Transaction Fees and Seigniorage in Proof-of-Work Cryptocurrencies*

Michele Fabi†

July 2022

Abstract

I study the optimal design of transaction fees and seigniorage for a Proof-of-Work cryptocurrency. Commodity merchants need blockchain miners to record their payments and secure the blockchain by remaining active. Fees make miners willing to process merchants' transactions by compensating for the risk that doing so slows down block transmissions making blocks invalid. Seigniorage convinces miners to operate when pending transactions are scarce. Both seigniorage and fees are necessary. As fees are distortionary for merchants, an optimal design uses them only as required by incentive-compatibility.

JEL Codes: C73, D47, E42, G11.

Keywords: cryptocurrencies, proof-of-work, blockchain, miners, seigniorage.

*I am especially thankful to Matthew Ellman for guiding me with care and dedication throughout the development of this paper. I also thank Yackolley Amoussou-Guenou, Jordi Caballé, Ramon Caminal, Jonathan Chiu, Xavier Cuadras, Marc Escrihuela, Hanna Halaburda, Gur Huberman, Sjaak Hurkens, Joachim Jungherr, Hannes Mueller, Chara Papioti, Amedeo Piolatto, Julien Prat, Hugo Rodriguez, Natkamon Tovanich for valuable discussions. More in general, I thank the research staff at ENSAE, IAE, UAB, and the audience at EARIE2021 and JEI2021. This paper benefited from the help of the online community at www.stackexchange.com which I sincerely thank. Finally, I gratefully acknowledge financial support from the Blockchain and Platform Chair, Barcelona GSE, and the Spanish Ministry of Science, Innovation and Universities (MCIU)—FPI fellowship 914886-79792965.

†Ecole Politechnique, CREST, IP Paris (michele.fabi@ensae.fr)

1 Introduction

A *cryptocurrency* is a digital currency governed by algorithms and managed by a decentralized, free-entry network of so-called *miners*. These subjects act as cryptocurrency accountants, keeping track of "who owns what".

Each cryptocurrency is joined at the hip to a *blockchain*, a public ledger secured by cryptographic techniques that miners share to record transactions and algorithmic instructions (smart contracts). To update the blockchain, miners use dedicated computers to form a *mempool* (memory pool) of pending operations and then record these operations on the blockchain, batched within time-stamped data blocks. The concatenation of blocks in chronological order establishes a timeline of events and an operation registry.

The blockchain is a distributed ledger in that it is split into separate copies spread across the miner network. Each miner is in charge of updating its own copy with new recorded blocks and communicating the new blocks it records to the other miners, allowing them to comply with updates. All blockchain communities rely on a *protocol*, a set of algorithmic rules programmed in miners' machines and written codes of conduct that miners are supposed to follow to reach consensus over the state of the ledger. However, given the lack of a central authority that can enforce them, miners' behavior outside hard-coded rules is purely driven by incentives.

Cryptocurrencies such as Bitcoin, its numerous spin-offs, and the current version of Ethereum run on Proof-of-Work (PoW) blockchains. To make PoW blockchains tamper-proof, miners have to solve a costly cryptopuzzle (a cryptographic puzzle) for each block they form before recording the block and transmitting it to the rest of the miners. The activity of solving these cryptopuzzles is termed *blockchain mining*. When a miner *mines a block* by solving its respective puzzle, it can show the solution as proof of the computational work it employed to secure the blockchain. In this sense, a cryptopuzzle solution is a "proof of work" for the block it refers to. The *block time*, i.e. the time it takes miners to find a proof of work, is kept constant on average by an algorithm based on the aggregate miner computational power (CPU). Therefore, the block rate, i.e. the average number of puzzle solutions per unit of time, is also constant.

To compensate miners for the costs of the PoW mechanism, the protocol rewards them with two sources of revenue for each block successfully recorded and transmitted: fees collected from the operations—that I will simply call 'transactions' afterwards—it records and seigniorage from the creation of new coins.[1] The structure of these operational revenues not only determines the miner network's size via free-entry but also how and when miners choose to stay active and record pending transactions on the blockchain. In turn, miner behavior affects whether merchants that use the cryptocurrency as a means

[1]The name 'miners' indeed comes from the fact that they are in some sense mining these new coins out of the blocks as if they were "digital gold".

of payment can rely on the blockchain to record trades.

In this paper, I characterize the properties of the optimal combination of miner seigniorage and fees. I use the Lagos-Wright (LW) framework in continuous time to model a single-cryptocurrency economy populated by miners and merchants, these last further divided into buyers and sellers.

Buyers meet sporadically with sellers of a special good to trade in a decentralized market (DM), but the only way they can do so is by holding the cryptocurrency in advance. So buyers must first work to produce generic good units and exchange them in a centralized market (CM) for crypto coins (or *crypto* in brief). As soon as buyers' payments are recorded on the blockchain, sellers exchange the received coins in the CM to consume the corresponding generic good value.

Miners' role in this story is to record buyers' payments to sellers so that trade can occur. Miner participation and activity will determine whether sellers trust accepting payments in crypto. The more miners actively join, the higher the computational power spent on the PoW security mechanism. Moreover, thorough the mining process, miners will also determine the time it takes buyers' payments to reach sellers. In fact, the confirmation discount factor that sellers apply to buyers' payments discounts time from the moment these payments enter the mempool to the time miners record them on the blockchain. This discount factor in turn pins down merchants' terms of trade and buyers' demand for crypto. But crypto demand also influences the mining process, creating a feedback system among the worlds of miners and merchants. The two facets of the economy link-up thanks to the endogenous steady state of the mempool, which allows to compute all the equilibrium variables for fixed design parameters.

Within this general equilibrium, transaction fees and seigniorage play a crucial role on the miners' side of the economy. The crux of my analysis is to understand how the protocol can use them to incentivize miners to operate most efficiently.

I identify three fundamental incentive problems that the cryptocurrency design must tackle: First, it has to induce miners to willingly record merchants' pending transactions; then, it has to convince new miners to join until secure participation is reached; and finally, it has to motivate miners to remain active steadily over time. The first task is needed for trade to take place in the first place (otherwise, payments are never settled). The remaining tasks guarantee blockchain security, which requires not only enough miners to join but also that they remain constantly active to prevent potential attackers from taking advantage of temporary inactivity. I consider this security requirement in the form of an exogenous constraint on minimum mining activity.

The constraints on blockchain protocol's tuning that emerge from these three challenges are what I call Recording Constraint (RC), Participation constraint (PC), and Activity Constraint (AC).

The issue tackled by the Recording Constraint is that non-empty blocks with at least

one transaction face a block invalidation risk that is instead (practically) absent in empty blocks.[2] This creates thus a record-keeping disincentive.

The invalidation risk stems from miners' competition to simultaneously extend the top of the blockchain: Only the quickest miner in mining and transmitting the next block to the other miners updates the ledger. The later transmitted blocks will still appear in the blockchain, but as part of *forks* (i.e., ramifications) of the registry that miners and merchants ignore.

In this competition for speed, a block's size (i.e., the amount of digital information stored) does not affect its block time, but it does slow down transmission speed. Therefore, a miner might want to deviate from an equilibrium where it adds transactions to blocks by instead keeping blocks empty and transmitting them faster. By doing so, the deviating miner gains a competitive advantage in the mining race while still earning seigniorage from new coins (which is also awarded to empty blocks). However, in this way it also forgoes the fees attached to the unrecorded transactions.

It turns out that miners will find this deviation profitable unless the amount of those fees is conspicuous relative to seigniorage. The protocol can then use transaction fees as the carrot for convincing miners to record transactions. Indeed, the Recording Constraint will require positive transaction fees and bounds from below the fees-to-seigniorage ratio. For tractability, I present the block-size tradeoff assuming binary blocks: empty blocks carry no transaction but transmit immediately; full blocks carry one transaction and require an exponentially distributed transmission time.

In Bitcoin, the invalidation risk is not really a concern since the block transmission delays are negligible compared to the average block time: Mining a block takes on average 10 minutes, while block transmission takes about 10 seconds (Decker and Wattenhofer, 2013). Hence, the impact of block size on the invalidation probability is negligible. Nevertheless, slow transmissions are more dangerous in faster blockchains such as Ethereum, which has an average block time of about 10 seconds. In this case, 10-second transmission times would cause about half of the transmitted blocks to get wasted (Buterin, 2014). Ethereum's developers already recognize the importance of using transaction fees to cover the invalidation risk, also referred to as "orphaning risk" (Buterin, 2021). Furthermore, faster blockchains that are now proliferating the crypto space to host Decentralized Finance (DeFi) technologies require short block times in order to be meaningful alternatives or complements to standard (centralized) financial technologies. The problem of coping with transmission delays is thus not just a mere intellectual curiosity; it has a very concrete practical relevance.

Moving on, the problem of inducing steady miner activity is different in nature. This time it can be solved by miner seigniorage. The origin of the incentive problem here is

[2]By 'empty block' I refer to a block that does not record any merchant transaction. Empty blocks still contain miner seigniorage rewards.

the stochasticity of fee revenues, which fluctuate depending on merchants' trade intensity. When few or no transactions are pending in the mempool, miners cannot earn much from fees, yet they consume energy to keep mining blocks. The only reason for them not to switch off is to gain seigniorage. As a result, the Activity Constraint imposes minimum positive seigniorage.

In practice, the big miners are private companies operating with large computer farms that are difficult to switch back and forth between off and on swiftly. Even under this technical limitation, providing an amount of seigniorage that is at least commensurate to the energy cost of mining is still needed to avoid that miners remain inactive during a long crypto bear market.

At this point, we can already see that the protocol must provide miners with both positive transaction fees and seigniorage. Taking care of each of the previous incentive problems required restricting appropriately one of the two revenue instruments. On the other hand, achieving secure miner entry requires the protocol to design miners' lifetime revenue stream as the number of participating miners will equal the ratio of these revenues to mining costs, both evaluated at present value. The Participation Constraint requires that entry so determined meets the exogenous safety threshold.

The last analysis section of this paper uses all the above to set up the optimal cryptocurrency design problem that a Social Planner solves by choosing how to best tax merchants with seigniorage and transaction fees to subsidize miners. The planner can set a fixed pro-rata transaction fee rate algorithmically as it occurs in the Ethereum blockchain after its EIP-1559 upgrade.

I find that seigniorage is a more efficient tax than transaction fees. While seigniorage is a non-distortionary transfer from buyers to miners, transaction fees worsen the terms of trade that buyers face in the decentralized market. The welfare cost of the inefficiency is fully accounted by the increased cost in generic good production that buyers have to pay to acquire the needed liquidity. The ideal miner revenue structure would then only use seigniorage, but this is not incentive-compatible. Hence the second-best option is to use fees only inasmuch as required by the binding Recording Constraint.

The model's parameters affect the relative weight of tax instruments directly through the Recording Constraint. Some are neutral, such as the frequency at which buyers and sellers meet. Others are not: an increased block rate shifts weight from seigniorage to fees owing to increased competition in miners' speed tournaments; quicker block transmissions raise the weight on seigniorage for the opposite reason. I highlight these comparative statics results with their implied welfare effects in a numerical exercise.

The paper concludes with a discussion of possible directions for future research and broader applicability of its results. Despite its focus on PoW blockchains, I argue that some of this paper's tradeoffs are to some extent applicable to blockchain of the Proof-of-Stake (PoS) kind. These last are based on an eco-friendly alternative to PoW that is

quickly becoming the dominant blockchain design paradigm.

.

2 Literature Review

This paper contributes to the recent, rapidly growing literature on blockchains and cryptocurrencies. The model I propose wraps miners' and merchants' sub-economies within a general equilibrium framework inspired by the continuous-time Lagos-Wright (LW) model proposed by Choi and Rocheteau (2020b), which is particularly suited for cryptocurrency economies. Lagos et al. (2014) provides a thorough review of the new monetarist economics literature that pioneered my approach.

