

Beyond Winner-Take-All Procurement Auctions

Pranav Garimidi*, Michael Neuder†, and Tim Roughgarden‡

Abstract. Blockchain protocols often seek to procure computationally challenging work from a decentralized set of participants. While there are simple procurement auctions that result in the minimal cost of acquisition and maximal efficiency, they also lead to concentration in the provider set due to the winner-take-all market structure. We design and analyze single-good procurement auctions that balance social-cost minimization (at the extreme, a winner-take-all auction) with decentralization (at the extreme, a uniform allocation). We first give a dominant-strategy incentive-compatible (DSIC) mechanism explicitly designed to implement non-winner-take-all allocations. Our allocation rule uniquely solves an optimization with respect to a modified social-cost metric that penalizes large, single-player concentrations and is parameterized with a curvature value, α , with $\alpha \rightarrow 0$ implementing the uniform allocation and $\alpha \rightarrow \infty$ implementing the winner-take-all allocation. We further quantify the loss in social cost of this mechanism as a function of α .

We then propose two alternative mechanisms, each addressing a limitation of the DSIC mechanism, namely a lack of Sybil-resistance and a complex payment rule. First, we examine a variation of Tullock contests to achieve a non-winner-take-all Sybil-proof procurement mechanism. Second, we consider a mechanism with the same allocation rule as the DSIC mechanism but with an alternative payment rule in which producers are simply paid proportionally to their bids. This provides a much simpler payment rule which, while not DSIC, still results in the mechanism being ex-post “safe” (where there exists a bidding strategy that is guaranteed to result in non-negative utility) for participating bidders. For both non-DSIC mechanisms, we characterize the equilibrium allocations and prove price of anarchy bounds.

1 Introduction

Many blockchain-related protocols face *procurement problems*, in which the protocol must secure services from a permissionless set of providers. For example, layer-one blockchain protocols require validators to secure the network, prover marketplaces gather provers to fulfill requests for succinct proofs, and decentralized inference networks recruit GPU providers to carry out queries. Procurement problems are well-studied in the traditional mechanism design literature, where

* a16z crypto. Email: pgarimidi@a16z.com.

† Princeton University. Author supported in part by the Ethereum Foundation: Grant FY25-2276. Email: michael.neuder@princeton.edu.

‡ Columbia University & a16z crypto. Author’s research at Columbia University supported in part by NSF awards CCF-2006737 and CNS-2212745. Email: tim.roughgarden@gmail.com.

the focus has generally been on *efficient* mechanisms that seek to allocate work to parties that can complete the work at the lowest cost. Procurement mechanisms that focus solely on minimizing costs lead to *winner-take-all* outcomes, where the single most efficient service provider consistently wins all the work.

Blockchain-related applications, however, often additionally require a “sufficiently decentralized” provider set. For example, a typical PBFT-style consensus protocol requires a sufficiently large number of distinct validators in order to be confident that fewer than one-third of them are faulty. Similarly, a prover or inference network requires a sufficient number of active providers to ensure that a single one cannot cause a liveness failure for the protocol’s consumers.

Winner-take-all procurement encourages centralization, with all but the most efficient provider leaving the market or consolidating with the most efficient one. An end state with a single provider would threaten the security assumptions of many blockchain-related applications, and more generally, can be economically undesirable for several reasons:

- Once a single service provider has crowded out their competition, it can price-gouge by charging monopoly prices.
- Having only a single service provider makes the procurer vulnerable to the provider failing for any number of reasons.
- With no competition, there are no incentives for a single service provider to make investments in decreasing their cost.
- In case the type of work demanded changes in the future, it is favorable to have multiple providers with different specialties available.
- Once a single service provider dominates the market, it can undercut costs whenever a new participant attempts to enter the market, defending their monopoly.

In practice, blockchain-related protocols are often designed to maintain a sufficiently decentralized set of providers. For instance, in addition to acting as Sybil-resistance mechanisms, proof-of-work and proof-of-stake convert earning consensus rewards into a Tullock contest, a mechanism in which many nodes participate at equilibrium instead of only the most efficient one [1]. Similarly, in prover markets, proposals have been made to explicitly trade off efficiency for a more robust set of providers [19]. These examples demonstrate the importance of understanding the design space of procurement mechanisms that go beyond pure cost-minimization. As blockchain protocols and their applications expand to new contexts, these mechanisms will become increasingly important. This work initiates the exploration of this design space through the study of three natural mechanisms and their properties, detailed in the next section.

1.1 Summary of results

This work is an applied modeling contribution that serves to motivate and discuss trade-offs in the study of procurement in blockchain settings. We model the procurement game and define a set of natural properties that practitioners may

find desirable, along with giving mechanisms for achieving different subsets of these properties. We study three mechanisms under slightly different models, which we discuss in more detail below, but summarize here for clarity. We start by considering an incomplete information setting where bidders' costs are private and examine a truthful mechanism that explicitly implements a non-winner-take-all allocation. This mechanism, however, has a couple of clear downsides: the payment rule is both difficult for bidders to interpret and for the protocol to implement, and is not Sybil-proof. Additionally, this mechanism can result in arbitrarily high payments for the protocol, but we give a modification that allows the protocol to bound its costs with the downside of a potentially more centralized allocation. To address these downsides, we consider two additional mechanisms. Both mechanisms are non-truthful, and thus we move from the incomplete to the complete information setting, where bidders' costs are common knowledge to analyze the mechanisms' pure Nash equilibria. The first of these mechanisms is essentially a Tullock contest which is Sybil-proof and allows the protocol to control the total payments it gives out. However, this mechanism may still be difficult for bidders to participate in and may even leave bidders with negative utilities for participating. The second mechanism we consider uses the same allocation rule as the truthful mechanism but with a paid-as-bid payment rule, where bidders are paid exactly proportional to their bids and the amount of work they were allocated. This mechanism is comparatively easier for bidders to reason about, with the downside of not being Sybil-proof and potentially resulting in arbitrarily high payments. A reserve price can be used to cap the total payments, but we only analyze the equilibria without a reserve.

We view the present work as an important first step in discussing the trade-offs and desiderata of decentralized procurement mechanisms, while stopping short of proposing a definitive procurement mechanism for protocols to use. The mechanisms we introduce are analyzed under different informational assumptions and thus not directly comparable. Nevertheless, each of the three mechanisms captures important features of interest to protocol designers and can serve as building blocks for the future. We leave a unified model and a characterization of an optimal mechanism for this setting as important future work.

Section 3 introduces α -proportional allocation rules (α -PARs) which implement allocations minimizing an α -scaled social cost metric that penalizes allocations with large single-player concentrations, explicitly optimizing for non-winner-take-all allocations. These allocation rules are monotone and thus truthfully implementable via the corresponding Myersonian payment rules. We characterize the trade-off between social cost (where winner-take-all, implemented by taking $\alpha \rightarrow \infty$, is maximally efficient) and “decentralization” (where the uniform allocation rule, implemented by taking $\alpha \rightarrow 0$, is maximally decentralized) for these mechanisms by showing the social cost scales in the number of players, n , and the protocol-selected α parameter as $1 + (n/\alpha)^{1/\alpha}$. This DSIC mechanism has two principal drawbacks that, depending on the details of the application, motivate considering alternative procurement mechanisms. First, the mechanism is not Sybil-proof. Second, its payment rule—the unique payment rule that can

be coupled with an α -PAR to produce a DSIC mechanism—is complex to implement and hard to interpret for bidders.

Section 4 presents a Sybil-proof mechanism that we call a Tullock procurement contest. We shift focus here to the complete-information setting (the standard setting for the analysis of Tullock contests), in which each supplier’s cost of production is common knowledge. This assumption is reasonable in the context of mechanisms that are run repeatedly with public bids (as in the case of many blockchain-related applications), as market participants may learn each other’s costs over time. (See [5] for a more extended discussion of this point in the context of the complete-information analysis of generalized second-price auctions for online advertising.) Here, we prove the existence and uniqueness of PNE in the game induced by a Tullock procurement contest. We then bound the social cost of the PNE, and show that as the auctioneer increases their budget of rewards, the social cost at equilibrium increases, but the allocation becomes more evenly distributed across providers. In particular, as the reward budget grows increasingly large, with n competing bidders, their allocations converge to $1/n$.

