Analyzing the Economic Impact of Decentralization on Users

S. Matthew Weinberg
Princeton University, NJ, USA

Amit Levy
Better Bytes & Princeton University, NJ, USA

Chenghan Zhou
Stanford University, CA, USA

Abstract

We model the ultimate price paid by users of a decentralized ledger as resulting from a two-stage game where Miners (/Proposers/etc.) first purchase blockspace via a Tullock contest, and then price that space to users. When analyzing our distributed ledger model, we find:

- A characterization of all possible pure equilibria (although pure equilibria are not guaranteed to exist).
- A natural sufficient condition, implied by Regularity (à la [34]), for existence of a "market-clearing" pure equilibrium where Miners choose to sell all space allocated by the Distributed Ledger Protocol, and that this equilibrium is unique.
- The *market share of the largest miner* is the relevant "measure of decentralization" to determine whether a market-clearing pure equilibrium exists.
- Block rewards do not impact users' prices at equilibrium, when pure equilibria exist. But, higher block rewards can cause pure equilibria to exist.

We also discuss aspects of our model and how they relate to blockchains deployed in practice. For example, only "patient" users (who are happy for their transactions to enter the blockchain under any miner) would enjoy the conclusions highlighted by our model, whereas "impatient" users (who are interested only for their transaction to be included in the very next block) still face monopoly pricing.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design; Applied computing → Digital cash

Keywords and phrases Blockchain, Cryptocurrency, Blockspace Markets, Decentralization, Distributed Ledgers, Equilibrium Analysis, Tullock Contests

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.34

1 Introduction

Following Nakamoto's creation of Bitcoin in 2008 [35], adoption of blockchain technology for various purposes has steadily grown.[1] More relevant to this paper is ongoing interest in so-called "Web3" or "Decentralized Apps", for which an estimated $5.4B USD in VC funding was raised in 2024.[2] This paper seeks contributions to a theoretical foundation for

[1] For example, Forbes reports a cryptocurrency market cap of $3.27T USD at time of writing. Source: https://www.forbes.com/digital-assets/crypto-prices/

[2] Source: https://cointelegraph.com/news/vc-roundup-web3-funding-5-4-billion-2024 . Crunchbase further estimates a cumulative $111 B USD in VC funding raised for Web3: https://news.crunchbase.com/web3-startups-investors/.

why users might (or might not) ultimately find value in decentralized services in comparison to centralized alternatives.

Classic vs. Modern Pitches for Decentralized Services. Aside from financial speculation, perhaps the dominant 'real' use case of blockchain technology is as a currency for users with no viable alternative. While compelling applications, the economic case for such users is relatively straight-forward because the competing product[3] is so dysfunctional that concerns about (say) Bitcoin's transaction fees, volatility, and UI become very second-order. A classical pitch for decentralization therefore emphasizes simply that decentralized services make it more challenging for authoritarian leaders (and law enforcement) to deny access, and this pitch is plenty convincing in comparison to the (functionally non-existent) alternatives.

A modern discussion on blockchain technology, however, includes applications targeting users in developed economies with highly developed alternatives. For example, the pitch for stablecoins to users with a hyper-inflating local currency looks very different than to users with access to Venmo, Paypal, and credit cards. Consider also decentralized services such as file storage (for which centralized services such as Dropbox are a reasonable substitute), social networks (for which centralized services such as Facebook or Twitter are a reasonable substitute), or gaming (for which centralized gaming services produced by Riot or Blizzard are a reasonable substitute) – what would cause users to prefer decentralized services over highly-developed centralized alternatives?

A natural answer is that perhaps the decentralized service might somehow be 'better' than the centralized competitor.[4] But it is initially confusing how that might possibly arise as centralized services can optimally coordinate to lower internal costs, whereas decentralized services must additionally manage incentives/trust across distributed entities.

Decentralizing Natural Monopolies. One well-understood source of inefficiency in centralized services is deadweight loss caused by a monopolist.[5] That is, a decentralized service might plausibly be desirable to a highly-developed centralized alternative simply because the decentralized service results in different prices, and this can still be the case even if the centralized infrastructure is more efficient. Therefore it is natural to target domains with a "natural monopoly" aspect (such as social networks, payment systems, marketplaces, etc.).

Indeed, independently of any blockchain discussions, [48] highlights natural monopolies for digital services as a growing challenge, and further poses several possible approaches (each with drawbacks). One approach is described as follows: "An alternative approach to full-scale regulation consists in insulating a natural monopoly (or bottleneck or essential facility) segment, as became popular in the late twentieth century. This segment remains regulated and is constrained to provide a fair and nondiscriminatory access to competitors in segments that do not exhibit natural monopoly characteristics and therefore can sustain competition."[6] One of two key drawbacks of this approach is as follows: "... one wants

[3] For example: a currency likely to be frozen by an authoritarian government, a hyperinflating currency, or a currency that can be tracked by law enforcement critical of your illicit activity

[4] See here for an example of this pitch: https://a16zcrypto.com/posts/article/how-stablecoins-will-eat-payments/ .

[5] The holdup problem is another – see Section 1.3 for a brief discussion.

[6] [48] cites several examples: electricity markets might insulate the natural monopoly (the grid) and enable open competition on generation, or rail travel might insulate the natural monopoly (the tracks/stations) and enable open competition on train operation. A prevalent digital example is Local Loop Unbundling, where many countries insulate the natural monopoly (the "local loop" – physical copper wires servicing telecommunications) by requiring its owners to lease access at nondiscriminatory prices to service providers.

to break up the incumbent without destroying the benefits of network externalities. For example, breaking a social network into two or three social networks might not raise welfare." [7]

One interpretation of Bitcoin is exactly through this lens, and [29] are the first to make this point.[8] Indeed, when viewed as a payment system, the natural monopoly segment is 'ledger maintenance' (where a consistent record of transactions is maintained), while the user-facing 'transaction processing' segment is not a natural monopoly. In the language of Bitcoin, substantial network effects arise from having consensus on a single consistent ledger, but minimal network effects arise from users transacting within the same block (or with the same miner). In the language of payment systems, substantial network effects arise from users transacting in the same 'currency', but minimal network effects arise from users using the same app to process those transactions.[9] Perhaps shockingly, this insulation is maintained without regulation,[10] and therefore provides a novel approach to insulating natural monopolies.

But does Decentralization Actually Help? The preceding paragraph highlights blockchain-style decentralization an innovative approach to insulate natural monopolies from derivative services, but should we expect users to ultimately be better off? How would the answer depend on market primitives? Moreover, how do we even draw conclusions on users' utility from decentralized services? Surprisingly few answers are known to questions like these, and surprisingly fewer frameworks are known to even approach them. The goal of this paper is therefore to provide a framework towards such questions in the core domain of distributed ledgers, with an emphasis on connecting users' ultimate utility to properties of the decentralized ledger.

1.1 Overview of Results

We consider the core setting of a ledger. Ultimately, users desire the service of writing their transaction to the ledger, and have some value for doing so. Inspired by the preceding discussion, we separate this service into an Upstream segment which is a natural monopoly, and a Downstream segment which is not.[11]

A centralized ledger would simply provide the entire service in a vertically integrated

[7] The second key challenge highlighted is identifying a core bottleneck to insulate. We argue in Section 4 that for many domains of interest, such a bottleneck can be identified and (in theory, at least) insulated. A final challenge highlighted is the actual process of unbundling an existing product, which is unrelated to our work.

[8] "We model this novel economic structure and show that the BPS's [Bitcoin Payment System's] decentralized design offers a prototype of a payment system in which users are protected from monopoly harm even if the payment system were a monopoly... Standard economic arguments suggest that weak competition among monopolistic firms calls for regulation to mitigate monopoly harm. Under the BPS, users are protected from abuses of monopoly power even without competition from other payment systems. Thus, the BPS addresses potential antitrust concerns in a novel, even revolutionary, way."

[9] The preceding sentence is necessarily clunky, as it is somewhat unnatural to imagine separating a centralized payment system into a back-end transaction processor (that stores data and moves money around, where network effects arise) and front-end transaction processor (that interfaces directly with users, and minimal network effects arise). One high-level contribution of the "Decentralized view" is as a lens to dis-integrate services without an obvious dis-integration.

[10] Our analysis does rely on Miners treating core aspects of the consensus protocol as exogenous, which bears conceptual similarity to regulation.

[11] See Section 2 for further discussion. Intuitively, the Upstream segment directly edits the ledger, which is a natural monopoly due to network effects of multiple users sharing access to the same ledger. The Downstream segment directly interfaces with users to solicit their transactions and pass to the Upstream segment, and exhibits minimal network effects.

manner.[12] A classic Industrial Organization exercise might consider dis-integrating the Upstream monopolist from separate Downstream firms that compete with one another.[13] A distributed ledger removes centralized control entirely – the Upstream segment is provided by a hard-coded Protocol with exogenously set parameters. Competitive Downstream providers then 'purchase' the Upstream resource according to the rules of the protocol and use it to process users' transactions.[14] We analyze this in Section 3, and in particular include a discussion of why the model captures key aspects of distributed ledgers.

