dots, 1}^{n \text{ times}}) - 2n}{\sqrt{n}} = -\frac{1}{\sqrt{\pi}} \quad (21)$$

Proof. Let X_i be the indicator variable whether the i 'th 2-bidder is in group A , and similarly Y_i for the i 'th 1-bidder. Let $Z_i = X_i - Y_i$. Let $X = \sum_{i=1}^n X_i$, $Y = \sum_{i=1}^n Y_i$ and $Z = \sum_{i=1}^n Z_i$.

Note that the revenue for any realization of the coin toss is exactly $2n - |Z|$. Suppose first that $Z \geq 0$, i.e., $X \geq Y$. In this case, the price offered to set A (which is the monopolistic

price computed on set B) is 1, and similarly the price offered to set B is 2. Thus, the revenue is $1 \cdot (X + Y) + 2 \cdot (n - X) = 2n - Z$. Similarly, if $Z < 0$, the revenue is $2n + Z$. This shows that in all cases, the revenue is $2n - |Z|$.

Since $\mathbb{E}[Z_i] = 0$ and $\text{Var}(Z_i) = \frac{1}{2}$, by the central limit theorem, $\frac{1}{\sqrt{n}}Z$ converges in distribution to $N(0, \sigma^2 = \frac{1}{2})$ as n goes to infinity. Recall that for a normal random variable $X' \sim N(0, \sigma)$, the distribution of $Y' = |X'|$ is called half-normal, and $E[Y'] = \frac{\sigma\sqrt{2}}{\sqrt{\pi}}$.

Therefore, $E[\frac{|Z|}{\sqrt{n}}] = \frac{\sigma\sqrt{2}}{\sqrt{\pi}} = \frac{1}{\sqrt{\pi}}$. This completes the proof, since $RSOP(\overbrace{2, \dots, 2}^{n \text{ times}}, \overbrace{1, \dots, 1}^{n \text{ times}}) = 2n - \mathbb{E}[|Z|]$. \square

An immediate consequence of Claim 17 is the following:

Corollary 2. $RSOP(\overbrace{2, \dots, 2}^{n \text{ times}}, \overbrace{1, \dots, 1}^{n \text{ times}}) = 2n - \sqrt{\frac{n}{\pi}} + o(\sqrt{n})$.

G Desiderata

In this appendix we provide the desired properties of a fee mechanism in Bitcoin. We note that some of the properties mentioned below may conflict with others, and it may well be that no single mechanism provides them simultaneously.

High social welfare The sum of valuations of accepted transactions should be large. In this sense, the system is providing a high level of utility to its users. In particular, an efficient allocation is warranted: transactions with higher values should be included before those with lower values (assuming both are available at the same time) as this clearly improves the social welfare.

Revenue extraction The amount of money transferred to the miners is high, which would buy more security for the system. Stated another way: the revenue of miners should be part of the social welfare measure of the system.

Truthful bidding An ideal mechanism would allow users to state their preferences clearly and would encourage honest reporting. The main advantage would be that users do not need to “overthink” how they should act. Furthermore, this removes computational burdens associated with such strategies: there is no need to monitor the network to obtain statistics on past bids, perform estimates of the threshold bid needed to get accepted into a block, and no need to perform bid updates for the transaction. Note that truthfulness should imply that various, more technical, manipulations, should not be profitable. For example, such manipulations include splitting a single transaction into two transfers of smaller amounts, adding more transactions between two bitcoin addresses of the same client, etc. A desired property that such manipulation would not be profitable.

Resistance to manipulation by miners The mechanism should be resistant to selfish behavior by the miners. Such behavior can include miners adding transactions of their own into their own blocks, miners withholding transactions and selecting other sets of transactions, etc.

Resistance to manipulation via side payments Some mechanisms may appear well at first sight, but in fact may encourage miners and users to circumvent the fee system altogether. For example, if a miner has to give half of the revenue from his block to the miner that creates the next block in the chain, then he can offer users a deal: The miner will include transactions with 0 fees in his block in exchange for a side payment that will be given to him via a separate and direct transaction. In this manner, he does not need to share the rewards with others. Of course, since Bitcoin provides a natural payment channel from the users to the miners, the mechanism should be resistant to such side payments.