Other authors already treated record-keeping and security aspects of PoW cryptocurrencies. Nevertheless, this work is the first one that bridges them coherently and studies the novel implications originating from their interplay.

The closest paper is Chiu and Koeppl (2019), to my knowledge the first proposing a general equilibrium model of cryptocurrency (i.e. Bitcoin) adoption. In Chiu and Koeppl's paper, buyers, who are also miners, can engage in fraud by reverting their payments to sellers back to themselves by means of double-spending attacks. As in my paper, Chiu and Koeppl show that security threats pose a minimum requirement on miner participation that the planner must satisfy raising an adequate level of miner revenues taxing merchants. A fundamental difference with respect to my approach is that I consider the secure level of miner participation (and activity) as exogenously given whereas these authors derive it within their model. On the other hand, I introduce explicit mempool dynamics and record-keeping incentives these authors do not consider.

The main normative finding of Chiu and Koeppl (2019) is that an optimal cryptocurrency design rewards miners purely with seigniorage. This result is based on two arguments. First, since the inflation tax resulting from miner seigniorage is charged on all coins in circulation while transaction fees are charged only on traded coins, seigniorage reaches a more extensive monetary base. Hence, it allows the planner to smooth merchants' tax burden over smaller, less distortionary payments than those implied by transaction fees.

Second, transactions fees do not disincentivize double spending: A malicious buyer pays no transaction fees if a double-spending attack succeeds since he controls the miner that ultimately receives those fees.

I also find that pure-seigniorage is optimal but through different model mechanics. The intuition for the inefficiency of fees in my story is that they distort merchants' terms of trade in the DM. On the contrary, terms of trade are unaffected by seigniorage which, except potentially for considerations outside my model, is just a lump-sum transfer from merchants (i.e. buyers) to miners. The advantage of lump-sum taxes is already

well understood from classical taxation theories, which point out that optimal taxes are inversely related to elasticity of demand (Ramsey, 1927). This rationale suffices to explain why fees are less effective than seigniorage. Furthermore, even though my analysis does not consider the effect of taxation on double-spending incentives, doing so would only reinforce my argument.

A striking feature of my model is that direct trade distortions are the *only* reason why fees are inefficient. It is initially surprising to see that the tax-base advantage of seigniorage encountered by Chiu and Koeppl completely disappears in my model. Yet, it can be easily explained by the fact that mempool's endogeneity neutralizes such effect. The precise key observation to make is that the mempool's steady-state links available fee income to miners through congestion in the queue of pending transactions. This results in the whole monetary base contributing in expectation to both fees and seigniorage income.

The recording constraint introduces a final twist with respect to Chiu and Koeppl's optimal design receipt. As the zero-fees solution is not implementable, the planner uses some transaction fees but keeps them at a minimum.

Pagnotta (2022) is also close to my work. In Pagnotta's LW model, the feedback system among bitcoin price and blockchain security gives rise to multiple stationary equilibria ranked by price-security pairs. Due to this feature, security can be understood as intrinsic to the bitcoin price, in contrast to standard (centralized) monetary systems in which security is extrinsic, i.e. independent of the price of money.

As in my paper, this related work takes into account aggregate security threats in the form of generic sabotage attacks, but there the security function that quantifies blockchain resilience varies continuously in miners' participating CPU (in non-trivial cases). In my formulation instead, security is dichotomous. Hence conditional on miner CPU satisfying my Participation Constraint, security can be treated *as if* it were extrinsic, resulting in a single non-banal stationary equilibrium sustaining a positive crypto price.[3]

Although most other related works including Pagnotta (2022) find large equilibrium sets featuring more sophisticated price dynamics, e.g. boom-burst, they also find at least one equilibrium consistent with steady inflationary prices, this way justifying my focus. Fernández-Villaverde and Sanches (2019) and Choi and Rocheteau (2020a) do so in models of competing private monies; Schilling and Uhlig (2019) in a two-currency economy with bitcoin and US dollar coexisting.

Huberman et al. (2019) and Easley et al. (2019) offer an in-depth analysis of Bitcoin transaction fees. In these papers, mempool dynamics are controlled by miners who record transactions giving priority to those carrying higher fees. While in these models users (equivalent to buyers in my model) determine these fees in a first-price auction, I instead assume that the protocol sets a unique transaction fee rate.

These papers' formalism is suited to address the determination of fees in Bitcoin but

[3]I ignore the no-adoption equilibrium with the crypto price at zero.

treats crypto demand as exogenous and uses linear impatience. These features preclude seigniorage from being a policy variable, plus prevent direct incorporation of the queuing-theory techniques they employ into infinite-horizon general equilibrium models that use exponential discounting. My model closes the divide by embedding the mempool's queue originating from a continuous-time miner game into the new-monetarist approach à la Choi and Rocheteau (2020b). As in Easley et al. (2019), I restrict the maximum block size to unity when investigating miners' optimal block size choice.

Houy (2016) provides the first game-theoretical examination of the tradeoff between block size and invalidation risks. He finds that block size is increasing in the ratio of fees to seigniorage. The analogous condition is needed in my model for miners to produce full blocks. Houy solves the Nash equilibrium block size explicitly treating it as a continuous variable only for the case of two miners. My paper uses binary blocks but solves in closed form the focal incentive constraint for an arbitrary number of miners together with free-entry participation. Lehar and Parlour (2020) and Malik et al. (2022) also endogenize block size but use miner collusion rather than invalidation risks as fundamental driver.

Me and Houy (2016) take for granted that miners act strategically within the PoW protocol when forming blocks strategically, thereby obeying the Longest Chain Rule (LCR) of mining on top of the faster longest branch of the blockchain. Biais et al. (2019) show that miners follow the LCR in one among multiple equilibria, with some featuring persistent forks. They also explain how a single transmission delay can lead to a (one block long) temporary fork. In my model, temporary forks can take place repeatedly owing to stochastic block transmission times. When a temporary fork occurs, these authors find that forking-out the last received block is a valid prediction, although an equilibrium with miners discarding instead the first-received block is equally valid. Again, this feature is not present in my contribution as I impose LCR exogenously.

Broadly related to my papers are those studying the blockchain mining industry. Prat and Walter (2018) estimate industry dynamics of Bitcoin mining and its relationship with the bitcoin price; Cong et al. (2019) demonstrate that risk-averse miners form mining pools. The presence of big mining pools turns the previously competitive miner market into a strategic one.

Finally, my paper relates to the computer science literature on blockchains. The PoW protocol is introduced by Nakamoto (2008), whose "Section 11: Calculations" has been subsequently corrected by Rosenfeld (2014) and later contributions. Liu et al. (2022) justifies empirically my approach with data showing that block size in Ethereum got heterogeneous after the introduction of EIP-1559, meaning that Ethereum miners started to choose block size strategically. Decker and Wattenhofer (2013); Neudecker and Hartenstein (2019) study block propagation and temporary forks in Bitcoin. Carlsten et al. (2016); Tsabary and Eyal (2018) anticipated the logic of my activity constraint, arguing that fee stochasticity creates security breaches. A solid research line on Nakamoto

consensus points out that block times have to be long relative to transmission delays in order to safeguard the blockchain against attacks. This assumption is required by the security microfoundation of my model. Pass and Shi (2017) and other early contributions prove this result assuming bounded transmission delays. Sankagiri et al. (2021) covers the exponentially distributed delays that I use.

3 Merchants' Economy

I start to build my model by describing the real sector where merchants trade. To begin with, consider an economy populated by infinitely-lived agents. Time $t \in \mathbb{R}_+$ is continuous and is discounted by all agents at rate r. Equilibria will be stationary or Markovian, so I will sometimes drop the time index when clear from the context.

The inhabitants of the economy are a fixed measure of buyers $N \in \mathbb{R}_+$, a unit measure of sellers, and an endogenous number of miners $M \in \mathbb{N}_0$ determined by free entry.[4] I will refer to a specific buyer, seller, or miner with the pronoun 'he', 'she', 'it' respectively.

Two types of perishable and divisible goods are available. The first is a generic, numéraire good denoted by $x_t \in \mathbb{R}$; $x_t > 0$ if consumed, $x_t < 0$ if produced. The second is a special good whose consumption and production is denoted by $y_t \in \mathbb{R}$ and follows the same sign convention. The special good can be interpreted as a physical or digital good that can be traded at favorable market conditions using crypto or as the access to a Web3 service running on the blockchain that can be only acquired by exchanging crypto for specific utility tokens.[5] All agents enjoy one-to-one utility from consuming the generic good and disutility from producing it. Production disutility can equally be interpreted as disutility from labor.

Preferences and technology for the special good instead differ across agents. Buyers cannot produce the special good but enjoy consuming it according to a regular utility function $u(y)$. Sellers do not derive utility from its consumption but can produce y units paying a one-to-one production cost. Miners neither produce the special good, nor value its consumption.

The functional form for $u(y)$ that I will use is the generalized logarithmic utility function

$$u(y) = \ln(1 + \eta y), \quad \eta \in \mathbb{R}_+ \tag{1}$$

This ensures $u(0) = 0$ and avoids a situation where even optimal trade would yield a negative utility via taste parameter η sufficiently large. Lagos and Wright (2005), Chiu

[4] There is no need to use a notation for the seller population since in equilibrium their rent will be fully taken by buyers.

[5] These services comprehend decentralized digital storage (Filecoin); blockchain data extraction (The Graph); tokenized carbon footprint trackers (Climate Trade); videogames and E-sport (Dogami, ConsenSys). Although financial applications of the blockchain are currently the most popular, non-financial applications are now mushrooming. Soon they will likely gain more popularity.

and Koeppl (2019), Pagnotta (2022) use similar functional forms.[6]

All agents can also obtain storable and perfectly divisible coins of a PoW cryptocurrency with no intrinsic consumption value. In Ethereum, the units of these coins range in value from "wei" to "ether" (10^{18} wei); in Bitcoin, from "satoshi" to "bitcoin" (10^{8} satoshis). I let z_t denote the value of an agent's crypto portfolio at time t in terms of the numéraire. The blockchain protocol employs a hard-coded inflationary money supply rule such that the new coins printed in each block debase the value of crypto by an inflation factor π; the portfolio value reduces by $z_t\pi$ when a new block recorded on the blockchain at time t.

Miners are the only agents that can produce the cryptocurrency. They do so whenever they record a new block that is accepted by the miner network's consensus. The counting of recorded blocks $B_t \in \mathbb{N}_0$ follows approximately a Poisson process at the rate $\mu \in \mathbb{R}_+$ at which miners find a PoW—I will clarify what I mean by 'approximately' in Section 4.4. The protocol keeps algorithmically the aggregate PoW rate constant regardless of miners' CPU.

Miners inject crypto in the economy through a continuously open Centralized Market (CM), e.g. a centralized exchange such as Binance, Coinbase, or Kraken. The CM acts as a market-maker taking crypto orders against the generic good. Each agent can use the CM to liquidate crypto and consume generic good, or to produce generic good and acquire crypto. (Remember that the generic good is non-storable so it has to be produced on the spot). The CM can process agents' operations instantaneously by endowing them with custodial wallets that process transactions outside the blockchain (off-chain).

Besides trading in the CM, merchants meet sporadically in a Decentralized matching Market (DM). Here, each buyer meets a generic seller at a Poisson rate $\alpha \in \mathbb{R}_+$ to make a take-it-or-leave-it offer for the special good. As a result of technological asymmetries, DM meetings are the only occasion for buyers to buy the special good from sellers.

To make the model interesting, merchants cannot produce the generic good while busy bargaining. Furthermore, the DM is anonymous (or better, pseudo-anonymous) so that buyers cannot take a credit to finance their operations with sellers.[7] As a result, the only way buyers can trade with sellers in the DM is by using crypto previously bought in the CM. As long as crypto has a positive equilibrium value, sellers are willing to accept it and subsequently liquidate it in the CM in exchange for the generic good. This feature makes the cryptocurrency essential, in the sense that it widens the frontier of welfare-improving trade arrangements.