We then draw the following conclusions in our model:

- To the extent that a quantitative "measure of decentralization" impacts the price faced by users, it is the *size of the largest miner*.[15]
- Block rewards have limited impact on users. Specifically, block rewards cannot impact the ultimate price users would face in equilibrium (provided equilibria exist),[16] but can cause equilibria to exist.[17]
- We also characterize all possible equilibria (Theorem 10) and provide sufficient conditions for desirable equilibria to exist (Theorem 11).
- Our model applies only to *patient* users (who are happy to have their transaction included in any block), whereas *impatient* users (who want their transaction included in the next block or not at all) instead face miners with monopoly power over the contents of that block.

1.2 Roadmap

In Section 1.3, we discuss related work. Section 2 overviews our model. We include a section that provides technical preliminaries for equilibria of simultaneous first-Price auctions in our full version.[18] Section 3 describes our Distributed Ledger Model, highlights key distinctions to a Centralized provider, highlights its connection to distributed ledgers in practice, and provides our main analysis. Section 4 concludes.

1.3 Related Work

Modeling Economic Impact of Decentralized Technologies. The most closely related works in terms of motivation also seek to understand potential economic benefits of aspects of decentralized technologies (although there is no technical overlap between our work and any of these). By far the most related in terms of motivation is [29], who also view distributed ledgers through the lens of insulating a natural monopoly. [29] considers users with a simple value for service (either High or Low), and who prefer not to wait for their transactions to be included. In their model, a monopolist excludes all Low users, but immediately processes

[12] Venmo is a good example to have in mind for this model – the Venmo backend is the Upstream segment, and the Venmo app is the Downstream segment. Users enjoy network effects due to the backend database, and minimal network effects from opening the same app on their phones.

[13] The authors are not aware of a live example matching this model. A hypothetical example to have in mind would be if Venmo allowed third-party apps to access its ledger, and charged access fees to those apps.

[14] In the language of Bitcoin, miners are Downstream providers. Miners solicit transactions from users, and 'purchase' the right to include transactions in a Bitcoin block by 'paying' in hashes.

[15] See Theorems 10 and 11 for precise statements.

[16] See Theorem 10 for a precise statement.

[17] See Propositions 12 and 13 for precise statements.

[18] Due to space constraints, we prioritize presentation of our model, statements of results, and implications in the body. This technical analysis is useful primarily for technical intuition, and so is moved to the full version due to space constraints.

all High transactions, causing deadweight loss. Bitcoin, on the other hand, processes all users, but with delay cost.[19] So their work highlights a tradeoff between a monopolist (deadweight loss) and Bitcoin (delay cost). In comparison, our work (a) focuses exclusively on the monetary cost paid by users, and (b) considers a richer model of user preferences (i.e. an arbitrary demand curve, sometimes subject to a standard regularity condition).

Other works analyze the economic impact of aspects of decentralized technologies from an orthogonal viewpoint. For example, [46, 41] view decentralization/tokenization as a commitment device by which a platform can cede control to users. In addition, [25] similarly view tokenization as a commitment device to future competitive pricing. These works address a similar high-level challenge (platforms with network effects), and also through novel approaches that arose recently alongside blockchain technology. However, these works still involve a rent-seeking platform (in comparison to our exogenous protocol), and cede control to users or external investors (in comparison to changing the market structure).

Tullock Contests. At a technical level, our work studies equilibria of a two-part game, one of which is a Tullock Contest and the second of which is an auction (see Section 3 for a precise specification). As such, much of our technical analysis concerns Tullock Contests [8, 28, 27], which are commonly used to capture the game played by Bitcoin miners to produce blocks (and also to capture related aspects of blockchain ecosystems) [3, 2, 16, 4]. The key technical distinction between our work and these works lies in our second-stage auction game, which will become clear in Section 3.

Industrial Organization Theory. Our model is inspired by 'textbook' Industrial Organization Theory models [47, 44]. Our model focuses on textbook settings (without demand uncertainty, and without costly marketing) to isolate the impact of the novel blockchain-inspired market structure. This may a fruitful aspect to consider as future work develops. In classical language, we model downstream producers that sell identical products (because the users are patient, and therefore indifferent to which block they get in). Impatient users (which are not the focus of our work, as they simply face a downstream monopolist) would instead be captured by perfectly differentiated downstream products (because impatient users want only to enter the next block).

Other Economic Aspects of Blockchains. Numerous other works consider economic aspects of blockchains. Several consider the economic incentives of protocol participants [18, 17, 12, 30, 43, 7, 22, 26, 37, 38, 10, 36, 20, 52, 39, 51, 54, 1, 5, 11, 19, 9]. These works uncover reasons why participants may not be incentivized to follow the protocol specifications. In comparison to these works, we assume the underlying blockchain protocol functions as intended. Several consider "transaction fee mechanism design" – the auction specified by the protocol for users to purchase transactions from miners [31, 42, 45, 50, 53, 23, 21, 15, 14, 13, 6, 32, 33, 24]. We model miners running a first-price auction with reserve, and discuss briefly in Section 3 the connection between our modeling decision and blockchains with alternate TFMs (such as Ethereum's EIP-1559).[20] Finally, [40] considers the pricing dynamics of serial monopolists selling blockspace to patient buyers. In comparison to our work, [40] considers Miners who produce only a single block and aim to maximize their revenue from that block in isolation, whereas our work considers Miners who aim to optimize their joint revenue from multiple blocks (and also models the Tullock contest by which Miners earn the right to

[19] [29] further analyze the delay as a function of Bitcoin protocol parameters.

[20] Briefly, what really matters for our model is the cost of including a transaction on-chain (which in EIP-1559 is the base fee, and in Bitcoin is zero), and how a profit-maximizing miner would choose to sell block space (given that cost) to users who can choose to instead purchase from other miners.

produce those blocks).

2 Preliminaries

Running Story. Our model is motivated by the concept of a ledger. Ultimately, the product consumed by an end-user is the ability to write information to the global ledger (which we call a WRITE).[21] That is, each end-user has a message they would like to write on the ledger, and purchases a WRITE to do so.

The entire value proposition of a global ledger is that there is ultimately a single consistent ledger. It is therefore crucial that *some* aspects of ledger maintenance are performed via a single entity/protocol/etc. (for example, centralized ledgers should maintain a single consistent back-end database. Decentralized ledgers should have a single protocol from which observers can conclude a single consistent ledger). Intuitively, these are operations that directly edit content in the ledger (and because there is a single consistent ledger, these operations must be carefully coordinated by a single entity/protocol/etc.). Other aspects of ledger maintenance can in principle be performed by competing entities (for example, end-users can in principle face different User Interfaces, pricing schemes, etc.). We abstract away precise details of the ledger maintenance process, and simply refer to operations that directly edit the single consistent ledger as *Upstream* (and refer to one unit of these operations as an APPEND), and those that could in principle be performed by competing entities *Downstream*.

It may help to have a few examples in mind. Imagine breaking a centralized ledger (i.e. Venmo) into its back-end database maintenance and front-end User Interface. The back-end database must ensure consistency on a single global ledger, and so is Upstream. Edits to the back-end database must be reliable and consistent (even if the database is replicated, distributed, etc.). One APPEND constitutes the resources necessary to add one entry to the back-end database (maintaing consistency, availability, etc.). The front-end User Interface is Downstream – the front-end UI interacts directly with consumers, and turns communication with end-users into a query to the Upstream back-end database. The front-end UI consumes APPENDs in order to produce WRITEs, and sells WRITEs to users. Note that, in principle, the centralized ledger could offer different front-end UIs to different consumers (with different pricing schemes, different communication protocols, different app layout, etc.) – doing so does not in principle interfere with the ability to maintain a single, consistent back-end database.

One could imagine instead a centralized back-end database that provides an API for third-party app access. The centralized back-end database again is a producer of APPENDs. Each third-party app is a consumer of APPENDs and a producer of WRITEs. End-users purchase WRITEs from a third-party app (who incurs costs both from interacting with the end-user, and from purchasing APPENDs). Again, each third-party app could in principle differ in pricing schemes, communication protocols, app layouts, etc., and purchase APPENDs from the same back-end database (that interacts with each third-party app in a manner that maintains a single consistent ledger).

One could also imagine a decentralized consensus protocol maintaining a decentralized ledger, allowing participation from "miners", "stakers", "proposers", etc.[22] The decentralized

[21] In order to focus on the relevant market primitives, we do not explicitly model ledger maintenance, consensus, cryptography, privacy, or reading. The service purchased by an end-user gets their message onto the ledger, and in a manner that can be read by the desired recipients.