Section 5 steps back and instead uses a simple to implement and interpret “paid-as-bid” payment rule with the same α -PAR allocations as the DSIC mechanism. Beyond simplicity, any participant who bids at least their true cost is guaranteed non-negative utility, no matter what the other participants bid (we call this property *ex-post safety*). Unfortunately, this mechanism is not Sybil-proof. We prove the existence and uniqueness of PNE for this mechanism. Further, we implicitly characterize the equilibrium allocations and the associated worst-case cost vectors. We parameterize this class of mechanisms by α , where as α grows increasingly large, the mechanism concentrates an increasing amount of the allocation between the lowest-cost bidders. We show that at equilibrium, every agent’s allocation is upper bounded by $1 - 1/\alpha$. Additionally, we provide numerical results demonstrating how the price of anarchy scales with the number of players and the disparity between their respective costs. For example, we show that with 16 participants that are no worse than 100% more costly than the best prover, the price of anarchy of the $\alpha = 4$ mechanism is no more than 1.4 (i.e., a 40% degradation in the social cost). Section 6 concludes and outlines future work.

1.2 Related Work

Procurement auctions. Procurement or “reverse” auctions have a rich literature, finding numerous applications in practice. While the breadth of topics is too deep to cover here, see [4] for a full treatment, we focus on settings where the auctioneer sacrifices minimizing their cost to ensure more equitable outcomes. [14] considers a setting similar to ours, where they model the cost to the auctioneer of agents leaving the auction if they are not allocated enough work. To address this, they study a threshold allocation rule where the lowest k producer bids are each allocated exactly $1/k$ of the production rights. They prove that setting a lower threshold (thus effectively increasing k) can lead to higher revenue in the repeated game despite not being revenue maximizing in the one-shot

game (where $k = 1$ would be optimal). [3] studies multi-dimensional procurement auctions, where the price and quality of the produced good are considered. This opened a line of literature considering non-price features when deciding how to award government contracts and was grounded in how the Department of Defense handled contracts and negotiations. [12] extend this to consider that the US government restricts entry in procurement auctions for many reasons: (i) regulatory restrictions on the producers (e.g., domestic vs. foreign), (ii) small business concerns (e.g., promoting and supporting allocation to businesses that otherwise wouldn't be competitive), and (iii) discretionary reasons, a catchall. The “small business concerns” are most relevant to our work and were also studied explicitly in [17], which demonstrated that small business participation in procurement auctions would be reduced by 40% without a significant portion of the budget being set aside for them. Additionally, they show that the procurement cost would increase significantly without the increased competition facilitated by the small business subsidy.

Supply-side bidding and supply function equilibria. Our work relates to the operations research literature of firms competing to produce some supply of a homogeneous good, also known as Cournot competition. Most notably, [15] generalizes Cournot’s model to allow producers to specify a more general class of “supply functions,” which specify the amount of output the firm will produce as a function of price. They characterize the equilibria of these supply functions under uncertainty of the realized demand shocks. In other words, they answer the question, “How do firms with identical costs but uncertainty around demand determine how to report their supply functions in equilibria?” [2] studies supply function equilibria when limiting the producers to linear supply functions, while [13] limits even further to only consider scalar supply functions. Both of these models are similar to ours in that the reported values resemble the per-unit cost of production that the producers strategize over, but our allocation rule is distinct.

[8] and [10] are the most relevant to our work. These works examine supply-side bidding in the “inelastic demand” setting, where the demand is relatively fixed, and the producers have asymmetric costs. [8] points out that the simple mechanism that pays each bidder exactly their bid and allocates the resource proportionally (exactly the reverse of the Tullock contest) cannot have a Nash equilibrium, as the bids can grow to infinity while each firm’s respective cost remains bounded, and uses this to motivate a constrained set of supply functions. They then solve for the optimal allocation if the firms are price takers, along with the Nash equilibrium when the firms are price anticipating, thus calculating the efficiency loss of the strategic setting. Our work uses the same problem statement but differs in our allocation and payment rules.

2 Model

We consider a standard procurement setting where there is a divisible unit of work to be assigned to suppliers. There are n competing agents, indexed by i ,

where each agent incurs a cost linear in their allocation parameterized by c_i . This cost can be interpreted as the resource cost agent i has to pay to complete a given amount of work. Let $\mathbf{c} = (c_1, c_2, \dots, c_n)$ denote the vector of agent costs. We consider both the *incomplete-information* case, where costs are only privately known with competing bidders having no prior knowledge of the costs of their competitors, and the *complete-information* case, where \mathbf{c} is common knowledge to all bidders.

A procurement mechanism \mathcal{M} consists of an allocation and payment rule (x, p) . The allocation rule determines how the work is split among the bidders, and the payment rule dictates how much each bidder gets paid. Each agent i simultaneously submits a bid b_i , or potentially multiple bids, to the mechanism. We denote the bid vector $\mathbf{b} = (b_1, b_2, \dots, b_n)$ and use \mathbf{b}_{-i} to denote the bid vector excluding bidder i 's bid. Agents have standard quasi-linear utilities.

Definition 1 (Player utility function). *The player utilities are their payments less the cost incurred from their allocation, $u_i(\mathbf{b}) = p_i(\mathbf{b}) - c_i x_i(\mathbf{b})$.*

Based on the information agents have about each other, we consider different equilibrium concepts. In the incomplete-information case, where bidders don't have information about each other's costs, we insist that the mechanism be dominant strategy incentive compatible (DSIC), as otherwise, agents lack information on how they should bid.

Definition 2 (Dominant-strategy incentive-compatible). *A mechanism is DSIC if for every bidder i , for all $b_i \geq 0$, and all bid vectors \mathbf{b}_{-i} , $u_i(c_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$*

In the complete-information setting, where the full cost vector \mathbf{c} is common knowledge to all the bidders, we relax our solution concept to pure-strategy Nash equilibria (PNE).

Definition 3 (PNE). *A bid vector $\tilde{\mathbf{b}}$ is a PNE, if for every bidder i , and all alternate bids $b'_i \geq 0$, $u_i(\tilde{\mathbf{b}}) \geq u_i(b'_i, \tilde{\mathbf{b}}_{-i})$*

We further include two ideal properties for mechanisms to satisfy. First, we target Sybil-proof mechanisms, where agents cannot increase their utility by submitting bids under multiple identities. While this problem can be addressed by the auctioneer knowing the bidder's identities in traditional settings, this property is especially important for permissionless mechanisms where the protocol might not have any information on who the bidders are. Formally, a mechanism is Sybil-proof if every bidding strategy that submits multiple bids is weakly dominated by a bidding strategy that only submits a single bid.

Definition 4 (Sybil-proof). *A mechanism is Sybil-proof, if for all bidding strategies $\sigma_i(\mathbf{c})$, potentially submitting multiple bids, there exists an alternative bidding strategy $\sigma'_i(\mathbf{c}) : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ such that for all other bids \mathbf{b}_{-i} , $u_i(\sigma'_i(\mathbf{c}), \mathbf{b}_{-i}) \geq u_i(\sigma_i(\mathbf{c}), \mathbf{b}_{-i})$*

Second, we desire mechanisms that are ex-post “safe,” where bidders have a non-trivial bidding strategy (i.e., the strategy doesn’t guarantee an agent zero utility regardless of others’ bids) under which they never regret participating in the auction. This is similar in spirit to individual rationality, where bidders are certain to have non-negative utility from reporting truthfully, but extended to settings where the bidders are playing strategically. For example, a standard first-price procurement auction is ex-post safe, since as long as an agent makes a bid above their cost, they are guaranteed to have a non-negative utility. Mechanisms that are not ex-post safe have higher barriers to entry, as bidders have to carefully reason about other bidders’ strategies to avoid losing money and thus may not enter at all.

Definition 5 (Ex-post safe). *A mechanism is ex-post safe, if for every agent i , there exists a bidding strategy $\sigma(c_i)$ such that for all bid vectors \mathbf{b}_{-i} , $u_i(\sigma(c_i), \mathbf{b}_{-i}) \geq 0$ and there exists a bid vector \mathbf{b}_{-i} such that $u_i(\sigma(c_i), \mathbf{b}_{-i}) > 0$.*

We judge the efficiency of a mechanism by the social cost it incurs, i.e., the total cost incurred by all agents who were allocated work. This is the objective procurement mechanism that typically seeks to minimize. While our mechanisms have goals beyond minimizing social cost, we still use it as a benchmark to quantify *how* much efficiency our mechanisms give up for the sake of a more robust supplier set.