[6]The generalized CRRA used in both Lagos and Wright (2005) and Pagnotta (2022) is equivalent to $u(y) = (1-\epsilon)^{-1}[(y+1/\eta)^{(1-\epsilon)} - (1/\eta)^{(1-\epsilon)}]$, which converges to the utility function (1) for $\epsilon \to 1$.

[7]Even in these cases merchants' identity can be retrieved indirectly by analyzing the blockchain tree. For example, Alyssa Blackburn and Erez Lieberman Aiden were able to reconstruct a map of Bitcoin transactions spanning from 2009 to 2011 by analyzing data leakages (Roberts, 2022). Also, some years before, the US Federal Bureau of Investigation (FBI) was able to trace the identity of most people involved in illicit trade through the website Silk Road.

Contrarily to CM transactions, those taking place in the DM have to be recorded on the blockchain (on-chain) by miners to be considered confirmed. Confirmation is not immediate and depends on miners' activity, so sellers discount the value of the crypto payments they receive using a *confirmation discount factor* $\beta \in [0, 1)$ that I will later determine endogenously in Section 5. Moreover, each time they pay a seller, buyers are charged with a transaction fee τ that goes to the miner that records the payment on the blockchain as a remuneration for its effort. The gain from trade of a seller from producing y units of special good and receiving a gross payment of z coins is therefore $\beta z(1-\tau) - y$. Setting this gain to zero determines the amount $y = \beta z(1-\tau)$ of special good a buyer with crypto balances z can acquire from sellers. Given that the cryptocurrency is inflationary, the only reason buyers hold it is to trade, so they give away all their coins in a DM meeting and replenish their pockets at once after when rejoining the CM.

3.1 Merchants' value functions and crypto portfolio

With the elements introduced in the preceding part of Section 3, we can now set up and solve the value function of a representative buyer n (sometimes I call him just "the buyer" for briefness). This representative buyer faces a standard optimal control problem of choosing how much generic good to produce and crypto to hold as to maximize his (expected) lifetime utility. His state variable is z_t and his control variable is x_t. I allow for both flow and lumpy production of generic good in the CM. Taking this into account and letting T_{trade} denote the stopping time of his next DM meeting, the buyer's value function $V_t^n(z)$ evaluated at $t = t_0$, $z_0 = 0$ reads

$$V_{t_0}^n(0) \equiv V_{t_0}^n = \max_z \mathbb{E}\left\{e^{-rT_{\mathsf{trade}}}\left[u\left(\beta z_{t_0+T_{\mathsf{trade}}}(1-\tau)\right) + V_{t_0+T_{\mathsf{trade}}}^n\right] - \int_0^{T_{\mathsf{trade}}} e^{-rt}\,\mathrm{d}x_t\right\} \quad (2)$$

$$\text{subject to } \mathrm{d}z_t = \mathrm{d}x_t - z_t\pi\,\mathrm{d}B_t \quad (3)$$

The positive term sums the buyer's utility from consuming the special good right after the DM meeting and the continuation value $V_{t_0+T_{\mathsf{trade}}}^n$ after he spent all his crypto. The negative term is the integral of generic good production costs paid to buy crypto in the CM. The change in the buyer's crypto balance over time follows the law of motion in Eq. (3): It increases as he produces generic good to acquire crypto in the CM and reduces owing to inflation as the blockchain expands.

I focus on a stationary solution to Eq. (2) such that the buyer keeps his token portfolio at a constant value $z_t = z$ for all t (except for a zero-measure set of dates). In this way he will always enter the DM with a targeted amount $y \equiv y_{T_{\mathsf{trade}}}$ no matter the random time T_{trade} of the meeting. The procedure that allows the buyer to achieve this stationary portfolio is the following: First, when the buyer is without tokens at t_0, he enters the CM and pays straight away $\mathrm{d}x_{t_0} = z$ to obtain the desired portfolio value. Then, while waiting

to find a seller, he supplies a flow of generic good to the CM that exactly compensates inflation. That is, $\mathrm{d}x_t = z_t \pi \, \mathrm{d}B_t$ keeps $\mathrm{d}z_t = 0$. To actuate this step, the buyer programs a bot (a robot program) that continuously scans the blockchain and places high-frequency trading orders when new blocks appear. Finally, the buyer uses all his coins in the DM meeting and restarts his cycle one instant after. For this plan to make sense, I assume that the buyer can observe the block counter B_t but not miners' mempool. This precludes any contingency on the queue of pending transactions.

By keeping his portfolio stationary, the buyer makes his value function stationary as well, i.e. $V^n_{t_0} = V^n_t = V^n \ \forall t$. To be precise,

Proposition 1. *The representative buyer's value function is given by*

$$V^n = \max_z \quad -z + \frac{1}{r}\left\{\alpha\left[u\left(\beta z\left(1-\tau\right)\right) - z\right] - \mu z \pi\right\} \tag{4}$$

Crypto demand $z(\pi,\tau)$ and special good production $y(z(\pi,\tau))$ solve simultaneously

$$u'\left(y(z(\pi,\tau))\right)\beta(1-\tau) = 1 + \frac{r+\mu\pi}{\alpha} \qquad\qquad y(z(\pi,\tau)) = \beta z(\pi,\tau)(1-\tau) \tag{5}$$

Proof *in Appendix B.*

In Eq. (4), the buyer considers the cryptocurrency as an asset. He first forms his portfolio by paying z upfront. This entitles the buyer to receive a perpetual stream of DM capital gains $u\left(\beta z\left(1-\tau\right)\right) - z$ occurring at rate α but also requires him to bear a perpetual stream of inflation-adjustment costs $z\pi$ at the PoW rate μ. The present value of these constant flows results from simply dividing them by r. Notice that the resulting value function is quasi-linear in z, a standard feature in the LW literature.

Eqs. (4) and (5) determine $z(\pi,\tau)$ by equating marginal revenues to marginal costs of investing in crypto and $y(\pi,\tau)$ by leaving sellers with zero surplus. With some elementary algebra, Eq. (1) readily transforms the equations in (4) into explicit solutions.

Corollary 1. *Crypto demand and special good production are*

$$z(\pi,\tau) = \left(1 + \frac{r+\mu\pi}{\alpha}\right)^{-1} - \frac{1}{\eta\beta(1-\tau)} \qquad y(\pi,\tau) = \beta(1-\tau)\left(1 + \frac{r+\mu\pi}{\alpha}\right)^{-1} - \frac{1}{\eta} \tag{6}$$

To close the analysis of merchants, I provide for completeness the value function of a representative seller s, which is simply the discounted sum of her null DM capital gains.

$$V^s = \frac{\alpha}{r}\left(z(\pi,\tau)\beta(1-\tau) - y\right) \equiv 0 \tag{7}$$

At this point we have everything needed to find the cryptocurrency demand. The analysis would be complete at this early stage were it not for β being still an exogenous obscure object. Characterizing β will take us through a journey within the internal

functioning of the blockchain wherein miners operate. Sections 4 to 5 go deep into the rabbit hole.

4 Miners' Blockchain Updating Game

In this section I construct and solve miners' sub-economy. In contrast to merchants, miners are strategic. Their interactions are best modeled as a stochastic game, which I call the *blockchain updating game*.

Entities $m \in \{1, 2, \ldots, M\}$ become miners after buying a *mining node*, i.e. a dedicated hardware, paying an entry cost F. Each miner m can be thought of as a firm controlling the mining node through a battery of computers rather than as a person running the mining node on a laptop. The brand-new miner connects right away the node to a network of geographically separated machines owned by the other miners, who coordinate their actions according to the PoW protocol.

The M participating miners are profit-oriented and honest; they stick to the prescriptions of the PoW protocol. Since honest miner participation is a defense against security attacks, at any point in time merchants distrust using the cryptocurrency unless $M \geq \underline{M}$ miners are active. (See Appendices A.1 to A.3 for further details on the PoW protocol and its security mechanism).

Miners store and update two digital objects that are of cardinal importance to keep the cryptocurrency alive. The first is a copy of the *blockchain*, a ledger that records all cryptocurrency movements as well as execution data of blockchain programs (smart contracts). Miners build the blockchain through sequential concatenation of time-stamped data structures called blocks. Each block is indexed by a block height $B \in \mathbb{N}_0$ that counts the number of predecessor blocks up to an initial genesis block with height 0. The blockchain height $B_t \in \mathbb{N}_0$ at time t is the height of furthest block away from the genesis.

The second important object that miners store is a *mempool* (memory pool) containing a queue of $Q_t \in \mathbb{N}_0$ pending transactions sent to the miner network by merchants. Miners' task is essentially to select transactions from the mempool and record them on the blockchain. In principle, mempool and blockchain copies can differ across miners, but miners follow the PoW protocol to reconcile their copies and reach consensus on a single version quickly. Mempools can differ only during a negligible amount of time so I will keep track of a single mempool.

The tuple $(Q_t, B_t) \in \mathbb{N}_0^2$ is the state variable of the blockchain updating game. Each round $(Q_t, B_t) = (Q, B)$ initiates a tournament in which miners compete to record a new block and collect the revenues that lie therein. The tournament winner is the first miner that manages to increase the blockchain height.

4.1 Miner actions

In each round (Q, B), the probability that a miner succeeds in extending the blockchain depends on its action $a^m_{(Q,B)}$ and on the actions of the competing miners $\mathbf{a}^{-m}_{(Q,B)}$. Miners can choose $a^m_{(Q,B)} \equiv a^m \in \{\texttt{OFF}, 0, \min(1, Q)\}$.

With $a^m = \texttt{OFF}$, miner m chooses to switch off its machine; with $a^m \neq \texttt{OFF}$, miner m is active and attempts to form a block that **extends** the blockchain at height B (i.e., with height B+1). By choosing $a^m = 0$, the miner works on an *empty block* that records no transactions. On the other hand, by $a^m = 1$, the miner chooses to work on a *full block* with one transaction. The maximum block size available in a given round is $\min(1, Q)$ since block size is constrained by mempool size. So $a^m = 1$ is feasible only if $Q > 0$.

In reality, blocks can include thousands of transactions of heterogeneous size (smart-contract data take much more space than simple token transfers). My restriction to binary block size serves the purpose of obtaining clean expressions.

A profile of actions $\boldsymbol{a}_{(Q,B)} = (a^1, a^2, \ldots, a^M)$ can be stated as a pair of numbers ($m0$,$m1$) indicating how many miners choose $a^m = 0$ and $a^m = 1$. Naturally, $m\texttt{OFF}$=M-$m0$-$m1$ miners play $a^m = \texttt{OFF}$. From the point of view of miner m, the profile $\mathbf{a}^{-m}_{(Q,B)}$ of *other* miners' actions can be expressed via $(m0_o, m1_o)$, $m\texttt{OFF}_o$=M-1-($m0_o$+$m1_o$).

Having well-defined actions, a miner m can form (pure) Markovian strategies $S^m \triangleq \bigcup_{(Q,B)\in\mathbb{N}_0^2} a^m_{(Q,B)}$ that specify $a^m_{(Q,B)}$ in each possible (Q, B) sequence. I will consider exclusively this kind of strategies. The strategy profile $\mathbf{S} = \{S^m\}_{m\in\{1,2,\ldots,M\}}$ fully determines miners' play.

4.2 Miner payoffs

In reward for successfully extending the blockchain, the winner of a miner tournament receives seigniorage from new minted coins (an inflation tax on merchants holding crypto) together with fees if the block records a transaction.

After obtaining these rewards, miners immediately exchange them in the CM to consume their value counterpart in the generic good. Participating to a miner tournament is however not a free lunch. Miners suffer from energy costs per unit of time during tournaments and have no guarantee of success. Indeed, a tournament ends with either a single winner or with no winner at all. This last case occurs if the mempool state changes before any miner is able to form a block.