[22] Throughout this paper, we adopt the language of Bitcoin and refer to these participants as miners.

protocol outlines a costly procedure by which miners receive Appends.[23] Each miner is a consumer of Appends (that they "purchase" by completing the costly procedure specified in the decentralized protocol), and a producer of Writes. End-users purchase Writes from a miner (who incurs costs both from interacting with the end-user, and "purchasing" an Append). The consensus protocol structures its "sale" of Appends so that all ledger updates contribute to a single consistent ledger.[24]

The subsequent paragraphs formalize our model in the abstract – the running story provides intuition for each concept.

Market Resources. We consider two types of resources. The Downstream resource, Write, is consumed by end-users. The Upstream resource, Append, is required to produce Writes (in one-to-one ratio). Production of the Upstream resource is a natural monopoly, and therefore will be produced by a single entity/protocol. Production of the Downstream resource is not a natural monopoly, therefore we model a Downstream market with multiple competing participants.

Market Participants. There are two types of market participants. End-users are the ultimate consumers, who desire Writes. Each end-user wants a single Write, and has some value v should they receive one. Downstream producers produce Writes, which necessitates consumption of Appends. The protocol (which is hard-coded and has no objective function or strategic decisions) produces Appends.

Market Primitives. There is a continuum of end-users, with $D(p)$ denoting the mass of consumers with value at least p for a Write.

We assume that $D(\cdot)$ provides finite revenue to a monopolist (that is, $\sup_p\{p \cdot D(p)\} < \infty$). Our main results require a standard regularity assumption on $D(\cdot)$.

▶ **Definition 1** (Regular). *A demand curve $D(\cdot)$ is* Regular *if:*

- *$D(\cdot)$ is differentiable and strictly decreasing. In this case, we use $-d(\cdot) := D'(\cdot)$.*
- *The function $\varphi_D(x) := x - \frac{D(x)}{d(x)}$ is monotone non-decreasing in x.*

Structure of the Game. We model the interactions between the Upstream protocol, Downstream providers, and End-Users as a three-stage game. First, the Upstream protocol sets the dynamics for selling Appends to Downstream providers. Next, with this protocol fixed, Downstream providers set their strategies both for purchasing Appends and for selling Writes to End-Users. Finally, with these strategies fixed, End-Users set their strategies for purchasing Writes from Downstream providers.

Equilibrium Analysis. Let DP denote the set of Downstream providers, EU denote the set of End-Users, and $P_i(\vec{a})$ denote the payoff to Player i when the action profile is $\vec{a}$.

An End-User Equilibrium fixes some actions $\vec{a}_{\text{DP}}$ by the Downstream providers, and is a Nash Equilibrium of the End-User game induced by $\vec{a}_{\text{DP}}$ (with payoff $P_i(\vec{a}_{\text{DP}}; \vec{a}_{\text{EU}})$ to Player $i \in \text{EU}$ on action profile $\vec{a}_{\text{EU}}$).

A Downstream Equilibrium[25] specifies, for each possible action profile $\vec{a}_{\text{DP}}$ of the Downstream providers, an End-User Equilibrium $E(\vec{a}_{\text{DP}})$ for the end-user game induced by $\vec{a}_{\text{DP}}$, and then (together with $E(\cdot)$) is a Nash Equilibrium among Downstream providers for the Downstream game induced by $E(\cdot)$ (which awards payoff $P_i(\vec{a}_{\text{DP}}, E(\vec{a}_{\text{DP}}))$ to Player $i \in \text{DP}$

[23] For example, Bitcoin miners receive Appends by repeated hashing (which costs electricity and hardware). Ethereum stakers receive Appends by locking up ETH in the Ethereum protocol (which costs capital).

[24] That is, this paper assumes that the consensus protocol functions as intended. See Section 1.3 for a brief discussion on on related work surrounding this assumption.

[25] We will sometimes simply call this an Equilibrium.

on action profile $\vec{a}_{\text{DP}}$). When $E(\cdot)$ is unique (or otherwise clear from context), we will abuse notation and simply refer to $\vec{\alpha}_{\text{DP}}$ as a Downstream Equilibrium. Moreover, we will also abuse notation and say that a Downstream Strategy α_i dominates α_i' if α_i dominates α_i' in the game among Downstream providers induced by $E(\cdot)$.

Notation. For a (not necessarily continuous) monotone non-increasing function $F(\cdot)$, we let $F^{-1}(y) := \{x \mid \lim_{z\to x^+} F(z) \le y \le \lim_{z\to x^-} F(z)\}$, $F^{-1}_{\text{inf}}(y)$ denote the infimum of $F^{-1}(y)$, and $F^{-1}_{\text{sup}}(y)$ denote the supremum of $F^{-1}(y)$. Observe that if $F(\cdot)$ is left-continuous, then $F(x) \ge y$ for any $x \in F^{-1}(y)$.[26] If $F(\cdot)$ is continuous and strictly decreasing, we simplify notation and define $F^{-1}(y) := F^{-1}_{\text{inf}}(y) = F^{-1}_{\text{sup}}(y)$.

First-Price Auction with Reserve. First-Price Auctions with Reserves are a common subgame in our market structures. With a continuum of bidders and a total supply of Q, a first-price auction with reserve r concludes as follows. First, let $B(q)$ denote the mass of bidders who submit a bid at least as large as q.[27] Next, if $B(r) < Q$, then every bidder who submits a bid at least as large as r wins and pays their bid. If $B(r) \ge Q$, then every bidder who submits a bid strictly exceeding $B^{-1}_{\text{sup}}(Q)$ wins and pays their bid, every bidder who submits a bid strictly below $B^{-1}_{\text{sup}}(Q)$ loses, a mass of $B(B^{-1}_{\text{sup}}(Q)) - Q$ bidders who submit a bid of exactly $B^{-1}_{\text{sup}}(Q)$ lose and the remainder win and pay their bid (and in this case a total mass of Q bidders win).[28] [29] Observe that every bid profile induces an effective price of $b := \max\{B^{-1}_{\text{sup}}(Q), r\}$ – every bidder who submits a bid exceeding b certainly wins (and pays their bid) and every bidder who submits a bid below b certainly loses.

Equilibria of Simultaneous First-Price Auctions. Simultaneous First-Price Auctions are another common subgame in our market structures. Below we overview Simultaneous First-Price Auctions and technical lemmas helpful to understand our results – full analyses and proofs are in the full version.

▶ **Definition 2** (Simultaneous First-Price Auctions). *In Simultaneous First-Price Auctions, there are n sellers. Each seller i has a Q_i mass of items for sale, and sets reserve r_i. We define $Q^{\le}(r) := \sum_{i,\ r_i \le r} Q_i$ to be the total mass of items for sale at reserve at most r,*[30] *$Q^{<}(r) := \sum_{i,\ r_i < r} Q_i$ to be the total mass of items for sale at reserve strictly less than r,*[31] *and $Q^{=}(r) := Q^{\le}(r) - Q^{<}(r)$ to be the total mass of items for sale at reserve exactly r.*

A continuum of unit-demand buyers each submit a (possibly 0) bid to each first-price auction. Each first-price auction executes exactly as defined in Section 2. An equilibrium of Simultaneous First-Price Auctions is simply a strategy profile where each bidder best responds.

For equilibria among bidders, fixing all Q_i, r_i, we establishes in the full version that all winning bidders pay the same price in equilibrium, and define the value of this as *clearing price*.

[26] Because for any $x \in F^{-1}(y)$, $y \le \lim_{z\to x^-} F(z) = F(x)$.

[27] Observe that $B(\cdot)$ is left-continuous. To see this, observe that all bidders who bid at least $b-\varepsilon$ contribute to $B(b-\varepsilon)$. So the bidders that contribute to $B(b-\varepsilon)$ for all $\varepsilon > 0$ are exactly those who bid at least b – the same bidders that contribute to $B(b)$.

[28] Recall that $B(B^{-1}_{\text{sup}}(Q)) - Q \ge 0$ as B is left-continuous.

[29] All of our analysis holds no matter how ties are broken to select the winning bidders among those who bid $B(B^{-1}_{\text{sup}}(Q))$.

[30] Observe that $Q^{\le r}$ is monotone non-decreasing, and right-continuous everywhere. To see this, observe that seller i contributes Q_i to $Q^{\le}(r)$ if and only if $r_i \le r$, and to $Q^{\le}(r+\varepsilon)$ for all $\varepsilon > 0$ if and only if $r_i \le r$. Therefore all sellers contribute the same to both $Q^{\le}(r)$ and $\lim_{\varepsilon\to 0} Q^{\le}(r+\varepsilon)$.