Definition 6 (Social cost). *The social cost of an allocation is the total cost incurred by the agents in an allocation, $SC(\mathbf{c}, \mathbf{x}) = \sum_{i=1}^n c_i x_i$.*

To quantify how efficient a given mechanism is, we benchmark the social cost under equilibrium bidding versus the optimal social cost achievable. This ratio is referred to as the price of anarchy. Note that in our model, the optimal social cost will always be to allocate all the work to the agent with the lowest cost.

Definition 7 (Price of anarchy (PoA)). *Let \mathcal{X} be the set of all allocations and \mathcal{X}_{eq} be the set of allocations induced at all pure Nash equilibria (PNE). The price of anarchy is the ratio of the highest-cost equilibrium allocation to the lowest-cost allocation,*

$$PoA = \frac{\max_{\mathbf{x} \in \mathcal{X}_{eq}} SC(\mathbf{c}, \mathbf{x})}{\min_{\mathbf{x} \in \mathcal{X}} SC(\mathbf{c}, \mathbf{x})}.$$

We evaluate the three procurement mechanisms in Section 3, Section 4, & Section 5, respectively, along these axes.

3 DSIC Mechanism

We first consider the incomplete-information setting, where agent costs are private. We start by defining a class of allocation rules parameterized by α , closely resembling those of a Tullock contest [21] with negative exponents, which seek to explicitly allocate work to higher-cost agents.

Definition 8 (α -PARs). An α -proportional allocation rule takes the form,

$$x_i(\mathbf{b}) = \frac{b_i^{-\alpha}}{\sum_{j=1}^n b_j^{-\alpha}},$$

for $\alpha > 0$.

Recall in the procurement setting that bids correspond to costs, so bidding *higher* results in a *lower* allocation, thus the need for a *monotone decreasing* allocation rule in b_i . Given that the allocation rule is monotone, we can implement it with a DSIC mechanism using Myerson's lemma [16] using the standard normalization that $p_i(\infty, b_{-i}) = 0$. Letting $D_i = \sum_{j \neq i} b_j^{-\alpha}$, the integral form of the payment rule is,

$$p_i(b_i, \mathbf{b}_{-i}) = \frac{b_i}{1 + D_i b_i^\alpha} + \int_{b_i}^{\infty} \frac{1}{1 + D_i t^\alpha} dt.$$

This payment rule represents the amount that the protocol needs to *pay* the agents to elicit truthful reports \mathbf{b} . Note that under this mechanism, the protocol can be forced to pay an arbitrarily large amount. To cap the total payment, we can modify the mechanism to only allocate work to the k lowest bids with a maximum bid of \bar{b} and the payment rule modified accordingly. This also improves the social cost of the mechanism as allocation gets moved only towards the agents with the lowest costs. Thus, social cost and total payments can be traded off with decentralization according to the mechanism designer's preference. The following analysis can then be interpreted as assuming $c_i \leq \bar{b}$ for all i and taking $n = k$.

We quantify exactly how these allocation rules trade off efficiency for distribution by showing that α -PARs are the unique solutions to minimizing the objective $\sum_{i=1}^n c_i x_i^{1+1/\alpha}$ via a standard Lagrangian argument (See Appendix B.1 for details).

Lemma 1 (Optimization problem). The allocation rule $x_i = c_i^{-\alpha} / \sum_{j=1}^n c_j^{-\alpha}$ is the solution to the following constrained optimization problem,

$$\begin{aligned} \min_{\mathbf{x} \in [0,1]^n} & \sum_{i=1}^n c_i x_i^{1+1/\alpha} \\ \text{s.t.,} & \sum_{i=1}^n x_i = 1. \end{aligned}$$

Since $x_i \leq 1$ and $\sum_{i=1}^n x_i = 1$, the additional $1/\alpha$ term in the exponent pushes the x_i closer together rather than placing all the weight on the lowest-cost agent. Another interpretation is that the function penalizes concentrated allocations by scaling their weight super-linearly in this objective. Thus, we can see how α can be tuned to trade off efficiency and a decentralized provider set since as $\alpha \rightarrow \infty$, we get that the objective converges to the standard social cost benchmark that is minimized by a winner-take-all allocation.

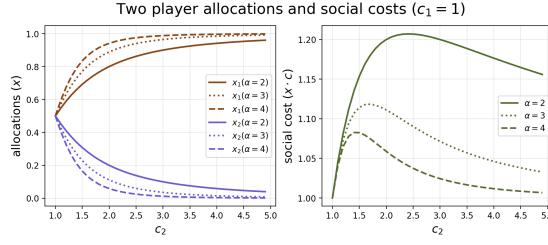


Fig. 1. Allocations (left) and social cost (right) with $n = 2$ and $c_1 = 1$ under the DSIC mechanism for various values of α .

We now bound the efficiency we trade off by optimizing for a more decentralized set of providers compared to allocating all the work to the lowest cost agent. Figure 1 shows the allocations (x_1, x_2) for the DSIC mechanism in the two-player game. We fix $c_1 = 1$ and vary c_2 . We also plot the social cost, which for the two-player game is simply $c_1x_1 + c_2x_2$. Notice that the minimal social cost in this setting would arise from fully allocating the work to player one, leading to a total incurred cost of 1. Also, notice how the value of α shapes the rate of concentration of the allocation to the low-cost player. Since the mechanism is truthful, the low-cost player will continue bidding $c_1 = 1$, while the high-cost player will bid c_2 . As c_2 increases, $c_1^{-\alpha}$ dominates the sum, and the full allocation goes to player one. Lastly, the social cost is not monotone increasing in c_2 .

For the n -player game, the social cost increases in n because a larger amount of the allocation is given to higher-cost players. The rate at which the social cost degrades is highly dependent on α . For large values of α , the allocation heavily concentrates around the lowest-cost agent regardless of how many other agents there are. Theorem 1 bounds this scaling precisely. We show that the instance with the worst-case social cost takes the form of $[1, r, r, \dots, r]$ for some $r > 1$ dependent on n and α . Intuitively, this setting implies that there is one “efficient” player and $n - 1$ “inefficient” players. From this characterization, we find the worst-case r to give a tight bound on the worst-case social cost.

Theorem 1. *The DSIC mechanism (the α -PAR allocation rule and the Myersonian payment rule) with n players has a worst-case social cost approximation of $1 + (\frac{n}{\alpha})^{\frac{1}{\alpha}}$ for $\alpha \geq 1$.*

Proof. Note that both the allocation rule and social cost-objective are homogeneous functions of degree zero in c . Thus WLOG assume for all agents i , $c_i \leq c_{i+1}$ with $c_1 = 1$. In this case, the optimal social cost is allocating the entire job to agent 1 for a cost of 1. Thus, to show the approximation, it suffices to bound the social cost of any instance with $c_1 = 1$. Since the mechanism is DSIC, we can take each $x_i = \frac{c_i^{-\alpha}}{\sum_{j=1}^n c_j^{-\alpha}}$. This implies the social cost under a cost vector of c is

$$SC(\mathbf{c}, \mathbf{x}) = \sum_{i=1}^n c_i x_i = \frac{\sum_{i=1}^n c_i^{1-\alpha}}{\sum_{i=1}^n c_i^{-\alpha}}.$$

For $\alpha = 1$, this gives us that $SC(\mathbf{c}, \mathbf{x}) = n/(1 + \sum_{j>1} c_i^{-1}) < n$. Notationally, we now refer to the social cost only in terms of the cost vector, $SC(\mathbf{c})$. Given this, we can show that for $\alpha > 1$, the cost vector maximizing the social cost must be of the form $[1, r, \dots, r]$ (See Appendix B.1 for details).

Lemma 2 (Worst-case social-cost vector). *Assuming $\alpha \neq 1$, there exists an $r \geq 1$ such that the worst-case social cost vector is of the form*

$$\mathbf{c}_{worst} = [1, \underbrace{r, r, \dots, r}_{n-1 \text{ players}}].$$

From this lemma, we have the cost vector $[1, r, r, \dots, r]$ for some $r > 1$ and the proof follows mechanically maximizing $f(r) = \frac{1+(n-1)r^{1-\alpha}}{1+(n-1)r^{-\alpha}}$ over all $r \geq 1$. See Appendix B.1 for full details.