Getting more formal, I let $c(a)$ denote the energy cost a miner pays as a function of its action: $c(a) = c > 0$ for $a \neq \texttt{OFF}$; $c(\texttt{OFF}) = 0$. I also let $R_{(Q,B)}(a^m, m^*)$ denote the revenues of miner m in tournament (Q, B) given winner $m^* \in \{1, 2, \ldots, M\} \cup \varnothing$. If the tournament has no winner or miner m is $\texttt{OFF}$, $R_{(Q,B)}(a^m, \varnothing) = R_{(Q,B)}(\texttt{OFF}, \cdot) = 0$.

Conversely, if miner m is actively working on a block of size $a \in \{0, 1\}$, it gets

$$R_{(Q,B)}(a, m') = 0 \text{ for } m' \neq m \qquad R_{(Q,B)}(a, m) = z(\pi, \tau)(B\pi + \tau a) \tag{8}$$

where $z(\pi, \tau)$ is the equilibrium crypto balance of the representative buyer. $R_{(Q,B)}(a, m)$ follows from the fact that all buyers are identical and the seigniorage rate is constant.

4.3 Miners' value function and Markov Perfect Equilibria

The equilibrium concept for the blockchain updating game is Markov Perfect Equilibrium (MPE): a profile of Markov strategies that makes each miner maximize profits $\Pi^m_{t_0}$ net of the entry cost,

$$V^m_{t_0} = \Pi^m_{t_0} - F \tag{9}$$

Mining profits obey the following recursion where T is the waiting time of a switch from the current state (Q_t, B_t) to the next state $(Q', B') = (Q_{t+T}, B_{t+T})$,

$$\Pi^m_t = \max_{S^m} \mathbb{E}\left[e^{-rT}\left(R_{(Q_t,B_t)}(a, m^*) + \Pi_{t+T}\right) - \int_0^T c(a)e^{-rt}\,\mathrm{d}t \,\middle|\, \left(a, \mathbf{a}^{-m}\right) \in S^m \times \mathbf{S}^{-m} \right] \tag{10}$$

The state changes when new transactions are *added* to the mempool. This takes place at the aggregate DM rate αN. It also changes when the blockchain grows by one block. As anticipated in Section 3, (in non-trivial equilibria) the blockchain growth follows roughly a Poisson process with rate μ.

Whenever a miner adds a block on top of the blockchain it is at its discretion whether to record a pending transaction or not. Hence the rate at which transactions are *removed* from the mempool depends on miners' actions. In any case, the state changes when either a new block or transaction arrives. This suffices to establish that T is an exponential random variable (ERV) with rate $\alpha N + \mu$.

The Marovian game structure can be exploited to break a miner's optimal dynamic program into a sequence of static problems, one for each round. The next lemma states this formally letting $P\left(a; \mathbf{a}^{-m}\right) \equiv \mathbb{P}_{a,\mathbf{a}^{-m}}(m^* = m \mid m^* \neq \varnothing)$ denote the probability that miner m wins in the current tournament conditional on the tournament having a winner.

Lemma 1. *The miner profit function at time t reads*

$$\Pi^m_t = \frac{1}{r + \alpha N + \mu}\left(\max_{a \in \{\mathsf{OFF}, 0, \min(1, Q_t)\}} v_{(Q_t,B_t)}\left(a, \boldsymbol{a}^{-m}\right) \right) + \mathbb{E}\left[e^{-rT}\Pi_{t+T}\right] \tag{11}$$

$$v_{(Q_t,B_t)}\left(a; \boldsymbol{a}^{-m}\right) \triangleq \mu R(a, m) P\left(a, \boldsymbol{a}^{-m}\right) - c(a) \tag{12}$$

Proof *in Appendix B.*

Assuming that miners hold passive beliefs on the distribution of Q_t (i.e., they take it as given), each miner m simply maximizes the expected profit rate $v_{(Q_t,B_t)}$ of each miner tournament independently. In a given tournament, the miner pays $c(a)$ in energy per unit of time and earns revenues $R(a,m)$ in Eq. (8) at the block creation rate μ weighted by the win probability. Notice that the tournament payoff enters the value function multiplied by the expected time-discounting for the round duration, $(r+\alpha N+\mu)^{-1}$.

As a consequence of Lemma 1, an action profile $\mathbf{a}^m(Q,B)$ is a substrategy of the full Markovian equilibrium if and only if it is a pure strategy Nash equilibrium (NE) of the tournament (Q,B). That is, for each $a^m_{(Q,B)}$ in $\mathbf{a}_{(Q,B)}$,

$$a^m_{(Q,B)} \in \arg\max_a \; v_{(Q,B)}\left(a;\mathbf{a}^{-m}\right)$$

4.4 Winner election in miner tournaments

The winning miner of a tournament is the first in finding a PoW *and* transmitting its block to the other miners. In symbols, the block of the winning miner has the shortest update time

$$T^m_{\mathsf{update}}(a) = T^m_{\mathsf{PoW}} + \Delta^m(a) \tag{13}$$

where a is the miner's action. The winning probability is thus

$$\mathbb{P}(m=m^*) = \mathbb{P}\left(T^m_{\mathsf{update}}(a) < \min_{m'\neq m} T^{m'}_{\mathsf{update}}\left(a^{m'}\right)\right) \tag{14}$$

The foundation of this winner election the Longest Chain Rule, which I present succinctly in Appendix A.2.

Win probabilities depend on the behavior of T^m_{PoW} and $\Delta^m(a)$. Being the PoW rate μ on aggregate, homogeneous miners form blocks at a rate μ/M; each miner creates a negative externality on the others when joining the network by increasing cryptopuzzles' difficulty. T^m_{PoW} is therefore an ERV at with that rate. In contrast, I make the following distribution assumptions on $\Delta^m(a)$:

Assumption 1. *$\Delta^m(\mathit{OFF})=+\infty$, $\Delta^m(0)=0$, $\Delta^m(1)\equiv\Delta$ is an ERV with rate $\theta>\mu$.*

Having more transmissions than mined blocks on average, i.e. $\theta>\mu$, is a necessary requirement for blockchain security as I explain in Appendix A.3.

From Assumption 1 and T^m_{PoW}, it follows that $T^m_{\mathsf{update}}(0)$ is an ERV with rate μ/M, while $T^m_{\mathsf{update}}(1)$ is hypo-exponential with parameters $(\mu/M,\theta)$.[8] Putting all together, the density and distribution function $f_a(t)$, $F_a(t)$ of $T^m_{\mathsf{update}}(a)\equiv t$ for size $a\in\{0,1\}$ are

[8]The hazard rate of an hypo-exponential distribution decreases over time, causing the memoryless property to get lost. With a broader strategy space, this allows miners to benefit from changing action over time *within* a mining tournament, for example by working on full blocks initially and then switching to empty blocks. The Markovian restriction precludes these more sophisticated strategies, making miner actions change *across* tournaments only.

$$
\begin{aligned}
f_0(t) &= \frac{\mu}{M} e^{-t\frac{\mu}{M}} & F_0(t) &= 1 - e^{-t\frac{\mu}{M}} \\
f_1(t) &= \theta \frac{\mu}{M} \frac{1}{\theta - \frac{\mu}{M}} \left(e^{-t\frac{\mu}{M}} - e^{-t\theta} \right) & F_1(t) &= 1 - \frac{1}{\theta - \frac{\mu}{M}} \left(\frac{\mu}{M} e^{-t\theta} - \theta e^{-t\frac{\mu}{M}} \right)
\end{aligned}
\tag{15}
$$

The probabilistic machinery in Eq. (15) allows a miner to calculate the win probability depending on its action.

Before going ahead computing these probabilities, a caveat is in order. As long as delays are sufficiently small relative to block times, their effect on time-discounting can be safely disregarded. Notwithstanding, transmission delays do affect miner tournaments' outcomes since winners are selected based on a lexicographic criterion. In this context, a tiny delay in transmission can cause a miner to lose fully its prize, which is usually substantial. This is why I isolate the effect of transmission delays on tournaments from the rest of the model (except Appendix A.3) by making the following approximation.

Assumption 2. $\min_{m\in\{1,2,\dots,M\}} T^m_{\mathit{update}}(a^m) \approx \min_{m\in\{1,2,\dots,M\}} T^m_{\mathit{PoW}}$

In this way, $T_{\mathsf{block}} = \min_{m\in\{1,2,\dots,M\}} T^m_{\mathsf{update}}(a^m) \approx T_{\mathsf{PoW}} \equiv \min_{m\in\{1,2,\dots,M\}} T^m_{\mathsf{PoW}}$. It is technically possible to develop the rest of the analysis dropping Assumption 2. Yet, keeping it avoids to substantially complicate the other parts of the model without any added important economic intuition.[9]

Win probability

The miner tournament for round (Q, B) has a winner if $T_{\mathsf{PoW}} < T_{\mathsf{trade}}$. Otherwise, the round ends without any blockchain progress. From Section 3, we know that T_{trade} is an ERV with parameter αN, so the tournament has a winner with probability $\frac{\mu}{\alpha N + \mu}$ (see Lemma B.1). Notice that this probability is unaffected by miners' actions.

The probability $P(a^m, \mathbf{a}^{-m}) \equiv P(m0_o, m1_o | a^m)$ that miner m wins conditional on the the tournament having a winner is zero if $a^m = \mathsf{OFF}$. For all other a^m it can be derived from formulae (15). By independence of update times, we have that $\mathbb{P}\left(\min_{m' \neq m} T^{m'}_{\mathsf{update}} > t \right) = (1 - F_0(t))^{m0_o} (1 - F_1(t))^{m1_o}$. Hence, integrating-out $t \equiv T^m_{\mathsf{update}}(a)$ using its density $f_a(t)$ leads to the final formula

$$
P(m0_o, m1_o | a) = \int_0^{+\infty} (1 - F_0(t))^{m0_o} (1 - F_1(t))^{m1_o} f_a(t)\, \mathrm{d}t \qquad a \in \{0, 1\} \tag{16}
$$

The next proposition summarizes its most important properties.

[9] The alternative approach would be to take explicitly into account delays for value functions and the mempool's steady state. The blockchain growth would then follow a two-states regime, with growth rates μ_0 if $Q = 0$; μ_1 if $Q > 0$. $\mu_0 = \mu$ since blocks must be empty, but $\mu_1 \leq \mu$. In the focal case where miners form full blocks for $Q > 0$, a variant of Ren (2019) with stochastic delays applies so that each new block is not forked-out (i.e., tailgated) with probability $\mathbb{P}(T_{\mathsf{PoW}} > \Delta) = (1 + \mu/\theta)^{-1}$. Hence $\mu_1 = \mu(1 + \mu/\theta)^{-1}$.

Proposition 2. *For a non-participating miner,* $P(\cdot,\cdot|\mathsf{OFF}) = 0$. *Conversely, for a participating miner with action* $a \in \{0,1\}$,

$$
\begin{array}{ll}
i & P(m0_o, m1_o|0) \geq P(m0_o, m1_o|1) \\
ii & P(m0_o, m1_o|a) < P(m0_o - m, m1_o + m|a) \quad 0 < m \leq m0_o \\
iii & P(m0_o, m1_o|a) < P(m0_o + m, m1_o + m'|a) \quad m, m' > 0 \\
iv & P(M-1, 0|0) = P(0, M-1|1) = 1/M
\end{array}
$$

The **Proof** is relegated to Appendix B. Hereafter I present the main logic. Proposition 2.i is a sanity check that is a direct consequence of the first-order stochastic dominance relation $f_1(t) \succeq_{\text{FOSD}} f_0(t)$. The successive points (ii) and (iii) state that a miner's winning probability increases if other miners playing $a = 0$ switch to $a = 1$ but decreases when more miners join the tournament. This is clear since empty blocks create a negative externality on miners working on full blocks and increased participation makes tournaments more competitive. The last property in point (iv) reveals that miners have equal probability of winning when they all choose the same action.