[31] Observe that $Q^{<}(r)$ is monotone non-decreasing, and left-continuous everywhere. To see this, observe that seller i contributes to $Q^{<}(r)$ if and only if $r_i < r$, and to $Q^{<}(r-\varepsilon)$ for some $\varepsilon > 0$ if and only if $r_i < r$. Therefore, all sellers contribute the same to both $Q^{<}(r)$ and $\lim_{\varepsilon\to 0} Q^{<}(r-\varepsilon)$.

▶ **Definition 3** (Clearing Price and Canonical Equilibrium). *For an equilibrium E of Simultaneous First-Price Auctions, we refer to its* clearing price $c(E)$ *as the single bid such that every winning bidder wins exactly one item at bid $c(E)$, or loses. We say an equilibrium is* canonical *if (i) a total supply of $\min\{D(c(E)), Q^{\leq}(c(E))\}$ items are sold,*[32] *and (ii) the clearing price is minimal across all equilibria.*[33] *Note that if $D(\cdot)$ is continuous and strictly decreasing, all equilibria are canonical.*

Downstream Equilibria in our market structures concern the behavior of sellers in Simultaneous First-Price Auctions (i.e. choosing a quantity Q_i according by participating in the Upstream protocol, and setting a reserve r_i). We refer to the reserve-setting aspect as a *price-setting equilibrium* (noting that a Downstream Equilibrium must both induce a price-setting equilibrium for fixed $\vec{Q}$, *and* be a joint equilibrium when considering both investment and reserves).

▶ **Definition 4** (Price-Setting Game). *A* Price-Setting Game *has the following structure:*

Players. *There are $n > 0$ sellers. Seller i has quantity Q_i of items.*

Action Space. *Each seller picks a reserve r_i to set in a first-price auction.*

Costs. *Each seller pays a cost of c_i per item sold.*

Payoffs. *On strategy profile $\vec{r}$, a continuum of buyers with values according to $D(\cdot)$ bids in equilibrium of the simultaneous first-price auctions with quantities $\vec{Q}$ and reserves $\vec{r}$, which induces a clearing price of $p(\vec{Q}, \vec{r})$. If Seller i sells a mass of Q'_i items in this equilibrium, their payoff is $Q'_i \cdot (p(\vec{Q}, \vec{r}) - c_i)$.*

Note that if $D(\cdot)$ is continuous and strictly decreasing, the clearing price $p(\vec{Q}, \vec{r})$ is unique. We refer to a Price-Setting Game as canonical if the equilibrium selected by buyers is canonical.

A concept throughout our analyses is whether a seller clears their entire inventory, and whether they determine the price at which the bidding equilibrium clears.

▶ **Definition 5.** *We say that Seller i is* saturated *in a strategy profile $\vec{r}$ if either (i) Seller i sells a mass of Q_i items or (ii) the clearing price $p(\vec{Q}, \vec{r}) \leq c_i$. We further say that an equilibrium is saturated if all sellers are saturated. Finally, we refer to Seller i as a* price-setter *in the strategy profile $\vec{r}$ if $Q_i > 0$ and the clearing price $p(\vec{Q}, \vec{r}) = r_i > c_i$.*

Finally, Proposition 6 characterizes that all potential price-setting equilibria are either a "market-clearing equilibrium" where no seller is sufficiently large to profit from price-setting, or have a unique price-setter.

▶ **Proposition 6.** *Let $Q := \sum_i Q_i$. Then every price-setting equilibrium takes one of two forms:*

- *Every seller is saturated and the clearing price is $D^{-1}_{\sup}(Q)$. An equilibrium of this form exists if and only if $Q_i \leq \frac{(x-c_i)(Q-D(x))}{x-D^{-1}_{\sup}(Q)}$ for all i and all $x > D^{-1}_{\sup}(Q)$.*

[32] That is, there does not exist a non-zero mass of buyers with value $D(c(E))$ and an unsaturated auction with reserve $c(E)$.

[33] Intuitively, what happens is the following. There is a demand curve $D(\cdot)$ defined by the users' demand. The sellers $\{(Q_i, r_i)\}$ jointly define a supply curve, with $S(q)$ denoting the quantity of WRITEs sold in some auction with reserve at most q. The supply curve has jump discontinuities, and so the demand may "meet" supply in a discontinuity. Still, any point where supply meets demand can be the effective price, and if $D(\cdot)$ is continuous and strictly decreasing there is a unique such point.

- *There is a single price-setter i, who sets price $r_i^* := \arg\max_{x \geq D_{\sup}^{-1}(Q)}\{(x - c_i)(D(x) + Q_i - Q)\}$. If an equilibrium of this form exists, it certainly exists with $i^* := \arg\max_i\{r_i^*\}$ as the price-setter (but equilibria with other price-setters are possible). Moreover, if $c_i = c_j$ for all i, j, then $\arg\max_i\{Q_i\} = i^*$.*

3 Distributed Ledger Model

We now formally present our Distributed Ledger model. After formally specifying the model, we overview key differences to classical market structures, and its connection to distributed ledgers in practice.

▶ **Definition 7** (Distributed Ledger Model). *The* Distributed Ledger Model *has the following properties. We refer to the Upstream provider as Protocol and to the Downstream providers as* Miners.

Upstream

Protocol. *The Upstream protocol produces a fixed amount of Q_A Appends and maintains consensus. The protocol runs a* Tullock Contest *[49] in some Resource to distribute the supply of Appends, with a block reward $B \geq 0$. Specifically, Miner i who invests q_i of Resource receives a fraction of the total Append supply proportional to their investment (i.e., $Q_A \cdot \frac{q_i}{\sum_j q_j}$). Miner i also receives $B \cdot Q_A \cdot \frac{q_i}{\sum_j q_j}$ payment.*

Payoffs. *Upstream protocol has no payoffs – it simply maintains consensus on the ledger.*[34] [35]

Downstream

Players. *There is a set of n Miners.*

Action Space. *Each Miner i chooses a quantity of investment q_i in the Upstream Tullock Contest, and a reserve price r_i for a first-price auction they will run among end-users.*

Costs. *Miner i pays cost c_i^R per unit of Resource, and c^W per Write. That is, if Miner i wishes to invest q_i in the Upstream game, they pay cost $q_i \cdot c_i^R$. If Miner i eventually sells $Q_i' \leq \frac{q_i}{\sum_j q_j} \cdot Q_A$ Writes, Miner i pays cost $Q_i' \cdot c^W$. W.l.o.g. we let $c_i^R \leq c_{i+1}^R$ for all $i \in [n-1]$.*

Payoffs. *For a Miner i who invests q_i in the Upstream game and eventually sells Q_i' Writes, their total cost is $c_i^R \cdot q_i + Q_i' \cdot c^W$. They receive a block reward of $B \cdot Q_A \cdot q_i / \sum_j q_j$, plus any additional revenue earned in their first-price auction. Therefore, if Miner i earns revenue R_i from their first-price auction, invests q_i in the upstream game, and sells Q_i' Writes, their payoff is $R_i - Q_i' \cdot c^W - q_i \cdot c_i^R + B \cdot Q_A \cdot \frac{q_i}{\sum_j q_j}$.*

End-User

[34] In the case of Bitcoin (and Ethereum, and most other Decentralized Ledgers), Miners are also participants in a consensus protocol. It may be helpful to think of Upstream providers as nodes that pass messages, verify authenticity, etc. in roles that would not also result in the ability to sometimes dictate contents of a block.

[35] We have intentionally modeled the decisions of the consensus protocol as exogenous to the game we study. Of course, *someone* decides on Q_A, B, and to use a Tullock Contest in the first place. In practice, these decisions happen on a *much* slower time scale than the game we model. For example, both Bitcoin and Ethereum (and all permissionless blockchains the authors are aware of) have used Tullock Contests since their creation. Bitcoin has never changed the formula for its block reward, and Bitcoin has technically not changed its blocksize either (although 'soft forks' have occasionally increased Bitcoin's functional blocksize). Still, it is also worthwhile for future work to study the processes by which protocol parameters are set.

Players. *There is a continuum of End-Users who follow a demand curve $D(\cdot)$ for* WRITEs.

Action Space. *Each end-user submits a bid to each First-Price Auction.*

Payoffs. *An end-user with value v has payoff $v - q$ for receiving at least one* WRITE *and paying total price q,*[36] *and payoff 0 if they do not get a* WRITE.

Before proceeding to analysis, some discussion is warranted on why the above model captures popular distributed ledgers, and what differentiates it from classic market structures.

Key Differences. One key difference between the Distributed Ledger Model and traditional market structures is the presence of a non-strategic protocol. Specifically, the Upstream game is hard-coded in the Distributed Ledger Model rather than endogenously optimized by a profit-maximizing Monopolist. This distinction is key.[37] Additionally, the decision to fix the quantity Q_A of APPENDs and run a Tullock contest is material, and meaningfully affects the analysis.[38] Finally, the decision to have a block reward (which directly rewards miners for purchasing APPENDs, even if they do not ultimately sell WRITEs) is material, although our analysis shows limited impact on end-users (see Section 3.3.3).