This bound shows that the social cost is not strictly increasing in each agent's costs. There is a certain threshold beyond which an agent's increasing cost decreases their allocation at a faster rate than their increased cost harms efficiency. For each n , the bound finds the worst-case r given that number of players and α , and thus we can examine the asymptotics only as a function of n and α .

While this parameterization lets a mechanism designer precisely trade off between decentralization and efficiency as they'd like, it has two key issues: (i) it is not Sybil-proof, and (ii) it is hard for bidders to reason about and difficult to implement on-chain. To see that the DSIC mechanism is not Sybil-proof, notice that a bidder can simply make infinite bids with a constant cost of C . Each additional bid increases their allocation at the same cost.

Regarding the complexity of the payment rule, while the mechanism is DSIC, there is a large empirical literature showing that bidders don't play dominant strategies in DSIC mechanisms even when the mechanism is relatively simple, such as a second-price auction [11, 7, 6]. Furthermore, the payment rules required to implement the allocation rules as described would be quite complex and require numerical approximations, leading to complex code for the auction mechanism. One would expect bidders to fail to understand that these mechanisms are truthful and, hence, bid sub-optimally. In Section 4, we first give a mechanism that is Sybil-proof but still hard for bidders to reason about; in Section 5, we give a mechanism that is also not Sybil-proof but easier for bidders to participate in.

4 Tullock Procurement Contests

To construct a Sybil-proof mechanism that avoids winner-take-all equilibria, we take inspiration from the literature on Tullock contests. Tullock contests refer to a class of games where agents compete for a good via making costly investments [21]. The good is typically allocated randomly proportional to agents' investments, but all agents have to pay regardless of whether they win or not. Usually, in Tullock contests, the loss in efficiency from a suboptimal agent winning the

good is a downside. However, in our case, we will use this fact to implement a non-winner-take-all equilibrium. We turn our setting into a Tullock contest by having the procurer put up a budget B as a reward for bidders to compete over, which is presented as a “Proof Contest” in [19]. Agents then submit bids, and both the work and budget are allocated proportionally to agents’ bids. However, agents get their bid deducted from their reward regardless of how much reward they are allocated. Under this mechanism, we allow bidders to have negative utilities where they potentially have to pay the auctioneer.

Definition 9 (Tullock procurement contest (see also [19])). *A Tullock procurement contest is given by the following allocation and payment rule parameterized by a budget B :*

$$x_i(\mathbf{b}) = \frac{b_i}{\sum_{j=1}^n b_j}, \quad p_i(\mathbf{b}) = Bx_i(\mathbf{b}) - b_i.$$

This implies that agents have utilities given by

$$u_i(\mathbf{b}) = \frac{b_i}{\sum_{j=1}^n b_j}(B - c_i) - b_i$$

From this, we can see that Tullock contests are Sybil-proof, as a bidder can always submit a single bid to get the same payoff they would have from multiple bids.

Lemma 3. *Tullock procurement contests are Sybil-proof.*

For proof, see Appendix B.2. We now characterize the equilibrium of these mechanisms and how they achieve a near-optimal price of anarchy (PoA). We start by recalling the standard Tullock contest setting and well-known results characterizing their equilibrium and PoA.

Definition 10 (Tullock contests (from [21])). *A Tullock contest is defined by a set of n agents, each having value v_i for some good. Each agent simultaneously submits a bid b_i , and the good is randomly allocated proportional to bids, with every agent paying their bid regardless of whether they won.*

Notice that a Tullock procurement contest with a budget of B is a standard Tullock contest where each agent i has value $\max\{B - c_i, 0\}$ for winning the item. Thus, we can use the following well-known results characterizing the PNE of Tullock contests to help analyze the equilibrium for Tullock procurement contests.

Lemma 4. *(from [20]) Tullock contests where at least two agents have strictly positive value for the good have unique PNEs.*

Lemma 5. *(from [9]) At the equilibrium allocation \tilde{x} induced by the PNE of a Tullock contest, $\sum_{i=1}^n v_i \tilde{x}_i \geq \frac{3}{4} \max\{v_i\}$.*

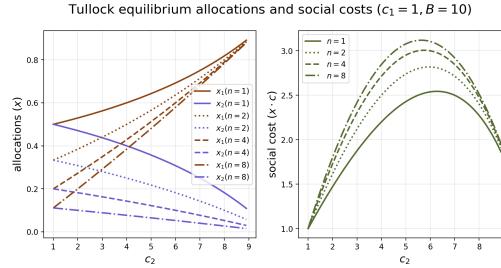


Fig. 2. Equilibrium allocations and the resulting social costs in the Tullock procurement mechanism as a function of c_2 with $c_1 = 1$ and various values of n . Note that we only show x_2 in the left subplot because $x_2 = x_3 = \dots = x_n$.

Lemma 6. (from [1]) Let $F(x) = \sum_i \max\left(1 - \frac{x}{v_i}, 0\right)$ and let v^* be the solution to $F(v^*) = 1$. Then the equilibrium allocation \tilde{x} of a Tullock contest is given by $\tilde{x}_i = \max\left(1 - \frac{v^*}{v_i}, 0\right)$

Figure 2 shows the equilibrium allocations and resulting social costs for the Tullock contest as a function of c_2 for the cost vector $\mathbf{c} = [1, c_2, \dots, c_2]$. Various values of n demonstrate the different rates at which the allocation concentrates on player 1. The right subplot shows that as n increases, the maximal social cost increases. Intuitively, this arises because larger portions of the allocation are given to the high-cost players.

From Lemma 4, we have that as long as there are two agents such that $B > c_i$, the Tullock procurement contest has a PNE. Then, applying Lemma 5, we can bound the PoA of Tullock procurement contests as a function of B and c_{\min} where c_{\min} is the lowest cost of any agent.

Lemma 7. Let $\gamma = B/c_{\min}$. Then the PoA of the Tullock procurement contest is $\frac{\gamma+3}{4}$

Proof. Let \tilde{x} be an allocation induced by the equilibrium of the Tullock contest. From Lemma 5, we have

$$\sum_{i=1}^n (B - c_i) \tilde{x}_i \geq \frac{3}{4} (B - c_{\min}) \implies \sum_{i=1}^n c_i \tilde{x}_i \leq \frac{B - 3c_{\min}}{4}$$

The lemma then follows from the optimal social cost being c_{\min} .

Hence, we see that the PoA is minimized when B is slightly above the second lowest cost out of all the agents. This corresponds to only the two bidders with the lowest costs bidding when the total reward is sufficiently small. On the contrary, as B grows large, even the highest-cost bidders participate. For sufficiently large B , all bidders receive an approximately equal allocation at equilibrium as the ratio of their values shrinks. Thus, a protocol can increase B to trade off

between efficiency (and total payment) and how many bidders they desire to be competitive at equilibrium (See Appendix B.2 for details).

Lemma 8. *In a Tullock procurement contest with equilibrium allocation \tilde{x} , for all agents i , $\lim_{B \rightarrow \infty} \tilde{x}_i = \frac{1}{n}$*

Note that while Tullock contests are Sybil-proof and implement near-optimal equilibria alongside decentralized allocations, bidders may find them problematic. In particular, Tullock contests are not ex-post safe. Regardless of a bidder's cost and the budget, for any non-zero bid b_i , there exists a bid vector \mathbf{b}_{-i} such that $(B - c_i)x_i(\mathbf{b}) = (B - c_i)\frac{b_i}{b_i + \sum_{j=1}^n b_j} < b_i$ hence giving them a negative utility. This means that bidders participating in such a mechanism must be sophisticated enough to reason about others' costs and bids to avoid losing money by participating.

In the following section, we will give a mechanism that is ex-post safe and has a simple payment rule, alleviating the difficulty of participation. However, this mechanism is no longer Sybil-proof, making it unsuitable for many decentralized applications.

5 Paid-as-Bid Mechanism

In this section, we consider a hybrid mechanism between the DSIC and Tullock contest mechanisms of the previous two sections. In particular, we use the same α -PAR allocation rule (Definition 8) but a simpler payment rule inspired by the all-pay-as-bid Tullock rule. Since α -PARs are reverse auction allocation rules, we consider the payment rule where the bidder is *paid* proportionally to their bid. These mechanisms are motivated by simplicity of implementation, where it is clear to users how their payment was derived from their bid and allocation. This is in contrast to the DSIC mechanism, where agents find it hard to reason about their payments even though truthful bidding is a dominant strategy.