Asymmetric action profiles can result in cumbersome probability expressions for a large number of active miners but remain otherwise clean. Lemma 2 computes $P(m0_o, m1_o|a)$ for $M \leq 3$ under asymmetric action profiles: $\{(m0, m1) : m0 \neq m1, m0+m1 \leq 3\}$. The Proof of Lemma 2 contains general closed-form formulae for any $(m0_o,m1_o)$ pair.

Lemma 2. *The closed-form solutions to Eq. (16) for* $\{(m0, m1) : m0 \neq m1, m0 + m1 \leq 3\}$ *are*

$$
\begin{aligned}
&P(0,0|a) = 1 \text{ for } a = 0,1 \\
&P(1,0|1) = \frac{\theta}{2\theta+\mu} && P(0,1|0) = 1 - \frac{\theta}{2\theta+\mu} \\
&P(2,0|1) = \frac{\theta}{3\theta+2\mu} && P(0,2|0) = 1 - \frac{6\theta(2\theta+\mu)}{(6\theta+\mu)(3\theta+2\mu)} \\
&P(1,1|1) = \frac{3\theta(2\theta+\mu)}{(6\theta+\mu)(3\theta+2\mu)} && P(1,1|0) = \frac{\theta+\mu}{3\theta+2\mu}
\end{aligned}
\tag{17}
$$

Proof *in Appendix B.* More subtle but equally important are the properties of the *likelihood ratio*

$$
L(m0_o, m1_o) \triangleq \frac{P(m0_o, m1_o|0)}{P(m0_o, m1_o|1)} \tag{18}
$$

$L(m0_o, m1_o)$ matters since the relative probability advantage of empty over full blocks drives miners' block size choice.

The likelihood ratio will soon enter into action (Section 4.5). For now, I just note two key properties.

Proposition 3. *The likelihood ratio* $L(m0_o, m1_o)$ *satisfies the following properties:*

i $L(m0_o - m, m1_o + m) \geq L(m0_o, m1_o) \quad 0 < m \leq m0_o$

ii $L(M-1,0) = 1 + \frac{\mu}{\theta}\left(1 - \frac{1}{M}\right)$ *is increasing in* M *and* μ/θ.

Proof *in Appendix B.*

Property (i) says that, fixed total miner participation, the relative disadvantage in win probabilities of working on a full block reduces with the number of other miners that do the same. That is, the likelihood ratio exhibits *strategic complementarity* in $a = 1$. Property (ii) instead solves the likelihood ratio in the most adverse situation for a miner producing a full block, i.e. when everyone else mines empty blocks. In this case, the disadvantage of a full block is exacerbated by larger miner participation while faster transmission relative to bock times (i.e., a lower μ/θ) bring the win probabilities associated with full and empty blocks closer.

4.5 Nash equilibrium activity and block size

I now characterize tournaments' NE. Taking $v_{(Q,B)}(\cdot)$ from Eq. (12), miner incentives are fully determined by the payoff inequality $v_{(Q,B)}\left(a; \mathbf{a}^{-m}\right) - v_{(Q,B)}\left(a'; \mathbf{a}^{-m}\right)$ evaluated at different $a, a' \in \{\texttt{OFF}, 0, 1\}$. The payoff inequality gives rise to four incentive-compatibility constraints $\text{IC}^{a,a'}$ for equilibrium action a and deviation a' that characterize Nash equilibria in non-trivial tournaments. Concretely, these are

$$1 + \frac{\tau}{N\pi} \geq L(m0, m1 - 1) \qquad (\text{IC}^{1,0})$$

$$1 + \frac{\tau}{N\pi} \leq L(m0 - 1, m1) \qquad (\text{IC}^{0,1})$$

$$\mu z(\pi, \tau)(N\pi + \tau) P(m0, m1 - 1|1) \geq c \qquad (\text{IC}^{1,\text{OFF}})$$

$$\mu z(\pi, \tau) N\pi P(m0 - 1, m1|0) \geq c \qquad (\text{IC}^{0,\text{OFF}})$$

I now study what conditions refine the equilibrium set in such a way to have only desirable equilibria for cryptocurrency design. To start with, in desirable equilibria all participating miners active, even when $Q = 0$ preventing $a = 1$. Thanks to this property Section 6 will configure the PoW protocol so to protect the blockchain with minimum defense spending; that is, attracting exactly $\underline{M} = M$ miners. To impose full activity it suffices to preclude miners from deviating to OFF in the equilibrium $(m0, m1) = (M, 0)$ because $\text{IC}^{0,\text{OFF}}$ implies $\text{IC}^{1,\text{OFF}}$. Symmetry on the equilibrium path makes things easier since the win probability of every miner is $1/M$ (Proposition 2.ii). It follows that all the M miners stay active if the *Activity Constraint*

$$z(\pi, \tau) N\mu\pi \geq cM \qquad (\text{AC})$$

is satisfied. Notice how AC requires seigniorage being positive and sufficiently high, being

seigniorage the only source of revenues when the mempool is empty.

Proceeding to the analysis of block size, I first solve all possible equilibria and then remove the undesirable ones. Here Proposition 3 plays a major role since the likelihood ratio drives incentives. Asymmetric action profiles cannot be NE since they require $IC^{1,0}$ and $IC^{0,1}$ to hold simultaneously, which is impossible under Proposition 3.i. So only profiles $(m0, m1) \in \{(M, 0), (0, M)\}$ are valid NE candidates. Whether any of these two profiles can actually be sustained as an equilibrium depends on the ratio of fees to seigniorage. An equilibrium with $(0, M)$ is possible only if $\frac{\tau}{N\pi} \geq L(0, M-1) - 1$; that is, transaction fees have to be positive and abundant in relation to seigniorage so to convince miners in taking the invalidation risk rather than gaining a transmission advantage with an empty block—here transmission delays create a marginal cost of recording transactions. On the other hand, an equilibrium where $(0, M)$ is played requires $\frac{\tau}{N\pi} \leq L(M-1, 0) - 1$ due to the converse logic. In the intermediate range $L(0, M-1) - 1 \leq \frac{\tau}{N\pi} \leq L(M-1, 0) - 1$, both symmetric profiles are NE.

Proposition 4. *Under AC, no equilibria involve $a =$ OFF. In equilibria with full activity, all miners choose the same block size determined by the ratio of fees to seigniorage.*

- *I* *For $\tau/(\pi N) < L(0, M-1) - 1$, all miners working on empty blocks is the unique equilibrium.*
- *II* *For $\tau/(\pi N) > L(M-1, 0) - 1$ all miners working on full blocks is the unique equilibrium.*
- *III* *For $L(0, M-1) - 1 \leq \tau/(\pi N) \leq L(0, M-1) - 1$, the game admits both equilibrium I and II.*

Now, equilibria of the type $(m0, m1) = (M, 0)$ are catastrophic for merchants since miners do not DM payments making crypto useless. The most robust design choice to avoid them is to set fees high enough so that $(M, 0)$ cannot be a stable equilibrium prediction. Using the closed-form expression for $L(M-1, 0)$ in Proposition 3.ii, the *Recording Constraint*

$$\frac{\tau}{N\pi} \geq \frac{\mu}{\theta}\left(1 - \frac{1}{M}\right) \tag{RC}$$

ensures that miners fulfill their basic record-keeping task.[10] The comparative statics Proposition 3.ii translate directly to the required safety markup of fees over seigniorage; it increases in M and μ/θ.

From the analysis of tournaments' NE, we can conclude that a PoW blockchain that supports trade requires both $\pi > 0$ and $\tau > 0$ if energy costs and transaction delays are not null.

[10] When RC binds, the mining tournament has multiple equilibria. However, in this cutting-edge case, equilibrium $(M, 0)$ is not stable since it breaks when some miner has arbitrarily small beliefs that other miners work on full blocks.

Proposition 5. *AC and RC require non-zero seigniorage and transaction fees.*

I will next use the parameter restrictions identified here to construct an ideal Markov equilibrium for the whole blockchain updating game.

4.6 Maximum-Mining Markov Perfect Equilibrium

The MPE of the blockchain updating game strategy profiles that induce NE in each tournament originating from the stochastic progression of states.

Under PC and RC, the blockchain updating game admits a unique, ideal *Maximum-Mining Markov Perfect Equilibrium* (MM-MPE) such that all miners remain active in every tournament and form blocks with the maximum available size. In brief, $m\mathsf{off} = 0$, $a^m(Q, 0) = 0$, $a^m(Q, B) = 1$ for $Q > 0$.

Lemma 3. *If conditions PC and RC hold, the Maximum Mining MPE $a^m(Q, B) = 1$ for $Q > 0$; $a^m(0, B) = 0$ for all m is the unique MPE of the blockchain updating game.*

Proof. Straightforward application of the one-shot deviation principle. ■

With $M = \underline{M}$, the MM-MPE makes the blockchain safe with minimum costs while processing merchants' transaction at maximum speed.

5 Mempool Dynamics and Miner Entry

Under the MM-MPE, the progression of states (Q_t, B_t) evolves according to a bivariate continuous Markov chain. Assumption 2 approximates the process of blocks B_t with a simple Poisson process with rate μ. The mempool Q_t follows instead a birth-and-death process with unit increments occurring at Poisson rate αN and unit reductions occurring at Poisson rate μ at dates where $Q_t > 0$.

5.1 Mempool dynamics

The probability distribution of the mempool $g_t(Q)$ is fully parametrized by the *load* $\rho \triangleq \frac{\alpha N}{\mu}$; the ratio of the rates at which transactions enter in and exit from the mempool. If $\alpha N > \mu$ the mempool grows unboundedly, so I rule out this case assuming that the block rate is higher than the transaction rate.

Assumption 3. $\mu > \alpha N$

Apart from the case ruled out by Assumption 3, the mempool has a unique geometric stationary distribution with mass function

$$g_t(Q) \equiv g(Q) = (1 - \rho)\,\rho^Q \qquad \text{with } \rho = \frac{\alpha N}{\mu} \in [0, 1) \tag{19}$$

The probabilities $g(Q), Q \in \mathbb{N}_0$ correspond (in the long-run) to the portion of time in which Q transactions are pending for confirmation. Therefore, ρ is the faction of time in which the mempool is non-empty.

5.2 Confirmation discount factor

In Section 3, the terms of trade set in DM meetings depend on β, the expected discounting over the confirmation period. We can now compute β for a given criterion that miners use for picking transactions out of the mempool.

I will assume that miners apply the random order of service (ROS) criterion, which prescribes them to select transactions uniformly at random. In other words, the probability that a pending transaction is included in the next full block recorded at time t is simply $1/Q_t$.[11] ROS induces a simple distribution for the blockchain growth during the confirmation period, i.e. the period starting at the time the transaction is sent to the mempool and ending the instant before miners insert the transaction into a block. Precisely,

Lemma 4. *The blockchain growth $b \in \mathbb{N}_0$ during the confirmation period follows a geometric distribution with mass function*

$$d(b) = \nu\,(1-\nu)^b \quad \textit{where } \nu \triangleq \frac{1-\rho}{\rho}\ln(1-\rho)^{-1} \tag{20}$$

is the probability that a pending transaction is included in the next block.

Appendix B contains the **Proof** with detailed derivations. The resulting ν is decreasing-concave in ρ, with $\nu = 1$ for $\rho = 0$, and $\nu \to 0$ as $\rho \to 1$.

Now we can use Eq. (20) to calculate the confirmation discount factor β. To keep the model clean, I assume that sellers have imperfect recall and discount each pending transactions independently; they do not suffer mempool congestion from their own transactions. As long as DM meetings are rare enough, this assumption has unimportant equilibrium implications.