Connecting the Distributed Ledger Model to Distributed Ledgers in Practice. First, we map aspects of Bitcoin onto the Distributed Ledger Model. Assuming that the Bitcoin protocol functions as intended,[39] let us first describe the interaction between Miners and Protocol. In order to produce a valid block, Miners must solve a "proof-of-work cryptopuzzle." Specifically, Miners trade one hash computation for one independent Bernoulli trial to create a valid block.[40] Moreover, the success rate of each Bernoulli trial is dynamically adjusted by the Bitcoin protocol so that one block is created amongst the entire network every ten minutes. Each block provides 1 MB of space for the Miner to include transactions, and awards the miner a block reward (currently 3.125 BTC). So in our model, the interaction between Miners and Protocol captures the following aspect: Resource is hash computations. Each Miner i has some cost c_i^R to perform one hash computation.[41] APPENDs are units of space in a valid block, and Protocol has hard-coded that $Q_A = 1$ MB per ten minutes are awarded in total and that each Miner receives a fraction of 1 MB blocks proportional to their hash computations (because the success probability of each hash dynamically adjusts to enforce a total quantity of 1 MB per ten minutes). Finally, the protocol hardcodes $B = 3.125$ BTC per 10 minutes as the total block reward,[42] which is also distributed proportionally to miners according to their hash computations.

Aside from its interaction with Miners, Protocol simply maintains consensus on the contents of the ledger. For example, Protocol verifies validity of contents of the ledger, resolves any conflicts using "Nakamoto consensus" [35], and widely disseminates the ledger

[36] That is, End-Users are unit-demand and only want a single WRITE – if they win multiple auctions they do not get additional utility. Still, they make a payment in any auction they win.

[37] As noted above, it is *certainly* relevant to *also* study the meta-game by which protocol rules are formed, but the game induced by a fixed protocol would still be relevant for the entirety of Bitcoin's existence.

[38] For example, conclusions would change if instead the protocol set a price p_A for APPENDs and sold whatever is demanded, even if p_A were determined exogenously. As previously noted, all blockchain protocols the authors are aware of run an Upstream Tullock contest, and it is not clear how to implement alternate Upstream market structures with a secure protocol.

[39] See Section 1.3 for a small subset of works describing manners by which the protocol may not function as intended.

[40] With extremely low probability of success – roughly 2^{-78} at the time of writing.

[41] This includes electricity, operational costs, amortized hardware costs, etc.

[42] Note that this quantity halves every four years, as pre-specified by the Bitcoin protocol.

itself. In particular, the protocol rules suffice to identify a unique consistent ledger to disseminate.

End-users get their transactions in a Bitcoin block by broadcasting to Miners. Each transaction includes a transaction fee, which is paid to whichever Miner includes that transaction in their block. Processing a transaction induces costs such as checking validity, and maintaining network connectivity (to hear about transactions in the first place), which are captured by c^W in our model. Miners are typically revenue-maximizing, and typically fill their blocks with transactions paying the highest fees (and, to the best of the authors' knowledge, typically without reserves). For a *patient* end-user, who wants their transaction in the Bitcoin ledger eventually but not necessarily immediately, each Miner is a potential seller running their own First-Price Auction, and the service offered by distinct miners is indistinguishable.

It is also worth highlighting which of these aspects are key to fit our model, and which are not. All permissionless distributed ledgers the authors are aware of run a Tullock Contest in *some* Resource. Some (including Bitcoin Cash, Litecoin, Ethereum Classic) also use hash computations as Resource ("proof-of-work"). Others (including Ethereum, Solana, Cardano) use locked capital as Resource ("proof-of-stake").[43] It is not material to our analysis which Resource is used, only that the Protocol ultimately awards block space proportional to that resource.

We note that [3] also model blockchain investment games as a Tullock contest, while other works [29] instead model it as perfect competition with free entry. We briefly note that free entry can also be captured arbitrarily well in our model by taking $n \to \infty$ Miners with identical c_i^R (see Theorem 11, for example).[44]

On the other hand, it is crucial for our model to accurately capture end-users that they are *patient* and therefore equally happy to be included in any valid block. *Impatient* users (such as those primarily motivated by DeFi applications) are not captured by our model. Instead, impatient users view the particular block offered by the next Miner as the only resource of interest, and therefore that Miner faces no competition. Therefore, a Miner selling blockspace primarily to impatient users is instead a monopolist.

Additionally, observe that Miners(/Stakers) are free to use whatever "off-chain" auction they like in order to sell space in their created blocks, independent of whatever "on-chain" mechanism is hardcoded. For example, it is immaterial to our model that Bitcoin's on-chain mechanism is pay-your-bid, whereas Ethereum's is a posted-price mechanism, because Miners(/Stakers) in both protocols can run a first-price auction with reserve off-chain to determine which transactions are included in the first place. On the other hand, the fact that Ethereum's EIP-1559 *burns*[45] revenue from the posted-price mechanism is material, and can be captured in our model via c^W. That is, our model adopts the perspective of [24] that EIP-1559 is really specifying a WRITE cost per included transaction (the burned base fee) on each *Proposer*, and the Proposer is then free to run whatever off-chain auction they like to build their block (rather than that EIP-1559 specifies *the* auction that Proposers *must* run when facing users).

Finally, while all mainstream protocols the authors are aware of currently have a single Proposer at each time slot, some protocols are now experimenting with "Multiple Concurrent

[43] That is, Miners are now called Stakers, who trade one unit of locked capital per Bernoulli trial.

[44] Specifically, the final bullet concludes that the natural "market clearing equilibrium" becomes an equilibrium with sufficiently-many identical Miners.

[45] That is, transactions included in an Ethereum pay a posted-price (set by the Ethereum protocol) that is destroyed, and not awarded to the miner.

Proposer (MCP)" protocols. In these protocols, there is no longer a monopolist for each block slot. Instead, multiple Proposers have the opportunity to insert transactions. Interestingly, our model also captures MCP protocols with either patient or impatient users.[46]

In summary, the Distributed Ledger Model captures key aspects of many mainstream blockchain protocols (they are Tullock contests, and the protocol can impact c^W), while not capturing others (such as impatient end-users, or blockchains with a heavy MEV ecosystem[47]). Indeed, this is in line with the motivation for our paper: Decentralization does not uniformly impact all users identically in all domains, and so our model necessarily picks one canonical domain of focus.

Throughout this section, we abuse notation and use $(\vec{q}, \vec{r})$ to refer to the Downstream Equilibrium $(\vec{q}, \vec{r})$ together with the mapping $E(\cdot)$ that takes $(\vec{q}', \vec{r}')$ to the canonical End-User Equilibrium. We repeat the key takeaways of this section below:

- While Equilibria may not always exist, Proposition 8 establishes an upper bound on the price any Miner will set: any strategy that sets a price exceeding the monopoly reserve for $D(\cdot)$ is a dominated strategy.
- Theorem 10 characterizes all potential Equilibria. In particular, there exists a single $\vec{Q}$ such that Miner i wins Q_i APPENDs in all pure equilibria. From here, there is a unique "market-clearing" potential equilibrium (where each miner sets reserve at most $D^{-1}(Q_A)$), and for each i a unique potential equilibrium where Miner i is a price-setter.
- Theorem 11 provides necessary and sufficient conditions for the "market-clearing" potential equilibrium to be an equilibrium.[48] Importantly, the condition *depends only on* $D(\cdot), Q_A,$ *and* $\max_i\{Q_i\}$. That is, to the extent that a "measure of decentralization" impacts the ultimate price paid by end-users, the correct "measure of decentralization" is the market share of the largest miner.
- The magnitude of the block reward (or whether there is a block reward at all) *does not impact* $\vec{Q}$*, nor any of the potential Equilibria identified by Theorem 10.* But, a larger block reward makes Equilibria more likely to exist (see Proposition 12).

3.1 Higher-than-Monopolist Reserves are Dominated Strategies

Before reasoning about equilibria, we first reason about what reserves a Miner might set in an undominated strategy. Formally, we say that Downstream Strategy (q_i, r_i) dominates (q_i', r_i') if *in the game among Downstream Providers,* $P_i((q_i, \vec{q}_{-i}), (r_i, \vec{r}_{-i})) \geq P_i((q_i', \vec{q}_{-i}), (r_i', \vec{r}_{-i}))$ for all $\vec{q}_{-i}, \vec{r}_{-i}$.