Definition 11 (Paid-as-bid payment rule). *The paid-as-bid rule pays each producer i , $p_i = b_i x_i$.*

Plugging this payment rule into the player utility function (Definition 1), we can rewrite i 's utility as

$$u_i(b_i, \mathbf{b}_{-i}) := (b_i - c_i) \cdot x_i(\mathbf{b}).$$

As the mechanism is paid-as-bid (rather than the Myersonian payment rule), it will not be DSIC. Thus, we consider the PNE as the solutions of interest for various cost vectors, \mathbf{c} , in the complete-information setting. Using first-order conditions and a standard fixed-point argument, we can show that a unique PNE always exists for this mechanism regardless of the private costs (See Appendix B.3 for details).

Theorem 2 (Unique PNE). *The mechanism defined by an α -PAR with $\alpha > n/(n-1)$ and the proportional-to-bid payment rule has a unique PNE characterized implicitly by the bid vector $\tilde{\mathbf{b}}$ where*

$$\tilde{b}_i = \frac{\alpha c_i(1-x_i)}{\alpha(1-x_i)-1} \quad \forall i \in [n]. \quad (1)$$

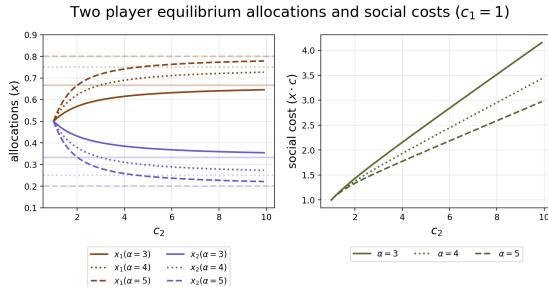


Fig. 3. Equilibrium allocations for the two-player game and corresponding social costs. The left plot shows the $1/\alpha$ and $1 - 1/\alpha$ bounds for the lower and higher allocations, respectively, as horizontal lines.

The exact equilibrium bids and allocation can be computed numerically, given this implicit characterization for different cost vectors. See Appendix B.3 for a depiction of this in the two-player game. A critical feature of the allocations at equilibrium is that they are bounded based on the size of α as $x_1 < 1 - 1/\alpha$ and $x_2 > 1/\alpha$, which we describe in the following corollary.

Corollary 1 (Bounded allocations). *Under the equilibrium bid vector, $\tilde{\mathbf{b}}$, each agent's allocation is bounded above as follows:*

$$x_i(\tilde{\mathbf{b}}) < 1 - 1/\alpha.$$

Proof. From the characterization of the equilibrium bids (1), $b_i > 0$ implies $a(1-x_i) - 1 > 0 \implies x_i < 1 - 1/\alpha$.

Figure 3 shows the equilibrium allocations and social costs for the two-player game. Again, we fix $c_1 = 1$ and increase c_2 . We show various values of α to demonstrate (i) the increase in the curvature of the α -PAR allocation rule resulting in faster concentration of the allocation on the lower-cost player, and (ii) the convergence of the allocations to $x_1 \rightarrow 1 - 1/\alpha$ and $x_2 \rightarrow 1/\alpha$ in the left subplot (shown as horizontal lines with matching line styles). The right subplot shows the scaling of the social cost in c_2 .

Note that, without additional constraints, the paid-as-bid mechanism can result in arbitrarily high payments. This is easy to see with Corollary 1; consider the two-player game with $\alpha = 3$, $c_1 = 1$, and $c_2 \rightarrow \infty$. We know that $x_1 \rightarrow$

$2/3, x_2 \rightarrow 1/3$, and by individual rationality, we must pay player two more than $c_2 \cdot x_2$. Thus, the total payment could be unbounded. As in the DSIC mechanism, this can be resolved by setting a reserve price with a maximal reportable cost of \bar{b} . However, this would change the equilibrium characterization given above. Characterizing the new equilibrium is not straightforward, as bidders with costs below \bar{b} might still have bid above \bar{b} at equilibrium. Thus, we leave characterizing the equilibrium with a reserve for future work and focus our analysis on the mechanism without a reserve.

We now turn our attention to estimating the worst-case scaling of the social cost, which is equivalent to the price of anarchy when we fix $c_1 = 1$. To bound the worst-case, we again characterize the worst-case cost vector \mathbf{c} as we did in Lemma 2. We first give a monotonicity lemma to show that for a fixed n and α , the social cost at equilibrium increases as any agent's cost increases (See Appendix B.3).

Lemma 9. *Given cost vectors \mathbf{c} and \mathbf{c}' that are identical except for some i where $c'_i > c_i$, the equilibrium social cost under \mathbf{c}' is higher than under \mathbf{c} .*

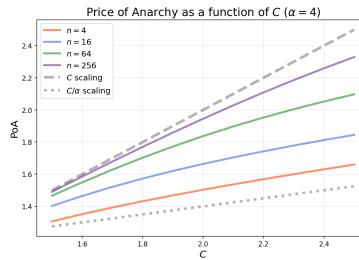


Fig. 4. Price of anarchy scaling in $C = c_{\max}/c_{\min}$ for $\alpha = 4$ and various values of n .

Combining this with the fact that both the equilibrium allocation and social cost are homogeneous in the cost vector, we can characterize the instance that gives the worst PoA parameterized by the ratio of the maximum-cost bidder to the minimum-cost bidder.

Corollary 2. *With n bidders where $c_{\max}/c_{\min} = C$, the worst-case PoA is realized with the cost vector $[1, C, C, \dots, C]$.*

From Corollary 1, we know the PoA will always be in the range $[C/\alpha, C]$ because the lower-cost players will always be allocated at least $1/\alpha$ in aggregate. We also know that as n increases, the PoA goes from C/α to C because each additional player is allocated a non-zero amount, which pulls some of the allocation from the low-cost player. We use this corollary along with the PNE characterization given in Equation (3) to numerically calculate the PoA as a function of C, α , and n . Figure 4 plots this PoA as a function of C for $\alpha = 4$ and various values of n . As n increases, the PoA shifts from the C/α scaling regime (shown as the dotted line) to C scaling (shown as the dashed line).

6 Conclusion

This work serves as a starting point for formalizing procurement mechanisms that do not optimize purely for efficiency. While we explore a range of mechanisms in this article, none is a panacea for all the desiderata presented in Section 2. The DSIC mechanism in Section 3 has serious drawbacks of not being Sybil-proof and having complex payment rules, each of which is addressed in isolation by the Tullock and the paid-as-bid mechanism of Sections 4 and 5, respectively. Tullock contests are Sybil-proof at the cost of not being ex-post safe; the paid-as-bid mechanism has a simple and interpretable payment rule structure, but loses Sybil-proofness. The DSIC and paid-as-bid mechanisms, by default, have arbitrarily high payments, but can both be modified to bound the protocol’s cost. As such, the immediate domain for future theoretical work is proving the existence or non-existence of a procurement mechanism that is Sybil-proof, and ex-post safe with non-winner-take-all equilibria.

We believe that this work is well-motivated as blockchains continue to mature and develop alternative roles for nodes participating in the network. For example, Ethereum is considering migrating to a ZK-proof-based execution environment, where each block is built along with a proof of validity, and attesters are only required to verify the proof rather than execute the transactions in their entirety. This paradigm shift could introduce the role of “provers” explicitly into consensus, and this work shows how different modeling choices could lead to different mechanisms. If the cost of generating a proof continues to drop (it is currently estimated to be about 5 cents per block) and is publicly known (e.g., through proof data aggregation layers like <https://ethproofs.xyz/>), then either the Tullock contest or the paid-as-bid mechanisms might be the best choice. When considering the possibility of a Sybil gadget (e.g., staking) that overlays the actual procurement mechanism, and the explicit goal of making proving decentralized, choosing an α -PAR might be well motivated. We hope that practitioners and builders contribute to the dialogue and help narrow down the right modeling decision for decentralized procurement.