Also, in practice, merchants do not consider a transaction confirmed as soon as miners reach consensus on the block containing it. They rather do so when that block reached a confirmation depth $k \in \mathbb{N}_0$ in the blockchain, i.e. is buried by at least k successor blocks. The longer the confirmation depth, the safer merchants can be that their transactions becomes irreversible after the extended confirmation period (see Appendix A.3 for a parenthesis on the logic behind the confirmation depth). Given that security is taken as a black box in this paper, I will normalize $k \equiv 0$ so that transactions are confirmed as

[11] Quoting Easley et al. (2019): "when a miner builds a block he selects from the mempool at random instead of taking the transaction in the pool that has been waiting the longest as in a standard first-in, first-out queue."

soon as they appear on the blockchain. This choice trades-off realism for the sake of a tidy exposition of my main arguments.

Getting back to the calculation of β, notice that a transaction recorded in the next block is time-discounted in expectation by $\mathbb{E}[e^{-rT_{\text{block}}}] = \frac{\mu}{r+\mu}$ (see Lemma B.1.I). Inflation also reduces the value of the pending transaction, but the effect of inflation here is so slight that for realism I will simply ignore it.[12]

Since a pending transaction is discounted incrementally for each block recorded during the confirmation period,

$$\beta \triangleq \mathbb{E}\left[\left(\frac{\mu}{r+\mu}\right)^{b+1}\right] = \sum_{n=0}^{\infty} d(b)\left(\frac{\mu}{r+\mu}\right)^{b+1} = \frac{\mu\nu}{r+\mu\nu} \tag{21}$$

We can see from Eq. (21) that the confirmation period follows an ERV with rate parameter $\mu\nu$.[13] This is natural given the stationary mempool and memoryless block times: Confirmations occur at the block rate μ scaled by the selection probability ν.

Given that violations of safety or incentive-compatibility completely undermine the value of crypto,

Lemma 5. *If either AC or RC is violated,* $\beta = 0$. *Otherwise,* $\beta = \frac{\mu\nu}{r+\mu\nu}$

5.3 Miner Entry

The stationary mempool distribution can be exploited to solve explicitly miners' value function (9). We can again consider a representative miner since all behave identically.

Lemma 6. *The value function of a representative miner under the Maximum-Mining MPE is*

$$W^m = \frac{1}{r}\left[\mu\left(\frac{z(\pi,\tau)\,(\pi N + \tau\rho)}{M}\right) - c\right] - F \tag{22}$$

Appendix B contains the **Proof** through full derivation.

The interpretation of Eq. (22) is straightforward. After paying F to enter the economy, miners incur energy costs at rate c and receive revenues from blocks. For each block, they gain in expectation $\frac{z(\pi,\tau)(\pi N+\tau\rho)}{M}$ as miners win tournaments with equal probability and earn the transaction fees only if the mempool is non-empty (with probability ρ). Each miner collects these revenues at the frequency μ at which the blockchain grows.

[12] It is unreasonable to think that the inflation paid over the minutes, hours at most, of the confirmation period can have any importance on the value of the traded coins. What instead does affect demand in a meaningful way is the inflation that buyers pay while looking for a seller. The length of the search time is on a much higher order of magnitude.

[13] Taking k into account, Eq. (21) becomes $\frac{\mu\nu}{r+\mu\nu}\left(\frac{\mu}{r+\mu}\right)^k$

Free-entry participation leaves no rents to a potential miner entrant. So miner entry is the highest integer M compatible with $W^m \geq 0$ at M and $W^m < 0$ at $M+1$. Using $\rho = \alpha N/\mu$, free-entry implies $M = \lfloor \tilde{M} \rfloor$, where $\tilde{M} = \dfrac{Nz(\pi,\tau)(\mu\pi + \alpha\tau)}{c + rF}$ is the ratio of aggregate miner revenues to costs such that $W^m = 0$ at $M = \tilde{M}$.[14] Since secure participation $\underline{M}$ is an integer, the blockchain is secured with $\tilde{M} = M = \underline{M}$.

The optimal blockchain design will not elicit miner participation beyond $\underline{M}$ since that would only create a dead-weight loss as will become clear in Section 6. Therefore,

Proposition 6. *Secure miner participation under free entry is guaranteed by the Participation Constraint*

$$\underline{M} = M = \frac{Nz(\pi,\tau)(\mu\pi + \alpha\tau)}{c + rF} \tag{PC}$$

Interestingly, after endogenizing the mempool, PC features N in front of both the seigniorage and the fee revenue component, thereby affecting only total revenues but not their relative composition. This happens because increasing the buyer population makes it less likely for miners to find the mempool empty. As a result, we can normalize N to 1 in the design section since its effect on the optimal mix of fees and seigniorage is equivalent to the one of shifting $\underline{M}$.

6 Cryptocurrency Design

This section takes a normative standpoint and introduces the problem of a welfare-maximizing social planner (SP) that taxes merchants with (π, τ) so that miner entry and strategies $(M, \mathbf{S})$ comply with constraints (AC, PC, RC). The SP can be thought of as a community of software developers proposing a hard fork of the cryptocurrency source code. It can also be considered as a decentralized autonomous organization (DAO).

The planner's objective is to maximize welfare W, the aggregate utility in the economy. The SP (she) can decide to either install the cryptocurrency or not to do it. In the latter case, $W = 0$. In the former more interesting case, all agents in the economy can benefit from miner participation reaching its secure level given that the ability to pay with crypto is a public good. However, running the PoW blockchain also creates a welfare loss caused by mining sunk costs (in particular energy consumption).

After normalizing N to 1 and setting directly $M = \underline{M}$ in the planner's constraints—I will prove shortly that is optimal—her cryptocurrency design problem can be framed as

$$W = \max_{\pi,\tau} \quad V^n + V^s + MV^m \qquad \text{subject to} \tag{23}$$

[14] The floor function $\lfloor \tilde{M} \rfloor$ outputs the highest integer $M \in \mathbb{N}_0$ that satisfies $M \leq \tilde{M} < M + 1$

$$W \geq 0$$
$$\underline{M}c \leq z(\pi,\tau)\mu\pi \tag{AC$'$}$$
$$\underline{M}(c+rF) = z(\pi,\tau)\left(\mu\pi + \alpha\tau\right) \tag{PC$'$}$$
$$\tau \geq \pi\frac{\mu}{\theta}\left(1-\underline{M}^{-1}\right) \tag{RC$'$}$$

Other constraints enter tacitly in the planner's problem. These are Eq. (5) for equilibrium demand $z(\pi,\tau)$; Eq. (21) for confirmation discounting β; Eqs. (19) and (20) for distributions' parameters ρ and ν.

It is handy to express W in terms of consumption and production of real goods rather than crypto balances because money exchanges are just lump-sum transfers across agents. In particular, merchants transfer money to miners through taxation via transaction fees and seigniorage, whereas buyers transfer money to sellers with their DM payments. Compensating these transfers,

Lemma 7. *Welfare in Eq. (23) corresponds to*

$$W = \max_{\pi,\tau} \ \frac{\alpha}{r}\left(u\left(y(\pi,\tau)\right) - \frac{y(\pi,\tau)}{\beta}\right) - \frac{y(\pi,\tau)}{\beta(1-\tau)} - M\left(\frac{c}{r}+F\right) \tag{24}$$

Proof *in Appendix B.*

In Eq. (24), welfare is the trade surplus generated in the DM minus the costs of liquidity acquisition in the CM and mining. β erodes part of the DM surplus due to the inefficiency of non-immediate confirmations; setting $\beta = 1$, the DM expression reduces to the full surplus $u\left(y(\pi,\tau)\right) - y(\pi,\tau)$.

Lemma 7 immediately provides a clear design clue. As M enters as a direct social cost, the planner attracts only $M = \underline{M}$ miners whenever $W \geq 0$, as required by PC$'$.

Corollary 2. *$M = \underline{M}$ is welfare-optimal.*

Also, notice that τ enters Eq. (24) directly while π does not. The fact that transaction fees distort the bargaining terms in the DM while seigniorage counts as a mere transfer to miners explains this observation. Specifically, τ makes the terms of trade less favorable, forcing buyers to produce more generic good and carry more liquidity to the DM.

A less evident property that I am about to prove is that the distortion of the terms of trade is the *only* detrimental effect of transaction fees on welfare. As a matter of fact, substituting fees for seigniorage only affects the cost of initial liquidity *acquisition* $\mathrm{d}x = z(\pi,\tau)$ that buyers pay when they are empty-handed. This is fully determined by sellers' contractual terms. Doing the substitution has no effect whatsoever on the *holding* cost of liquidity that buyers pay to preserve their portfolio's value over time. This is the same as saying that the change in policy would only affect $V^n(0)$ but not $V^n\left(z(\pi,\tau)\right)$.

The subtlety at play here is mempool's endogeneity, which removes the larger-tax-base advantage that seigniorage would otherwise have as identified by Chiu and Koeppl (2019), in turn eliminating the inefficiency of fees in terms of holding cost.

To see this more carefully, it is helpful to recognize that total welfare also corresponds to aggregate buyers' utility since buyers capture fully the surplus in the economy; i.e., $W{=}\max_{\pi,\tau}\ (V^n{+}V^s{+}\ MV^m){=}\max_{\pi,\tau} V^n$. Consequently, the total liquidity cost in the economy $\Phi(\pi,\tau,y(\pi,\tau)) \triangleq \frac{y(\pi,\tau)}{\beta(1-\tau)}\left(1+\frac{\alpha+\mu\pi}{r}\right)$ is the sum of all negative terms in Eq. (4) (expressed in terms of special good units). At this point, I proceed with the decomposition

$$\begin{aligned} \Phi(\pi,\tau,y) &\equiv \Phi^{\mathsf{a}}(\pi,\tau,y)+\Phi^{\mathsf{h}}(\pi,\tau,y) \\ \Phi^{\mathsf{a}}(\pi,\tau,y)=\frac{y}{\beta(1-\tau)} &\qquad \Phi^{\mathsf{h}}(\pi,\tau,y)=\frac{y}{r\beta}\left(\frac{\alpha+\mu\pi}{1-\tau}\right) \end{aligned} \tag{25}$$

$\Phi^{\mathsf{a}}(\pi,\tau,y)$ is the acquisition cost; $\Phi^{\mathsf{h}}(\pi,\tau,y)$ is the holding cost.

The surprising observation that follows is that, for $y(\pi,\tau)\equiv y$, $\Phi^{\mathsf{h}}(\pi,\tau,y)$ is invariant to the composition of (π,τ) as long as tax rates move along the PC$'$ curve. To see this, consider a change in the planners' policy from (π,τ) to $\tau'<\tau$, $\pi'>\pi$ that keeps miner average revenues $\frac{y}{\beta}\left[\frac{\alpha\tau+\mu\pi}{1-\tau}\right]$ constant. The new policy must be such that

$$\pi'=\pi+\epsilon_\pi \qquad \tau'=\tau-\epsilon_\tau \qquad \mu\epsilon_\pi=\epsilon_\tau\left(\frac{\alpha+\mu\pi}{1-\tau}\right) \tag{26}$$

leading to the change in liquidity cost

$$\begin{aligned} \Phi(\pi',\tau',y)-\Phi(\pi,\tau,y) &= \underbrace{\frac{y}{\beta}\left(\frac{1}{1-\tau'}-\frac{1}{1-\tau}\right)}_{\Phi^{\mathsf{a}}(\pi',\tau',y)-\Phi^{\mathsf{a}}(\pi,\tau,y)}+\underbrace{\frac{y}{r\beta}\left(\frac{\alpha+\mu\pi'}{1-\tau'}-\frac{\alpha+\mu\pi}{1-\tau}\right)}_{\Phi^{\mathsf{h}}(\pi',\tau',y)-\Phi^{\mathsf{h}}(\pi,\tau,y)} \\ &= \frac{y}{\beta}\left(\frac{1}{1-\tau'}-\frac{1}{1-\tau}\right) \\ &= -(z-z')<0; \qquad \text{with } z=y\beta(1-\tau),\ z'=y\beta(1-\tau') \end{aligned} \tag{27}$$

The difference in Φ^{h} before and after the policy change is identically zero by the definition of (π',τ') in Eq. (26). Only $\Phi^{\mathsf{a}}(\pi',\tau',y)$ is affected and lowered by the new policy.