▶ **Proposition 8.** *Let* $r^*(D, Q_A) := \arg\max_{r \geq D^{-1}_{\inf}(Q_A)}\{(r - c^W) \cdot D(r)\}$*. Then for all Miners* i*, all* $q_i > 0$*, and all* $r_i > r^*(D, Q_A)$*,* $(q_i, r^*(D, Q_A))$ *dominates* (q_i, r_i)*.*

[46] To be extra clear, our model verbatim captures an MCP protocol where the block space in each block is partitioned according to stake (i.e. a 10% staker gets 10% of every block). Most MCP proposals instead sample a discrete number of Proposers proportional to stake, and allow each such Proposer an equal fraction of the block. Our model does not capture this verbatim, as the sampling process would meaningfully complicate analysis, but it would be a natural direction for follow-up work to modify our model to capture these MCP protocols verbatim.

[47] In blockchains with a heavy MEV ecosystem, the process of turning APPENDs into WRITEs itself is cost-intensive and meaningfully asymmetric.

[48] In addition, we extend the necessary and sufficient conditions for the "market-clearing" equilibrium (when they exists) to the scenario where different Miners having different cost per WRITE in the full version.

Proposition 8 establishes an upper bound on what price could possibly arise, even out of equilibrium – it can be no worse than the price that would be set by a single Miner who produces all blocks.

3.2 Characterizing Equilibria (when they exist)

End-User Equilibria are analyzed in *Equilibria of Simultaneous First-Price Auctions* of Section 2 , which describes how to determine the clearing price as a function of the Miners' strategies. This section focuses on analyzing Downstream Equilibria. While Pure Equilibria do not always exist (even when $D(\cdot)$ is Regular – see Section 3.3.1), we are able to cleanly characterize all *potential* Pure Equilibria. Below, we outline the characterization.

- Taking investments $\vec{q}$ as fixed, Miners have quantities $\vec{Q}$, which necessarily satisfy $\sum_i Q_i = Q_A$ (as Q_A is exogenously set). In any Pure Equilibrium, it therefore must be that $\vec{r}$ is a price-setting equilibrium for $\vec{Q}$.
- Taking reserves $\vec{r}$ as fixed and the total quantity Q_A as fixed, $\vec{q}$ must be an investment equilibrium. Knowing from previous analysis that any price-setting equilibrium has at most one price-setter, we can instead study: taking the clearing price r, the price-setter i, and the total quantity Q_A as fixed, $\vec{q}$ must be an investment equilibrium. This is *almost* like asking for an equilibrium in the Tullock contest defined by costs c_i^R and total reward $(r - c^W) \cdot D(r) + B \cdot Q_A$.
  - But, the game is not *exactly* a Tullock contest. Indeed, no one gets reward $((r - c^W) \cdot D(r) + B \cdot Q_A) \cdot q_i / \sum_j q_j$ – all $j \neq i$ receive reward $(r - c^W + B) \cdot Q_A \cdot q_i / \sum_j q_j$, and the price-setter i receives reward $(r - c^W + B) \cdot Q_A \cdot q_i / \sum_j q_j - (r - c^W) \cdot (Q_A - D(r))$. However, after inspecting both reward formulas, *the marginal change in each Miners' payoff is identical to the marginal change in* $(r - c^W + B) \cdot Q_A \cdot q_i / \sum_j q_j$. Therefore, the same local optimality conditions that must be satisfied by an equilibrium of a Tullock contest with total reward $(r - c^W + B) \cdot Q_A$ must be satisfied by any equilibrium with clearing price r.
  - Importantly, however, while the equilibrium *investments* in a Tullock contest certainly depends on the total reward split, the equilibrium *resulting market shares* do not.
- The previous two bullets suggest the following as necessary conditions for a pure equilibrium:
  - The resulting market share of Appends won must match those in equilibrium of a Tullock contest where Miner i incurs cost c_i^R per Resource. Importantly, this equilibrium is unique and well-defined. Call this vector of quantities $\vec{Q}^*(\vec{c})$.
  - Let $Q_i^{\text{OPT}} := \arg\max_{Q \in [\sum_{j \neq i} Q_j^*, Q_A]} \{(Q - \sum_{j \neq i} Q_j) \cdot D^{-1}(Q)\}$. Then some Miner i is a price-setter at $D^{-1}(Q_i^{\text{OPT}})$ (this includes the possibility that $Q_i^{\text{OPT}} = Q_A$ for all i, and the unique possible Pure Equilibrium saturates all Miners).
  - Indeed, Theorem 10 confirms these characterize all potential Pure Equilibria.
- Ultimately, in order to be a Pure Equilibrium, the question is whether whether the pair (q_i, r_i) *which can be changed in tandem* is a best-response to $\vec{q}_{-i}, \vec{r}_{-i}$. The previous bullets expound upon necessary conditions for this to plausibly occur – if (q_i, r_i) is to be a best response to $(\vec{q}_{-i}, \vec{r}_{-i})$, q_i must be a best response to $(\vec{q}_{-i}, \vec{r}_{-i}; r_i)$, and r_i must be a best response to $(\vec{q}_{-i}, \vec{r}_{-i}, q_i)$. Section 3.3.1 contains an example demonstrating the possibility of no Pure Equilibria.

We execute this outline in our full version – Theorem 10 characterizes all possible Pure Equilibria.

▶ **Definition 9.** *Define* $c^*(\vec{c}^R)$ *to be the unique solution to* $\sum_{i=1}^{n} \max\{0, 1 - \frac{c_i^R}{c^*(\vec{c}^R)}\} = 1$.[49] *Further define* $x_i^*(\vec{c}^R) := \max\{0, 1 - \frac{c_i^R}{c^*(\vec{c}^R)}\}$.

▶ **Theorem 10.** *Let* $(\vec{q}, \vec{r})$ *be an Equilibrium in the Distributed Ledger Model, and let the clearing price for End-Users be* r. *Then:*

- $\sum_j q_j = Q_A \cdot (r + B - c^W)/c^*(\vec{c}^R)$.
- *For all Miners* i, $\frac{q_i}{\sum_j q_j} = x_i^*(\vec{c}^R)$.

Moreover, $r \geq D_{\sup}^{-1}(Q_A)$, *and:*

- *If* $r = D_{\sup}^{-1}(Q_A)$, *then* $Q_A \cdot x_i^*(\vec{c}^R) \leq \frac{(x - c^W)(Q_A - D(x))}{x - r}$ *for all* i *and all* $x > r$.
- *If* $r > D_{\sup}^{-1}(Q_A)$, *then there is a single price-setter* i^*, *who sets a price equal to* $r_{i^*} := \arg\max_{x > D_{\sup}^{-1}(Q_A)}\{(x - c^W) \cdot (D(x) + Q_A \cdot x_{i^*}^*(\vec{c}^R) - Q_A)\}$. *If an equilibrium of this form exists, one certainly exists with* $i^* = 1$ *as the price-setter (equilibria with other price-setters are possible).*

Importantly, observe in Theorem 10 that:

- Block rewards play no role in the ultimate clearing price, nor the resulting market share of each Miner.[50]
- A necessary condition for an equilibrium to exist with clearing price $D_{\sup}^{-1}(Q_A)$ (the smallest possible clearing price) is only a function of x_1^*, Q_A, and $D(\cdot)$. That is, to the extent that a quantitative measure of decentralization plays a role in the ultimate price paid by end-users, it is the size of the largest Miner. Moreover, x_1^* can be determined *only as a function of* $\vec{c}^R$.
- That is, there is one term that depends only on $\vec{c}^R$ ($x_1^*(\vec{c}^R)$), and another that depends only on c^W and $D(\cdot)$ ($\sup_{x>r}\{\frac{(x-c^W)\cdot(Q_A - D(x))}{x-r}\}$), and an Equilibrium that clears all Q_A WRITEs made available by the protocol can plausibly exist if and only if $x_1^*(\vec{c}^R) \leq \sup_{x>r}\{\frac{(x-c^W)\cdot(Q_A - D(x))}{x-r}\}$.

3.3 A Sufficient Condition for Pure Equilibria

Theorem 10 characterizes all possible pure Equilibria, and so it is tempting to proceed with equilibrium analysis under these conditions. Unfortunately, pure Equilibria are not guaranteed to exist in the Distributed Ledger Model. In Section 3.3.1 we provide an example demonstrating this, and identify the barrier. This motivates a natural sufficient condition that we analyze in Section 3.3.2. Along the way, we also discuss the impact of block rewards on equilibria in Section 3.3.3.

3.3.1 An Example with with Non-Existence

Consider a demand curve $D(\cdot)$ with $D(x) = 1 - x$ for all $x \in [0, 1]$, $Q_A = 1$, $B = 0$, and $c^W = 0$. Consider also $n = 3$ miners each with $c_i^R = 1$. Then Theorem 10 concludes:

- In any potential equilibrium, it must hold that $Q_i(\vec{q}, \vec{r}) = 1/3$ for all i.
- Fixing $Q_1 = 1/3$, $\arg\max_{x \geq 0}\{x \cdot (1/3 - x)\} = 1/6$. Therefore, it is not possible to have an equilibrium with clearing price $r = D^{-1}(1) = 0$.[51]

[49] [3] establish that $c^*(\vec{c}^R)$ is well-defined – the proof is straight-forward.