The blockchain industry has well-established norms for open-source software and public data, which affords excellent visibility into mechanisms being implemented in production today with real money on the line – an excellent empirical testbed. As the nascent prover markets (e.g., Succinct, RISC Zero, and =Nil;) evolve, measuring the practical bidding strategies and resulting allocations will provide a key empirical grounding for this work. Measurements for the number of bidders participating, the size of available rewards, and the hardware costs will all inform the theoretical study of permissionless procurement mechanism design. This data will also help ground which of the presented mechanisms may be most practical. For example, if it turns out that the prover set is relatively small and non-anonymous, using a non-Sybil-proof mechanism (or adding a Sybil-proofness gadget like staking) may be good enough in practice. We give an overview of future modeling work in Appendix A.

References

1. Arnosti, N., Weinberg, S.M.: Bitcoin: A natural oligopoly. *Management Science* **68**(7), 4755–4771 (2022)
2. Baldick, R., Grant, R., Kahn, E.: Theory and application of linear supply function equilibrium in electricity markets. *Journal of regulatory economics* **25**, 143–167 (2004)
3. Che, Y.K.: Design competition through multidimensional auctions. *The RAND Journal of Economics* pp. 668–680 (1993)
4. Dimitri, N., Piga, G., Spagnolo, G.: *Handbook of procurement*. Cambridge University Press (2006)
5. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* **97**(1), 242–259 (2007)
6. Garratt, R.J., Walker, M., Wooders, J.: Behavior in second-price auctions by highly experienced ebay buyers and sellers. *Experimental Economics* **15**, 44–57 (2012)
7. Harstad, R.M.: Dominant strategy adoption and bidders' experience with pricing rules. *Experimental economics* **3**, 261–280 (2000)
8. Johari, R.: Efficiency loss in market mechanisms for resource allocation. Ph.D. thesis, Massachusetts Institute of Technology (2004)
9. Johari, R., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* **29**(3), 407–435 (2004)
10. Johari, R., Tsitsiklis, J.N.: Parameterized supply function bidding: Equilibrium and efficiency. *Operations research* **59**(5), 1079–1089 (2011)
11. Kagel, J.H., Levin, D.: Independent private value auctions: Bidder behaviour in first-, second-and third-price auctions with varying numbers of bidders. *The Economic Journal* **103**(419), 868–879 (1993)
12. Kang, K., Miller, R.A.: Winning by default: Why is there so little competition in government procurement? *The Review of Economic Studies* **89**(3), 1495–1556 (2022)
13. Karpowicz, M., Malinowski, K.: Characterization of scalar strategy mechanisms for efficient resource allocation: supply side case. In: 2007 46th IEEE Conference on Decision and Control. pp. 150–155. IEEE (2007)
14. Khodabakhsh, A., Nikolova, E., Pountourakis, E., Horn, J.: Threshold mechanisms for dynamic procurement with abandonment. In: Algorithmic Game Theory: 16th International Symposium, SAGT 2023, Egham, UK, September 4–7, 2023, Proceedings. p. 383–400. Springer-Verlag, Berlin, Heidelberg (2023). https://doi.org/10.1007/978-3-031-43254-5_22
15. Klemperer, P.D., Meyer, M.A.: Supply function equilibria in oligopoly under uncertainty. *Econometrica: Journal of the Econometric Society* pp. 1243–1277 (1989)
16. Myerson, R.B.: Optimal auction design. *Mathematics of operations research* **6**(1), 58–73 (1981)
17. Nakabayashi, J.: Small business set-asides in procurement auctions: An empirical analysis. *Journal of Public Economics* **100**, 28–44 (2013)
18. Nash, J.F.: The bargaining problem. *Econometrica* **18**(2), 155–162 (1950)
19. Roy, U., Guibas, J., Kulkarni, K., Pai, M., Robinson, D.: Succinct network: Prove the world's software. <https://pdf.succinct.xyz/> (2024), accessed: 30 January 2025
20. Tullock, G.: Efficient Rent Seeking Revisited, chap. 7, pp. 91–101. Kluwer Academic Publishers (1988)
21. Tullock, G.: Efficient rent seeking. In: 40 Years of Research on Rent Seeking 1, pp. 105–120. Springer (2008)

A Future Work

Model extensions While α -PARs allow the designer to trade off between distribution of the allocation and efficiency, it remains a subject of debate as to the exact metric of fairness to optimize for. Lemma 1 demonstrates that α -PARs minimize a scaled social cost, but there are other notions of fairness (e.g., Nash Social Welfare [18]) that may be of interest. More generally, micro-founding fairness in blockchain settings remains important in justifying the value of decentralization that is often cited as the goal of existing systems. Whether the decentralization offers unique properties to the system (e.g., censorship resistance) or aims to achieve better user outcomes (e.g., more competition and thus lower prices), it is often accepted as an unquestioned axiom in blockchain protocol design without proper justification.

Extending the model to the Bayesian setting is also an open problem. We justify the use of the complete-information setting and the study of PNE because the open, repeated nature of the auctions may allow providers to learn about their competitors’ private costs. While this may be true, modeling this knowledge in the incomplete-information setting as defining a prior over the competitors’ costs and studying the Bayes-Nash Equilibria would yield different outcomes that are worth considering. Further, while we use the repeated nature of the game to motivate trading off efficiency for a more distributed allocation, explicitly modeling the dynamics in a repeated game would be an important way to sanity-check the effectiveness of the mechanisms we propose.

B Supplementary Material

B.1 Section 3 Proofs

Lemma 10 (Optimization problem). *The allocation rule $x_i = c_i^{-\alpha} / \sum_{j=1}^n c_j^{-\alpha}$ is the solution to the following constrained optimization problem,*

$$\begin{aligned} \min_{\mathbf{x} \in [0,1]^n} & \sum_{i=1}^n c_i x_i^{1+\alpha} \\ \text{s.t.,} & \sum_{i=1}^n x_i = 1. \end{aligned}$$

Proof. Let

$$F(\mathbf{x}) = \sum_{i=1}^n \frac{c_i x_i^{1+\alpha}}{1 + 1/\alpha}.$$

With the constraint $\sum_{j=1}^n x_j = 1$, we have the Lagrangian,

$$\begin{aligned} \mathcal{L}(\mathbf{x}) &= F(\mathbf{x}) + \lambda \left(1 - \sum_{j=1}^n x_j \right) \\ &= \sum_{i=1}^n \frac{c_i x_i^{1+\alpha}}{1 + 1/\alpha} + \lambda \left(1 - \sum_{j=1}^n x_j \right). \end{aligned}$$

Taking partials, we have

$$\frac{\partial \mathcal{L}}{\partial x_i} = c_i x_i^{1/\alpha} - \lambda = 0 \implies x_i = \left(\frac{\lambda}{c_i} \right)^{\alpha}$$

Combining these with the full allocation constraint, we have

$$\sum_{j=1}^n \left(\frac{\lambda}{c_j} \right)^{\alpha} = 1 \implies \lambda = \left(\frac{1}{\sum_{j=1}^n c_j^{-\alpha}} \right)^{1/\alpha}.$$

Solving for the final allocation, we have,

$$x_i = \left(\frac{\lambda}{c_i} \right)^{\alpha} = \left(\left(\frac{1}{\sum_{j=1}^n c_j^{-\alpha}} \right)^{1/\alpha} / c_i \right)^{\alpha} = \frac{c_i^{-\alpha}}{\sum_{j=1}^n c_j^{-\alpha}}.$$

Lemma 11 (Worst-case social-cost vector). *Assuming $\alpha \neq 1$, there exists a $r \geq 1$ such that the worst-case social cost vector is of the form*

$$\mathbf{c}_{worst} = [1, \underbrace{r, r, \dots, r}_{n-1 \text{ players}}].$$

Proof. Let the cost maximizing vector be given by \mathbf{c}^* . Consider the function $g_k(s)$ representing the social cost as bidder k 's cost varies, holding the other bidders' costs fixed, i.e. $g_k(s) = SC([c_1, \dots, c_{k-1}, s, c_{k+1}, \dots, c_n])$. Note that for $k > 1$, we must have $g'_k(c_i^*) = 0$. Let $G_{-i} = \sum_{j \neq i} c_j^{-\alpha}$ and $H_{-i} = \sum_{j \neq i} c_j^{1-\alpha}$ and. Then, we can calculate

$$\begin{aligned} g'_i(x) &= \frac{(1-\alpha)x^{-\alpha}(x^{-\alpha} + G_{-i}) + (x^{1-\alpha} + H_{-i})(\alpha x^{-\alpha-1})}{(x^{-\alpha} + G_{-i})^2} \\ &= \frac{x^{-2\alpha} + (1-\alpha)G_{-i}x^{-\alpha} + \alpha H_{-i}x^{-\alpha-1}}{(x^{-\alpha} + G_{-i})^2} \\ &= \frac{x^{\alpha-1}(x^{1-\alpha} + (1-\alpha)G_{-i}x + \alpha H_{-i})}{(1 + G_{-i}x^\alpha)^2} \end{aligned}$$