Adaptivity to changing demand, network conditions, block sizes, etc. It is important to avoid hard-coded magic numbers (such as hard-coded minimal fees) in the protocol as much as possible. Given that the protocol is hard to update, as it requires wide adoption of new code, any hard-coded number is difficult to adjust. A minimal fee, for example, which may need adjusting from time to time (e.g., when the exchange rate fluctuates, or demand increases) would insert inefficiency when it is not set in accordance with market conditions.

Accounting for temporal considerations Users may have different levels of urgency for their transactions. A good fee mechanism will take into account the willingness of users to delay the acceptance of their transactions, e.g., in exchange for paying lower fees. In Bitcoin, transactions that were not added to the blockchain persist with the same bid and may be included in other blocks (without a discount). Our own analysis in this paper considers only impatient users who only desire that their transactions will be included in the next block. We believe that such temporal considerations are an important issue to tackle in future work.

Consensus based mechanism One of the key features of cryptocurrencies such as Bitcoin is that it *can* be assumed that data that is part of the longest blockchain is agreed upon all miners and users. Therefore, the mechanism could make decisions based on previous blocks. On the other hand, there is no consensus between the miners regarding data which is not part of the blockchain. For example, different miners might have different views regarding the mempool (which is the set of transaction waiting to be included in the next block). In particular, it is impossible to take the statistics of *all* bids in previous rounds into account, since the bids that were not included in the block are not part of the consensus.

H Nomenclature

\mathbf{b} The bid vector $\mathbf{b} = (b_1, \dots, b_n)$ satisfying $b_1 \geq \dots \geq b_n$, page 5

$\Delta_n^{average}$	$= \mathbb{E}_{(v_1, \dots, v_n) \sim F} [\delta_1(v_1, \mathbf{v}_{-1})]$, page 6
$\delta_i(v_i, \mathbf{b}_{-i})$	The discount ratio, page 6
$\delta_{max}(\mathbf{v})$	$\delta_{max}(\mathbf{v}) \equiv \max_i \delta_i(v_i, \mathbf{v}_{-i})$, page 7
Δ_n^{GT}	$\Delta_n^{GT} \equiv \mathbb{E}_{(v_2, \dots, v_n) \sim F} [\delta_n^{GT}(\mathbf{v}_{-1})]$, page 8
Δ_n^{max}	$\Delta_n^{max} \equiv \mathbb{E}_{(v_1, \dots, v_n) \sim F} [\delta_{max}(\mathbf{v})]$, page 7
Δ_n^{RSOP}	$\Delta_n^{RSOP} \equiv \mathbb{E}_{(v_1, \dots, v_n) \sim F} \left[\frac{R(\mathbf{v})}{RSOP(\mathbf{v})} - 1 \right]$, page 20
ϵ -BNE	An ϵ -Bayesian Nash Equilibrium, page 9
F	All true values are drawn i.i.d. from some distribution F on $\mathbb{R}_{>0}$, page 6
inverse distribution	$F(x) = 1 - \frac{1}{x}$, $x \in [1, \infty)$, page 14
i^*	$i^* \equiv \operatorname{argmax}_{i=1, \dots, n} v_i$, page 10
$k^*(\mathbf{b})$	The index k that maximizes $k \cdot b_k$. In case of ties, k^* is chosen to be the maximal such index, page 5
k^*	$k^* \equiv k^*(\mathbf{v}_{-i^*})$, page 10
$p^M(\mathbf{b})$	The monopolistic price, page 5
$p^{multibid}(\mathbf{b}_{-i})$	A generalization of p^S to capture the possible benefit of splitting a single bid to u different bids, page 12
p^S	$p^S \equiv p^S(\mathbf{v}_{-i^*})$, page 10
$p^S(\mathbf{b}_{-i})$	The strategic price: The lowest possible bid for i to be included in the block, when the other bids are \mathbf{b}_{-i} , page 6
$R(\mathbf{b})$	The monopolistic revenue, $R(\mathbf{b}) \equiv k^*(\mathbf{b}) \cdot b_{k^*(\mathbf{b})}$, page 5
RSOP	The Random Sampling Optimal Price Auction, Definition 1, page 17
v_{max}	$v_{max} \equiv \max \operatorname{Support}(F)$, page 8