[50] As noted previously, Block rewards do play a role in determining whether Pure Equilibria exist – see Section 3.3.3.

[51] Because in such an equilibrium, all three Miners earn profit zero, whereas any Miner could deviate to set a price of 1/6 and instead earn profit 1/36 (letting the other Miners earn profit 1/18).

- $c^*(1,1,1) = 3/2$. Therefore, the only potential equilibria have one Miner as a price-setter at price 1/6, with $\sum_j q_j = \frac{1/6}{3/2} = 1/9$ (and therefore each Miner has $q_i = 1/27$).
- However, this is not an equilibrium. In this strategy profile, the price-setter earns revenue 1/36 from the simultaneous first-price auctions, but pays 1/27 in Resource cost, yielding negative payoff. The price-setter would be better off not investing at all.

In particular, what stands out about this example is that a fully-saturated equilibrium has absolutely no shot (because such an equilibrium would generate zero revenue in total). This suggests a natural sufficient condition: that the fully-saturated equilibrium generate *some* fraction of the optimal revenue a monopolist could earn.

3.3.2 A Sufficient Condition

Theorem 11 below states an interpretable sufficient condition for an equilibrium with clearing price $r = D^{-1}(Q_A)$, and a technical condition that is necessary and sufficient when $B = 0$.

▶ **Theorem 11.** *Consider a potential equilibrium $(\vec{q}^*, \vec{r}^*)$ such that:*

- *The clearing price is $r = D^{-1}_{\sup}(Q_A)$.*
- *$\sum_j q^*_j = Q_A \cdot (r + B - c^W)/c^*(\vec{c}^R)$.*
- *For all Miners i, $\frac{q_i}{\sum_j q_j} = x^*_i(\vec{c}^R)$.*

Then:

- *If $D(\cdot)$ is Regular and $x^*_1(\vec{c}^R) \leq 1 - \frac{1}{D(0)/Q_A - 1}$, then $(\vec{q}^*, \vec{r}^*)$ is an Equilibrium.*
- *If $B = 0$, define $k(z) := \frac{D^{-1}_{\sup}(z \cdot Q_A) \cdot z \cdot Q_A}{D^{-1}_{\sup}(Q_A) \cdot Q_A}$. Then $(\vec{q}^*, \vec{r}^*)$ is an Equilibrium if and only if $x^*_1(\vec{c}^R) \leq 1 - \sup_{z \in [0,1]} \left\{ \frac{k(z)-1}{2 \cdot \left(\sqrt{k(z)/z} - 1\right)} \right\}$*

As previously noted, Theorem 11 provides interpretable sufficient conditions for Q_A WRITEs to clear in Equilibrium (in Bullet One), and necessary and sufficient conditions (when $B = 0$) in Bullet Two. In both cases, the condition depends only on $D(\cdot), Q_A$, and $x^*_1(\vec{c}^R)$, highlighting the "size of the largest miner" (which can be computed as a function only of $\vec{c}^R$) as the relevant "measure of decentralization" for determining the economic impact on end-users (in our model).

We can also use Theorem 11 to reason about modifications to the example in Section 3.3.1. Consider the case of n Miners each with $c^R_i = 1$, $D(x) = 1 - x$ for all $x \in [0,1]$, $B = 0$ and $c^W = 0$, but we will vary Q_A.

- Then Bullet One of Theorem 11 confirms it is an Equilibrium for Q_A WRITEs to clear (at price $1 - Q_A$) as long as $1/n \leq 1 - \frac{1}{1/Q_A - 1}$, which can be rewritten as $Q_A \leq \frac{n-1}{2n}$. For this particular example in Section 3.3.1, this is not particularly impressive, as even a Miner controlling blockchain would sell all Q_A APPENDs as long as $Q_A \leq 1/2$.
- The more precise Bullet Two of Theorem 11 could be applied. In this case, $k(z) = \frac{z \cdot (1 - z \cdot Q_A)}{1 - Q_A}$, and for example when $Q_A = 3/4$, $\sup_{z \in [0,1]} \left\{ \frac{k(z)-1}{2 \cdot \left(\sqrt{k(z)/z} - 1\right)} \right\} = 2/3$. Therefore, $n = 3$ identical Miners are sufficient in order for a Protocol with $Q_A = 3/4$ to have an Equilibrium where all Q_A WRITEs are sold (whereas $n = 1$ Miner is insufficient, as such a Miner would choose to monopoly-price and sell only 1/2).

3.3.3 Block Rewards Support Existence of Pure Equilibria

In this section, we reason about the role of block rewards on equilibria. This section has two main results, both below.

▶ **Proposition 12.** *Let $(\vec{q}^*, \vec{r}^*)$ induce a clearing price of r, and be an equilibrium of the Distributed Ledger Model with Resource costs $\vec{c}_R$, WRITE cost c^W, demand curve $D(\cdot)$, and block reward B. Then $((1 + \frac{B'-B}{Q_A \cdot (r - c^W + B)}) \cdot \vec{q}^*, \vec{r}^*)$ is an equilibrium of the Distributed ledger Model with Resource costs $\vec{c}_R$, WRITE cost c^W, demand curve $D(\cdot)$, and block reward $B' > B$.*

Proposition 12 establishes that higher block rewards support the existence of pure equilibria (although Theorem 10 establishes that block rewards do not impact potential pure equilibria themselves). This is of standalone interest for understanding the impact of modeling parameters on end-users, and also a technical ingredient in the proof of Theorem 11.

Additionally, we show that if $\vec{c}^R$ does not immediately rule out a market-clearing equilibrium, there is a sufficiently large block reward so that a market-clearing equilibrium exists.

▶ **Proposition 13.** *Let $\vec{c}^R, D(\cdot), Q_A, c^W, Q_A$ be such that $Q_A \cdot x_i^*(\vec{c}^R) < \inf_{x > D_{\sup}^{-1}(Q_A)} \{\frac{(x - c^W) \cdot (Q_A - D(x))}{x - D_{\sup}^{-1}(Q_A)}\}$. Then, there exists a sufficiently large $B < \infty$ such that a market-clearing equilibrium exists in the market defined by $\vec{c}^R, D(\cdot), B, Q_A, c^W$.*

4 Conclusion

We investigate decentralization as a means to insulate a natural monopoly from its derivative services. We then draw conclusions regarding the ultimate utility of end-users as a function of protocol parameters. We further highlight the impact of various aspects on our conclusions: (a) our analysis applies to patient users (who are content with purchase from any miner) and not impatient users (who view each miner as a monopolist anyway), (b) the relevant "measure of decentralization" for impact on users' price is the size of the largest miner (which can be determined exclusively as a function of the profile of Resource investment costs, as in a pure Tullock contest), (c) block rewards don't impact users' price at equilibrium, but can influence whether equilibria exist.

Within distributed ledgers, our model considers the basic setup where users directly interact with miners to include a transaction on the blockchain. Of course, many blockchain ecosystems have evolved and now include additional parties (Builders and Layer-2s are two notable examples). Our work provides a framework through which to ask: how does the presence of these parties ultimately impact the service users receive?

Beyond our model, it is also important to endogenize aspects that our model treats as exogenous. For example, our model treats as exogenous the fact that Bitcoin/Ethereum run a Tullock contest in Computation/Stake. While this is an accurate representation of all major blockchains since their inception, and Upstream protocol mechanics change *much* more slowly than strategic Downstream decisions, these protocols are not truly exogenous – protocol rules are set by some governance process (perhaps formally specified, perhaps not). It is therefore an important direction for future work to additionally model the dynamics by which protocol mechanics are determined. Additionally, our model treats the "lines" between distinct miners/stakers are exogenous. While it is again the case that miners/stakers merge at a *much* slower pace than adapting strategic Downstream decisions, miner/staker identities are not exogenously fixed – parties might certainly merge and jointly strategize if they find it beneficial. In particular, if all parties were to merge/collude or otherwise jointly strategize, they could profit by setting monopoly prices, so it is important to additionally model the process by which miners/stakers might merge, collude, or otherwise jointly strategize.

Beyond distributed ledgers, there are many natural monopolies for digital services that pose challenges for traditional regulatory approaches [48]. Our results provide theoretical

foundations for exploring decentralized protocols as a tool that *might* prove useful to insulate such natural monopolies. Future work could, for example, consider decentralized protocols that manage marketplaces (allowing competition among matchmakers, search, etc.) or store social network data (allowing competition among UIs, content moderation, etc.). Our work motivates further investigation of decentralized protocols as a means to insulate natural monopolies in other domains in a detailed manner that can both (a) identify which natural monopolies might be amenable to insulation by a decentralized protocol, and (b) provide a complete analysis linking design choices of the decentralized protocol to impact on users.