Thus for all $i > 1$, taking $g'_i(c_i) = 0$ gives us

$$c_i^{1-\alpha} + (1-\alpha)c_i \sum_{j \neq i} c_j^{-\alpha} + \alpha \sum_{j \neq i} c_j^{1-\alpha} = 0.$$

Now let $G = \sum_{j > 1} c_j^{-\alpha}$ and $H = \sum_{j > 1} c_j^{1-\alpha}$ then we can rewrite

$$g'_i(c_i) = c_i^{1-\alpha} + (1-\alpha)c_i(G - c_i^{-\alpha} + 1) + \alpha(H - c_i^{1-\alpha} + 1).$$

Note that $c_i^{1-\alpha} - (1-\alpha)c_i \cdot c_i^{-\alpha} - \alpha c_i^{1-\alpha} = 0$ simplifying

$$g'_i(c_i^*) = (1-\alpha)(1+G)c_i^* + \alpha(1+H) = 0.$$

Since we assume $\alpha \neq 1$ and this holds for all $i > 1$, we have that for all $i, j > 1$, $c_i^* = c_j^*$.

B.2 Section 4 Proofs

Lemma 12. *Tullock procurement contests are Sybil proof.*

Proof. Consider a bidding strategy $\sigma_i(c) = (b_i^1, \dots, b_i^k)$. Then let $\sigma'_i(c) = \sum_{i=1}^k \sigma_i(c)$. Let the sum of the other agents' bids be β . Then we have that

$$u_i(\sigma_i(c), \mathbf{b}_{-i}) = \sum_{j=1}^k \left(\frac{b_i^j}{\sum_{\ell=1}^k b_i^\ell + \beta} (B - c_i) - b_i^j \right) = u_i(\sigma'_i(c), b_{-i}).$$

Lemma 13. *In a Tullock procurement contest with equilibrium allocation \tilde{x} , for all agents i , $\lim_{B \rightarrow \infty} \tilde{x}_i = \frac{1}{n}$*

Proof. We first show that for large enough B , all agents bid and hence have a positive allocation at equilibrium. Consider the function $F(B, x) = \sum_{i=1}^n \max\{1 - \frac{x}{B - c_i}, 0\}$ modified from Lemma 6 for our setting. Define $x^*(B)$ to be the solution to $F(B, x^*(B)) = 1$. It suffices to show $B - c_{max} > x^*(B)$ for sufficiently large B . Note that for all i , $\lim_{B \rightarrow \infty} \frac{B - c_{max}}{B - c_i} = 1$, thus $\lim_{B \rightarrow \infty} \sum_{i=1}^n 1 - \frac{B - c_{max}}{B - c_i} < 1$

and $B - c_{max} > x^*(B)$ since $F(B, x)$ is monotone decreasing in x . From this, for sufficiently large B , we have that all bidders are active, allowing us to explicitly solve for $x^*(B)$:

$$\sum_{i=1}^n 1 - \frac{x^*(B)}{B - c_i} = 1 \implies x^*(B) = \frac{n-1}{\sum_{i=1}^n \frac{1}{B-c_i}}.$$

It follows that $\lim_{B \rightarrow \infty} \frac{x^*(B)}{B - c_i} = \frac{n-1}{n}$. Then again using Lemma 6 again, we have that

$$\lim_{B \rightarrow \infty} \tilde{x}_i = \lim_{B \rightarrow \infty} 1 - \frac{x^*(B)}{B - c_i} = \frac{1}{n}.$$

B.3 Section 5 Proofs

Figure 5 shows the equilibrium bids and allocations, denoted \tilde{b} and \tilde{x} , respectively, for the two-player game. With $c_1 = 1$, we increase c_2 to demonstrate how the equilibrium bids and allocations evolve. For comparison, we include the honest bids and allocations (shown as solid lines in each subplot) arising from the truthful bidding in the DSIC mechanism from Section 3.

Two player honest and paid-as-bid equilibrium bids and allocations ($c_1 = 1, \alpha = 5$)

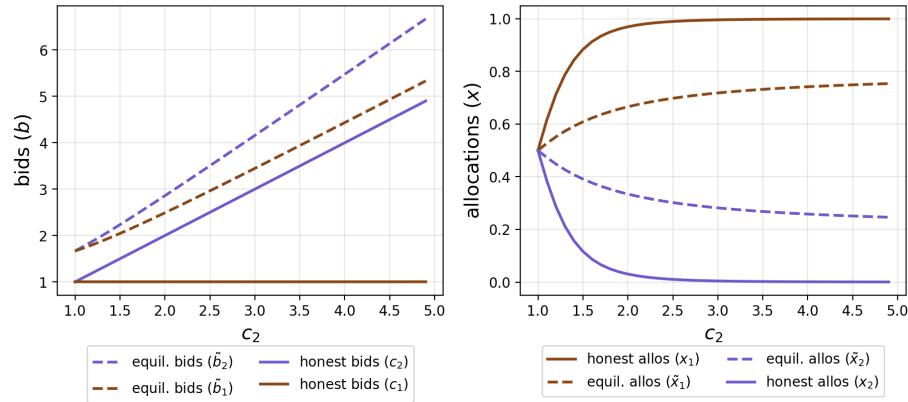


Fig. 5. Honest and equilibrium-induced bids (for the paid-as-bid payment rule, Definition 11) and allocations for the two-player game with $\alpha = 5$. We see that the equilibrium allocations approach their limits of $1/5$ and $1 - 1/5$ as c_2 increases. Also note that the equilibrium bids for player one increase in c_2 , while honest bids remain constant at $c_1 = 1$.

Theorem 3 (Unique PNE). *The mechanism defined by an α -PAR with $\alpha > n/(n-1)$ and the proportional-to-bid payment rule has a unique PNE characterized by the bid vector $\tilde{\mathbf{b}}$.*

Proof. Let $\mathbf{b} = (b_1, \dots, b_n)$ be an bid vector and let $o_i = \sum_{j \neq i} \frac{1}{b_j^\alpha}$. Then we can rewrite $u_i(b_i, \mathbf{b}_{-i}) = (b_i - c_i) \left(\frac{1}{1 + b_i^\alpha o_i} \right)$. From this, we can calculate the first order conditions for \mathbf{b} to be an equilibrium bid vector:

$$\begin{aligned} \frac{\partial u_i(b_i, \mathbf{b}_{-i})}{\partial b_i} &= \frac{1 + b_i^\alpha o_i + \alpha b_i^{\alpha-1} o_i (c_i - b_i)}{(1 + b_i^\alpha o_i)^2} = 0 \\ \iff 1 + b_i^\alpha o_i + \alpha b_i^{\alpha-1} o_i (c_i - b_i) &= 0. \end{aligned}$$

Note that this expression is strictly increasing for $0 \leq b_i < c$ and strictly decreasing for $b_i > c$. Since $u'(0, \mathbf{b}_{-i}) = 1$ this implies $\frac{\partial u_i}{\partial b_i}$ has a unique zero corresponding to a global max for $u_i(b_i, \mathbf{b}_{-i})$. This implies satisfying the first-order conditions is sufficient for b to be an equilibrium. Then, from the definition of o_i and our allocation rule x , we have

$$\frac{1}{1 + b_i^\alpha o_i} = x_i \iff b_i^\alpha o_i = \frac{1}{x_i} - 1.$$

Substituting this into our first-order condition gives

$$\frac{1}{x_i} + \frac{\alpha c_i}{b_i} \left(\frac{1}{x_i} - 1 \right) - \frac{\alpha}{x_i} + \alpha = 0 \iff b_i = \frac{\alpha c_i (1 - x_i)}{\alpha (1 - x_i) - 1} \quad (2)$$