In light of the previous considerations, the solution (π^*,τ^*) to a reduced SP problem that does not take into account neither AC$'$ nor RC$'$ sets $\tau^*=0$.

Proposition 7. *Transaction fees are inefficient relative to seigniorage. That is,* $\tau^*=0$.

The **Proof** (in Appendix B) is along the lines as Chiu and Koeppl (2019)'s proof of their Proposition 8.

With Proposition 7's insight at hand, seeking a solution (π^{**},τ^{**}) additionally constrained by AC$'$ and RC$'$ is just a matter of looking for the lowest τ that is compatible

with the two additional constraint. This is an easy check: Inspecting their expressions, it becomes clear that reducing τ relaxes AC$'$ since crypto demand rises but it does tight RC$'$, which therefore acts as a lower bound. The planner then sets τ as low as RC$'$ allows her to, making it eventually bind at

$$\tau^{**} = \pi^{**} \frac{\mu}{\theta}\left(1 - \frac{1}{\underline{M}}\right) \tag{28}$$

The corresponding seigniorage is (the lowest) π^{**} that solves PC$'$ $|_{\tau=\tau^{**}}$. Precisely, it is the solution to

$$\underline{M}\,(c + rF) = z\left(\pi^{**}, \pi^{**}\frac{\mu}{\theta}\left(1 - \frac{1}{\underline{M}}\right)\right)\mu\pi^{**}\left[1 + \frac{\alpha}{\theta}\left(1 - \frac{1}{\underline{M}}\right)\right] \tag{29}$$

Proposition 8. *In a cryptocurrency design, the planner uses the minimum incentive-compatible transaction fee; that is,* (π^{**}, τ^{**}) *solves jointly Eqs. (28) and (29).*

Fig. 1 illustrates graphically the constraints and solutions of the cryptocurrency design problem in the (π, τ)-space. The blue segment highlights the feasible design region where PC$'$ binds and AC$'$, RC$'$ are satisfied. W reduces as its violet level curves move away from the maximum welfare point $\pi = \tau = 0$ in any positive (π, τ) direction; lighter violet in the figure stands for lower welfare. The orange asterisk identifies the solution to the reduced problem $(\pi^*, \tau^*) = (\pi^*, 0)$, with π^* satisfying PC$'$. Conversely, the tax vector $(\pi_{\text{RP}}, \tau_{\text{RP}}) = (\pi^{**}, \tau^{**})$ corresponding to the more restricted solution makes RC$'$ bind at the lowest feasible τ and highest π. The last graphic element is the diametrically opposite design solution $(\pi_{\text{AP}}, \tau_{\text{AP}})$, which is not optimal but plays a role for cryptocurrency implementability as will see shortly.

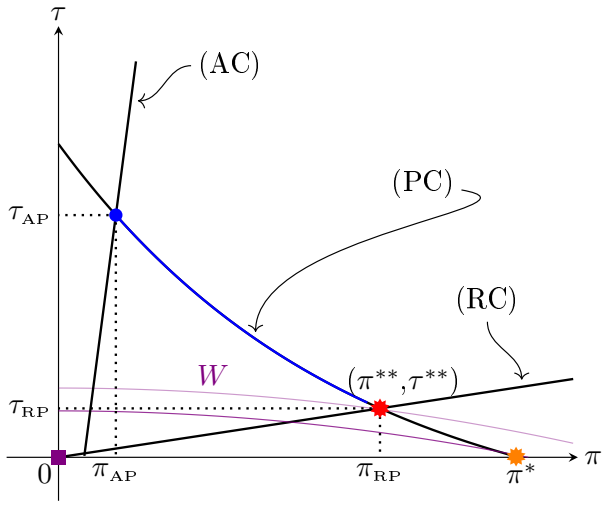

Figure 1: The cryptocurrency design problem

Proposition 8 gives a design receipt for the case in which the planner implements crypto. The only missing pieces of the optimal design puzzle are the conditions that

make implementing the cryptocurrency actually feasible and welfare improving relative to a trivial zero-welfare economy without crypto.

Starting from the latter aspect, the planner avoids the null design if buyers' taste for special-good consumption η justifies the mining costs from a welfare standpoint. A nice feature of the generalized log-utility (1) is that there always exist sufficiently large values of η that make the cryptoeconomy worthwhile. Raising η increases welfare benefits unboundedly while keeping welfare costs bounded. This can be clearly seen from Eq. (24) in combination with Eq. (5): As η grows large, non-mining welfare costs are proportional to $y(\pi,\tau)$, which converges to $\beta(1-\tau)\frac{\alpha}{\alpha+r+\mu\pi} \leq 1$, whereas $u(y(\pi,\tau)) = \ln\left(\eta\beta(1-\tau)\frac{\alpha}{\alpha+r+\mu\pi}\right)$ goes up to infinity.

On the other side of the coin, feasibility depends crucially on the Activity Constraint. AC′ plays no role for the solution identified by Proposition 8 as it remains slack (or binds simultaneously with PC′). However, it can prevent the planner to implement crypto if violated. Subtracting each side of AC′ to the corresponding side of PC′ gives the condition $\alpha z(\pi^{**},\tau^{**})\tau^{**} \leq rF\underline{M}$ that guarantees implementability. Rearranging after using Eq. (28) gives

$$\alpha z\left(\pi^{**},\pi^{**}\,\frac{\mu}{\theta}\left(1-\frac{1}{\underline{M}}\right)\right)\pi^{**}\frac{\mu}{\theta}\left(\frac{1}{\underline{M}}-\frac{1}{\underline{M}^2}\right) \leq rF \tag{30}$$

Eq. (30) counter-intuitively requires miner income from transaction fees to be low relative to the (actualized) entry cost. The source of the problem here are again transmission delays that force the planner to use inefficient transaction fees. Was RC′ absent, the planner would fully cover fixed and flow costs of mining with seigniorage. Yet RC′ forces the planner to collect some income from transaction fees depending on the magnitude of transmission delays to block times. The total tax income from seigniorage and fees combined is fixed, which means increased fees necessarily reduce seigniorage. Since the planner has to at least cover energy costs with seigniorage, fees cannot exceed the mining fixed cost. Therefore, when RC′ requires fees to be even higher, the planner cannot provide anymore sufficient seigniorage and breaks AC′.

It is clear now that keeping μ/θ low solves this problem. Indeed, condition (30) is trivially satisfied for $\theta \to \infty$ so that $\tau^{**} \to \tau^{*} = 0$. Since $z \leq 1$ and $(\frac{1}{\underline{M}} - \frac{1}{\underline{M}^2}) \leq 1/4$, Eq. (30) gives a loose upper bound of $\frac{\mu}{\theta} \leq 4\frac{rF}{\alpha}$. This observation closes the characterization of the social planner's problem.

Proposition 9. *The social planner implements the cryptocurrency only if η is sufficiently high and μ/θ is sufficiently low.*

6.1 Numerical comparative statics

In these last bits of analysis, I show how the optimal design reacts to parameters numerically. Fig. 3 displays the result of my numerical exercise on α, μ and θ.

To understand how the relative weight of τ^{**} and π^{**} changes with the parameters, it is worth to keep in mind RC′ while looking at Fig. 3. In panel (a.i), DM frequency reduces the overall tax burden but does not change the relative composition of taxes on merchants; α is neutral to RC′, which keeps the ratio of fees to seigniorage is constant. The other parameters do affect the relative weight of the two tax instruments. In panel (b.i), μ tightens RC′ by intensifying miner competition. Consequently the planner increases the weight of fees relative to seigniorage. This effect maintains τ^{**} almost constant for the parameters of my simulation. On the contrary, in panel (c.i) θ alleviates the restrictions of (RC′) so the planner can substitute fees with the more efficient seigniorage.

Moving to welfare effects, panels (b.ii) and (c.ii) show that W increases in μ and θ. θ reduces the needed fees boosting welfare. μ instead raises the relative weight of fees but also unclogs the queue of pending transactions and boosts the effectiveness of seigniorage as a tax instrument; in my simulation the positive effects prevail. The effect of α on welfare is non-monotone as evident from (a.ii). Higher DM frequency allows buyers and sellers to trade more often producing more value. But more frequent trade makes the mempool more congested thus reducing welfare when trade frequency becomes excessive.

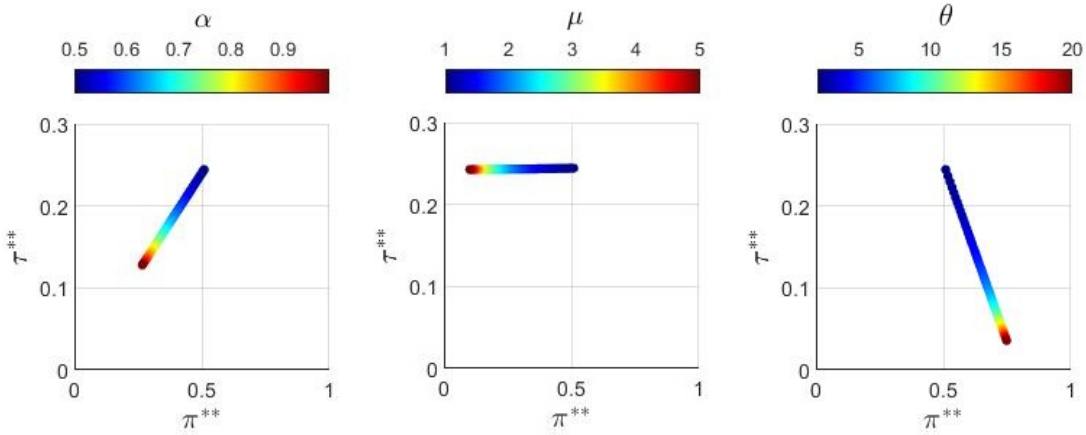

(a.i) (b.i) (c.i)

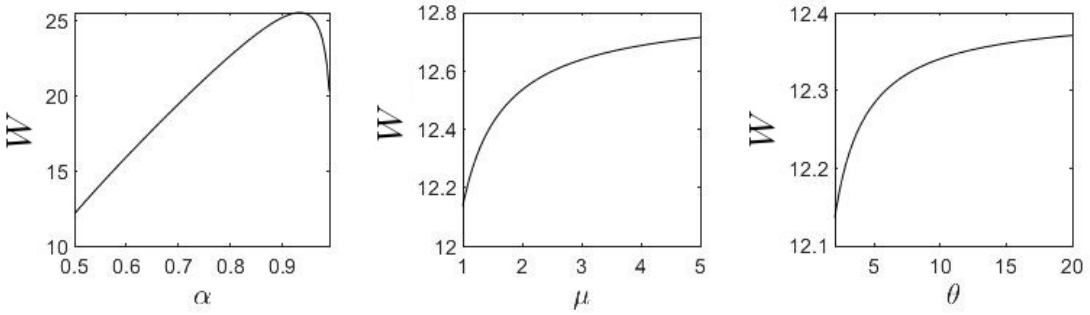

(a.ii) (b.ii) (c.ii)

Figure 3: Welfare maximizers in the (π, τ)-space for $u(y) = \ln(1 + 100y)$, $(N,M)=(1,25)$, $(c,F)=(0.001,0.1)$, $r = 0.1$. $\alpha \in [0.5, 1)$, $\mu=1$, $\theta=2$ in (a.i,ii); $\alpha=0.5$, $\mu \in [1, 5]$, $\theta=2$ in (b.i,ii); $\alpha=0.5$, $\mu=1$, $\theta \in [2, 20]$ in (c.i,ii).

7 Concluding Remarks

I presented a model for the plurality of interactions involved in the economy of a PoW cryptocurrency. Miners play the essential role of record-keepers by storing blocks of transactions on the blockchain. They also contribute to the security of the ecosystem and resolve inconsistencies among their ledger copies. The seigniorage and transaction fees they receive in reward for doing their job introduce fundamental tradeoffs that this paper brought to light.