Beyond insulating natural monopolies, our work contributes to an emerging line of works seeking theoretical foundations for the impact of decentralized systems on users [29, 46, 25, 41]. Significant further work along these lines is necessary in order to understand domains where decentralized systems have a shot at providing lasting value, even in the presence of highly-developed incumbents.

References

1. Kaya Alpturer and S Matthew Weinberg. Optimal randao manipulation in ethereum. In *6th Conference on Advances in Financial Technologies*, 2024.
2. Humoud Alsabah and Agostino Capponi. Pitfalls of bitcoin's proof-of-work: R&d arms race and mining centralization. *Available at SSRN 3273982*, 2020.
3. Nick Arnosti and S Matthew Weinberg. Bitcoin: A natural oligopoly. *Management Science*, 68(7):4755–4771, 2022.
4. Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. Transaction fee mechanism design with active block producers. In *International Conference on Financial Cryptography and Data Security*, pages 85–90. Springer, 2024.
5. Maryam Bahrani and S Matthew Weinberg. Undetectable selfish mining. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1017–1044, 2024.
6. Soumya Basu, David Easley, Maureen O'Hara, and Emin Gün Sirer. Towards a functional fee market for cryptocurrencies. *arXiv preprint arXiv:1901.06830*, 2019.
7. Jonah Brown-Cohen, Arvind Narayanan, Alexandros Psomas, and S Matthew Weinberg. Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 459–473, 2019.
8. James McGill Buchanan, Robert D. Tollison, and Gordon Tullock. Toward a theory of the rent-seeking society. *Southern Economic Journal*, 48:823, 1982. URL: `https://api.semanticscholar.org/CorpusID:153517516`.
9. Eric Budish, Andrew Lewis-Pye, and Tim Roughgarden. The economic limits of permissionless consensus. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 704–731, 2024.
10. Eric B Budish. Trust at scale: The economic limits of cryptocurrencies and blockchains. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (83), 2022.
11. Linda Cai, Jingyi Liu, S Matthew Weinberg, and Chenghan Zhou. Profitable manipulations of cryptographic self-selection are statistically detectable. In *6th Conference on Advances in Financial Technologies (AFT 2024)*, pages 30–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.
12. Miles Carlsten, Harry Kalodner, S Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 154–167, 2016.
13. Xi Chen, David Simchi-Levi, Zishuo Zhao, and Yuan Zhou. Bayesian mechanism design for blockchain transaction fee allocation. *Operations Research*, 2025.
14. Hao Chung, Tim Roughgarden, and Elaine Shi. Collusion-resilience in transaction fee mechanism design. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1045–1073, 2024.

15 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899. SIAM, 2023.

16 Nicola Dimitri. Bitcoin mining as a contest. *Ledger*, 2:31–37, 2017.

17 Ittay Eyal. The miner's dilemma. In *2015 IEEE symposium on security and privacy*, pages 89–103. IEEE, 2015.

18 Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. *Communications of the ACM*, 61(7):95–102, 2018.

19 Matheus VX Ferreira, Aadityan Ganesh, Jack Hourigan, Hannah Huh, S Matthew Weinberg, and Catherine Yu. Computing optimal manipulations in cryptographic self-selection proof-of-stake protocols. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 676–702, 2024.

20 Matheus VX Ferreira, Ye Lin Sally Hahn, S Matthew Weinberg, and Catherine Yu. Optimal strategic mining against cryptographic self-selection in proof-of-stake. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 89–114, 2022.

21 Matheus VX Ferreira, Daniel J Moroz, David C Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 86–99, 2021.

22 Amos Fiat, Anna Karlin, Elias Koutsoupias, and Christos Papadimitriou. Energy equilibria in proof-of-work mining. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 489–502, 2019.

23 Yotam Gafni and Aviv Yaish. Barriers to collusion-resistant transaction fee mechanisms. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1074–1096, 2024.

24 Aadityan Ganesh, Clayton Thomas, and S Matthew Weinberg. Revisiting the primitives of transaction fee mechanism design. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 703–703, 2024.

25 Itay Goldstein, Deeksha Gupta, and Ruslan Sverchkov. Utility tokens as a commitment to competition. *The Journal of Finance*, 79(6):4197–4246, 2024.

26 Guy Goren and Alexander Spiegelman. Mind the mining. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 475–487, 2019.

27 Mark Gradstein. Intensity of competition, entry and entry deterrence in rent seeking contests. *Economics & Politics*, 7(1):79–91, 1995.

28 Arye L Hillman and John G Riley. Politically contestable rents and transfers. *Economics & Politics*, 1(1):17–39, 1989.

29 Gur Huberman, Jacob D Leshno, and Ciamac Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *The Review of Economic Studies*, 88(6):3011–3040, 2021.

30 Aggelos Kiayias, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis. Blockchain mining games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 365–382, 2016.

31 Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin's fee market. *ACM Transactions on Economics and Computation*, 10(1):1–31, 2022.

32 Stefanos Leonardos, Barnabé Monnot, Daniël Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras. Dynamical analysis of the eip-1559 ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 114–126, 2021.

33 Stefanos Leonardos, Daniël Reijsbergen, Barnabé Monnot, and Georgios Piliouras. Optimality despite chaos in fee markets. In *International Conference on Financial Cryptography and Data Security*, pages 346–362. Springer, 2023.

34 Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

35 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Satoshi Nakamoto*, 2008.

36 Joachim Neu, Ertem Nusret Tas, and David Tse. The availability-accountability dilemma and its resolution via accountability gadgets. In Ittay Eyal and Juan A. Garay, editors, *Financial Cryptography and Data Security - 26th International Conference, FC 2022, Grenada, May 2-6, 2022, Revised Selected Papers*, volume 13411 of *Lecture Notes in Computer Science*, pages 541–559. Springer, 2022. doi:10.1007/978-3-031-18283-9_27.

37 Michael Neuder, Daniel J. Moroz, Rithvik Rao, and David C. Parkes. Defending against malicious reorgs in tezos proof-of-stake. In *AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY, USA, October 21-23, 2020*, pages 46–58. ACM, 2020. doi:10.1145/3419614.3423265.

38 Michael Neuder, Daniel J. Moroz, Rithvik Rao, and David C. Parkes. Low-cost attacks on ethereum 2.0 by sub-1/3 stakeholders. *CoRR*, abs/2102.02247, 2021. URL: https://arxiv.org/abs/2102.02247, arXiv:2102.02247.

39 Stephen H Newman. Decentralization cheapens corruptive majority attacks. *arXiv preprint arXiv:2310.01546*, 2023.

40 Noam Nisan. Serial monopoly on blockchains. *CoRR*, abs/2311.12731, 2023. URL: https://doi.org/10.48550/arXiv.2311.12731, arXiv:2311.12731, doi:10.48550/ARXIV.2311.12731.

41 Marco Reuter. *Platform Precommitment via Decentralization*. International Monetary Fund, 2024.

42 Tim Roughgarden. Transaction fee mechanism design. *Journal of the ACM*, 71(4):1–25, 2024.

43 Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *Financial Cryptography and Data Security: 20th International Conference, FC 2016, Christ Church, Barbados, February 22–26, 2016, Revised Selected Papers 20*, pages 515–532. Springer, 2017.

44 Richard Schmalensee, Mark Armstrong, and Robert D Willig. *Handbook of industrial organization*, volume 3. Elsevier, 1989.

45 Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design? In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2023.

46 Michael Sockin and Wei Xiong. Decentralization through tokenization. *The Journal of Finance*, 78(1):247–299, 2023.

47 Jean Tirole. *The theory of industrial organization*. MIT press, 1988.

48 Jean Tirole. Competition and the industrial challenge for the digital age. *Annual Review of Economics*, 15(1):573–605, 2023.

49 Gordon Tullock et al. Efficient rent seeking. *Toward a theory of the rent-seeking society*, 97:112, 1980.

50 Ke Wu, Elaine Shi, and Hao Chung. Maximizing miner revenue in transaction fee mechanism design. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2024.

51 Aviv Yaish, Gilad Stern, and Aviv Zohar. Uncle maker:(time) stamping out the competition in ethereum. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 135–149, 2023.

52 Aviv Yaish, Saar Tochner, and Aviv Zohar. Blockchain stretching & squeezing: Manipulating time for your best interest. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 65–88, 2022.

53 Andrew Chi-Chih Yao. An incentive analysis of some bitcoin fee designs. *arXiv preprint arXiv:1811.02351*, 2018.

54 Roi Bar Zur, Ittay Eyal, and Aviv Tamar. Efficient mdp analysis for selfish-mining in blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 113–131, 2020.