Letting $B = \sum_{i=1}^n b_i^{-\alpha}$, and noting $x_i = \frac{b_i^{-\alpha}}{B}$ $\implies b_i = (x_i B)^{-\alpha}$ this gives

$$\frac{\alpha c_i x_i^{\frac{1}{\alpha}} (1 - x_i)}{\alpha (1 - x_i) - 1} = \frac{1}{B^{\frac{1}{\alpha}}} \quad \forall i \in [n].$$

It follows that for x to denote an equilibrium allocation

$$\frac{c_i x_i^{\frac{1}{\alpha}} (1 - x_i)}{\alpha (1 - x_i) - 1} = \frac{c_j x_j^{\frac{1}{\alpha}} (1 - x_j)}{\alpha (1 - x_j) - 1} \quad \forall i, j \in [n]. \quad (3)$$

Assume that such an allocation \tilde{x} satisfying (3) exists, then we show that setting \tilde{b} according to (1) with $x_i = \tilde{x}_i$ is consistent with Definition 8. Let $\frac{\alpha c_i \tilde{x}_i^{\frac{1}{\alpha}} (1 - \tilde{x}_i)}{\alpha (1 - \tilde{x}_i) - 1} = \kappa$. Then from manipulating (1) we get $\tilde{x}_i = (\alpha \kappa / \tilde{b}_i)^\alpha$. Combining this with $\sum_{i=1}^n \tilde{x}_i = 1$ gives $(\alpha \kappa)^\alpha = 1 / \sum_{i=1}^n \tilde{b}_i^{-\alpha}$ in turn showing $\tilde{x}_i = \tilde{b}_i^{-\alpha} / \sum_{j=1}^n \tilde{b}_j^{-\alpha}$ as desired.

Now, to show that a unique PNE exists, it is sufficient to show that there is a unique allocation \tilde{x} satisfying (3). Define the function

$$f(x) := \frac{x^{\frac{1}{\alpha}} (1 - x)}{\alpha (1 - x) - 1}. \quad (4)$$

Then an allocation vector satisfying equation 3 is equivalent to finding a vector in the unit simplex $x \in \Delta^{n-1}$ such that $c_i f(x_i) = c_j f(x_j) \forall i, j \in [n]$. Further

note that for $x \in [0, 1 - 1/\alpha]$, $f(x)$ is strictly increasing with $f(0) = 1$ and $\lim_{x \rightarrow 1-1/\alpha} f(x) = \infty$. It follows that the inverse of f , $f^{-1} : [0, \infty) \rightarrow [0, 1-1/\alpha]$ is well defined. Thus for a fixed $d \in [0, 1 - 1/\alpha]$, let $g_i(d) = f^{-1}\left(\frac{c_1}{c_i} f(d)\right)$ with $g(d) = \sum_{i=1}^n g_i(d)$. f being continuous and strictly increasing implies that g is continuous and strictly increasing. Thus, $g(0) = 0$ and $g(1/n) \geq 1$ implies there exists a unique $d^* \in (0, 1/n]$ such that $\sum_{i=1}^n g_i(d^*) = 1$. Note that $g(1/n)$ is well defined since $\alpha > n/(n-1)$. We can then take $\tilde{x}_i = g_i(d^*) \forall i \in [n]$ to get the unique equilibrium allocation.

Lemma 14. *Given cost vectors \mathbf{c} and \mathbf{c}' where \mathbf{c} and \mathbf{c}' are identical except for some i where $c'_i > c_i$, the equilibrium social cost under \mathbf{c}' is higher than under \mathbf{c} .*

Proof. Let \mathbf{x} be the equilibrium allocation corresponding to a cost vector \mathbf{c} . WLOG scale \mathbf{c} so that $c_1 = 1$, since the equilibrium allocations are invariant under linear scalings of \mathbf{c} . Define the function $H(\mathbf{c}) = \sum_{i=1}^n c_i x_i$ to be the social cost of the equilibrium allocation under \mathbf{c} . Then it suffices to show $\frac{\partial H(\mathbf{c})}{\partial c_i} > 0$ for all $i \in [n], i \neq 1$. Recall from the proof of equilibrium existence, we can define $f(x) := x^{1/\alpha}(1-x)/(a(1-x)-1)$ and $g_i(x) := f^{-1}(f(x)/c_i)$ so that letting x^* be the solution to $\sum_{i=1}^n g_i(x^*) = 1$, at the equilibrium allocation \tilde{x} , $\tilde{x}_i = g_i(x^*)$ for all agents i . Thus, we can rewrite the social cost at equilibrium as

$$H(\mathbf{c}) = \sum_{i=1}^n c_i g_i(x^*)$$

We then have

$$\frac{\partial H(\mathbf{c})}{\partial c_i} = g_i(x^*) + \sum_{j=1}^n c_j \left(\frac{\partial g_j(x^*)}{\partial c_i} + \frac{\partial g_j(x^*)}{\partial x} \frac{\partial x^*}{\partial c_i} \right).$$

We can derive the following partial derivatives for $i > 1$,

$$\begin{aligned} \frac{\partial g_j(x^*)}{\partial c_i} &= \frac{f'(x^*)}{c_j f'(x_i)} \frac{\partial x^*}{\partial c_i} + \left(-\mathbf{1}_{i=j} \frac{1}{c_j} \right) \frac{f(x^*)}{c_j f'(x_i)}, \text{ and} \\ \frac{\partial g_j(x^*)}{\partial x} &= \frac{f'(x^*)}{c_j} (f^{-1})' \left(\frac{f(x^*)}{c_j} \right) = \frac{f'(x)}{c_j f'(x_j)}. \end{aligned}$$

Plugging these into the previous expression and simplifying gives us

$$\begin{aligned} \frac{\partial H(\mathbf{c})}{\partial c_i} &= g_i(x^*) + \sum_{j=1}^n c_j \left(\frac{f'(x^*)}{c_j f'(x_j)} \frac{\partial x^*}{\partial c_i} + \frac{f'(x^*)}{c_j f'(x_j)} \frac{\partial x^*}{\partial c_i} \right) \\ &= x^* + \frac{\partial x^*}{\partial c_i} \sum_{i=1}^n \frac{f'(x^*)}{f'(x_j)}. \end{aligned}$$

Since $f'(x) > 0$ for all $x \in [0, 1 - 1/a]$, all that remains to show is $\frac{\partial x^*}{\partial c_i} > 0$. We have that x^* is defined implicitly by the equation

$$\sum_{i=1}^n g_i(x^*) = 1$$

Taking the derivative of this with respect to c_i gives us

$$\sum_{j=1}^n \frac{\partial g_j(x^*)}{\partial c_i} + \frac{\partial g_j(x^*)}{\partial x} \frac{\partial x^*}{\partial c_i} = 0 \implies \frac{\partial x^*}{\partial c_i} = -\frac{\sum_{j=1}^n \frac{\partial g_j(x^*)}{\partial c_i}}{\sum_{j=1}^n \frac{\partial g_j(x^*)}{\partial x}}$$

From our previous derivations, we can compute

$$\begin{aligned} \sum_{j=1}^n \frac{\partial g_j(x^*)}{\partial c_i} &= \frac{\partial x^*}{\partial c_i} \sum_{j=1}^n \frac{f'(x^*)}{c_j f'(x_j)} - \frac{f(x^*)}{c_i^2 f'(x_i)}, \text{ and} \\ \sum_{j=1}^n \frac{\partial g_j(x^*)}{\partial x} &= \sum_{j=1}^n \frac{f'(x^*)}{c_j f'(x_j)}. \end{aligned}$$

Finally, we have

$$\begin{aligned} \frac{\partial x^*}{\partial c_i} &= \frac{\frac{f(x^*)}{c_i^2 f'(x_i)} - \frac{\partial x^*}{\partial c_i} \sum_{j=1}^n \frac{f'(x^*)}{c_j f'(x_j)}}{\sum_{j=1}^n \frac{f'(x^*)}{c_j f'(x_j)}} \\ &= \frac{f(x^*)}{c_i^2 f'(x_i) f'(x_1)} \left(\sum_{j=1}^n \frac{1}{c_j f'(c_j)} \right)^{-1} - \frac{\partial x^*}{\partial c_i} \\ \implies \frac{\partial x^*}{\partial c_i} &= \frac{f(x^*)}{2c_i^2 f'(x_i) f'(x_1)} \left(\sum_{j=1}^n \frac{1}{c_j f'(c_j)} \right)^{-1} > 0 \end{aligned}$$