As a final discussion, I will give some food for thought on possible extensions and the generality of main results.

A natural direct extension of this paper it to cover more elaborate types of PoW than the classic version that I analyzed. One of these variants is Ethereum's Greedy-Heaviest-Observed-Subtree (GHOST). The added feature is that miners earn rewards also in blocks that end up in secondary branches of the blockchain as long as they are not too far away from the consensus chain's endpoint. This variation is meant to cope with the high rate of forks and harmful block invalidation intrinsic to blockchains with a short average block time. Extending my work to cover GHOST would require miners to evaluate tournaments' payoffs on the grounds of a more general blockchain ramification process rather than only based on the branching of a single parent block.

Future research could also address an important alternative motive for miners to avoid forming blocks at ideal capacity: verification costs. In practice, each time a miner receives a block, it cannot mine on top of that block until it verified the block's validity (assessing whether it contains no double-spent transactions and carries a valid proof of work among other checks). The PoW protocol prescribes the miner to put the block in a dead fork if found invalid.

As for transmissions, verification times of bigger blocks are longer, especially when blocks contain smart-contract information. So verification delays are another reason why

miners could give up on record-keeping. The way miners can partially aid the slowdowns of verification times and earn more profits is by carrying out verifications while in parallel extending the blockchain with empty blocks, which cannot clash with the information to be verified. Nevertheless, nothing in principle prevents miners to act in more cheaty ways, for example by simply refusing to extend blocks that require a long verification or by extending them without verifying their content.

The strategic implications of costly verification are surely different from the ones of of costly transmission. The cost of verification is not paid by the miner that creates the block, but rather by those who *receive* it. These other miners do not earn transaction fees in compensation for validating, so incentivization via transaction fees seems not workable. Verification is also hardly possible to sustain by honest miners punishing cheaters with a reputation loss, e.g. with grim-trigger strategies, because a miner's mempool is private information; miners do not know exactly which blocks each of them is supposed to verify. The question of how to sustain a cryptocurrency equilibrium inducing record-keeping in the presence of costly verification is yet to be addressed.

Finally, a big open question is to what extent the insights of this paper generalize to other cryptocurrency archetypes. Most importantly, whether they apply to cryptocurrencies that use Proof-of-Stake blockchains. The basic principle of PoS is that miners (sometimes termed 'validators', 'minters', or 'bakers') are now identified through public addresses and miner tournaments are run by the protocol via plutocratic elections. The protocol execution follows a sequence of pre-determined time steps. In each step, it samples a subset of coins among those held by miners and gives their owners the possibility to propose the next block. The protocol then uses a rule to select a single block among all proposals.

A plethora of PoS variants coexist under this general principle, yet they all fall within two sub-categories: Longest-Chain PoS (LC-PoS), such as Tezos' Emmy+ or Cardano's Ouroboros; Byzantine-Fault-Tolerant PoS (BFT-PoS) like Algorand. For both types of protocols, the elimination of CPU-based tournaments reduces drastically the energy cost. So the importance of the Activity Constraint is not fully eliminated but much reduced. Block invalidation risks are practically absent from BFT-PoS blockchains: The ledger grows on a single chain though a multi-step voting process. On the flip side, they are still very much relevant for LC-PoS, which like PoW admit concurrency among simultaneous chains. In these protocols, forks can originate if more than one miner is elected to propose the next block in the same round. Once again, LCR resolves forks under the principle that the first transmitted block on top of the longest chain (or chains) defines the new state of the ledger. Hence also for these blockchains some variant of the Recording Constraint requiring minimum positive transaction fees is needed.

Looking at the broader picture, PoS blockchains were borne with the prerogative of solving the big environmental cost associated with the PoW mechanism. Removing re-

liance on energy-intensive cryptopuzzles makes PoS run faster than PoW, in turn making it also more suited for elaborate financial applications that require speed.

Despite the PoW protocol's pollution problem and heavy security mechanism, it can still boat some unique virtues among the spectrum of existing blockchain solutions: simplicity, robustness and full anonymity of its participants. On these aspects, PoW outperforms PoS. Indeed, LC-PoS protocols introduce many more design criticalities that a careful planner has to tackle—most importantly, PoS have to maintain miner tournaments unpredictable while limiting Nothing-at-Stake attacks (Brown-Cohen et al., 2019).

All things considered, a thorough understanding of optimal PoS cryptocurrency systems is nowadays becoming increasingly important. My personal view is that, in a near future, modern regulation-friendly PoS blockchains will complement central bank digital currencies (CBDCs) in the emerging internet of value.

References

Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.

Brown-Cohen, J., Narayanan, A., Psomas, A., and Weinberg, S. M. (2019). Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 459–473.

Buterin, V. (2014). Toward a 12-second block time. https://blog.ethereum.org/2014/07/11/toward-a-12-second-block-time/. Accessed: 2022-03-22.

Buterin, V. (2021). Eip 1559 FAQ. https://notes.ethereum.org/@vbuterin/eip-1559-faq#Might-EIP-1559-run-the-risk-of-over-stressing-nodes-and-miners-during-periods-of-high-usage. Accessed: 2022-03-22.

Carlsten, M., Kalodner, H., Weinberg, S. M., and Narayanan, A. (2016). 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.

Chiu, J. and Koeppl, T. (2019). The economics of cryptocurrencies—bitcoin and beyond. Staff Working Papers 19-40, Bank of Canada.

Choi, M. and Rocheteau, G. (2020a). Money mining and price dynamics. *Available at SSRN 3336367*.

Choi, M. and Rocheteau, G. (2020b). New monetarism in continuous time: Methods and applications. *Available at SSRN 3435889*.

Cong, L. W., He, Z., and Li, J. (2019). Decentralized mining in centralized pools. Technical report, National Bureau of Economic Research.

Decker, C. and Wattenhofer, R. (2013). Information propagation in the bitcoin network. In *IEEE P2P 2013 Proceedings*, pages 1–10. IEEE.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*.

Fernández-Villaverde, J. and Sanches, D. (2019). Can currency competition work? *Journal of Monetary Economics*, 106:1–15.

Gaži, P., Kiayias, A., and Russell, A. (2020). Tight consistency bounds for bitcoin. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 819–838.

Houy, N. (2016). The bitcoin mining game. *Ledger*, 1:53–68.

Huberman, G., Leshno, J., and Moallemi, C. C. (2019). An economic analysis of the bitcoin payment system. *Columbia Business School Research Paper*, (17-92).

Kleinrock, L. (1975). Queueing systems: Volume 1: Theory, 1975. *A Wiley-Interscience Publication*.

Lagos, R., Rocheteau, G., and Wright, R. (2014). The art of monetary theory: A new monetarist perspective. *forthcoming, Journal of Economic Literature*.

Lagos, R. and Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of political Economy*, 113(3):463–484.

Lehar, A. and Parlour, C. A. (2020). Miner collusion and the bitcoin protocol. *Available at SSRN 3559894*.

Liu, Y., Lu, Y., Nayak, K., Zhang, F., Zhang, L., and Zhao, Y. (2022). Empirical analysis of eip-1559: Transaction fees, waiting time, and consensus security. *arXiv preprint arXiv:2201.05574*.

Malik, N., Aseri, M., Singh, P. V., and Srinivasan, K. (2022). Why bitcoin will fail to scale? *Management Science*.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Cryptography Mailing list at https://metzdowd.com*.

Neudecker, T. and Hartenstein, H. (2019). Short paper: An empirical analysis of blockchain forks in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 84–92. Springer.

Pagnotta, E. S. (2022). Decentralizing money: Bitcoin prices and blockchain security. *The Review of Financial Studies*, 35(2):866–907.

Pass, R. and Shi, E. (2017). Rethinking large-scale consensus. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 115–129. IEEE.

Prat, J. and Walter, B. (2018). An equilibrium model of the market for bitcoin mining.

Ramsey, F. P. (1927). A contribution to the theory of taxation. *The economic journal*, 37(145):47–61.

Ren, L. (2019). Analysis of nakamoto consensus. *Cryptology ePrint Archive*.

Roberts, S. (2022). How 'trustless' is bitcoin, really? *The Ney York Times*.

Rosenfeld, M. (2014). Analysis of hashrate-based double spending. *arXiv preprint arXiv:1402.2009*.

Sankagiri, S., Gandlur, S., and Hajek, B. (2021). The longest-chain protocol under random delays. *arXiv preprint arXiv:2102.00973*.

Schilling, L. and Uhlig, H. (2019). Some simple bitcoin economics. *Journal of Monetary Economics*, 106:16–26.

Tsabary, I. and Eyal, I. (2018). The gap game. In *Proceedings of the 2018 ACM SIGSAC conference on Computer and Communications Security*, pages 713–728.

Appendix A Primer on Blockchain Technology

This appendix provides an elementary introduction to three fundamental aspects of the classical PoW protocol. Each aspect is treated in a separate but not independent subsection.

A.1 Block mining

To ensure tamper-proofness, PoW blockchains require miners to apply a security mechanism to each block they concatenate. Concretely, a miner is eligible to extend a parent block only after solving a cryptographic puzzle through an energy-consuming process.

The candidate solution works as an identifier, or hash, of the miner's new block. The miner computes candidate solutions by applying a hash function to the combination of (i) the hash of the parent block; (ii) the transactions it wants to record; (iii) a nonce, i.e. a randomly drawn number only used once.[15] Fig. 4 illustrates the resulting hash chain.

[15]The specific hash function is SHA256 in Bitcoin and Keccak256 in Ethereum.

A hash for the new block solves the cryptographic puzzle if it starts with a protocol-specified number of initial zeros. The longer the string of leading zeros, the harder the cryptographic puzzle. Since hash functions' output is practically unpredictable based on inputs, the miner can only attempt solving the cryptopuzzle by generating a nonce per unit of time until the hash validity condition is met. When that happens, the nonce becomes the miner's proof of work for the new block.

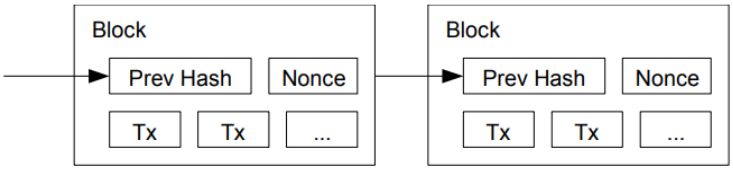

Figure 4: Hash chain (Nakamoto, 2008)

A.2 The Longest Chain Rule

Obtaining a valid PoW and extending the blockchain does not yet guarantee a miner that the update contained in its block is taken into account by the other miners. Due to asynchronous and decentralized communication among miners, it is possible that more than one miner finds valid a PoW for the same predecessor block, in this way proposing two or more distinct chronological successors. These blocks differ at least on the miner that receives the seigniorage payment.

The problem arising when this happens is the *forking* (bifurcation) of the blockchain into in several conflicting chains, each providing a different version of the ledger up to a common point of agreement. Forking turns the blockchain from a linear registry to a directed tree.

The presence of multiple active chains creates ambiguity on the state of the ledger and hinders merchants' trust in the blockchain (e.g., because their cryptocurrency balance becomes ambiguous). To avoid the endurance of multiple chains over time, the PoW protocol ensures that miners eventually reach consensus on a unique chain to follow and extend, allowing for simultaneous active chains only temporarily. For this reason, whenever a miner finds a valid PoW to extend the blockchain, it has to broadcast the new block to the other miners. In case there is no fork, the block is accepted. If instead there is a fork, miners will have to decide which branch to accept.

In PoW blockchains, miners follow the chain selection criterion prescribed by Nakamoto (2008); namely, the Longest Chain Rule. LCR resolves forks by selecting the longest chain (i.e., the chain with most blocks on it) a miner is first aware of. Assuming that block transmissions reach all miners simultaneously when completed, LCR de facto initiates a new tournament among miners at each (Q_t, B_t). The quickest miner in finding a PoW

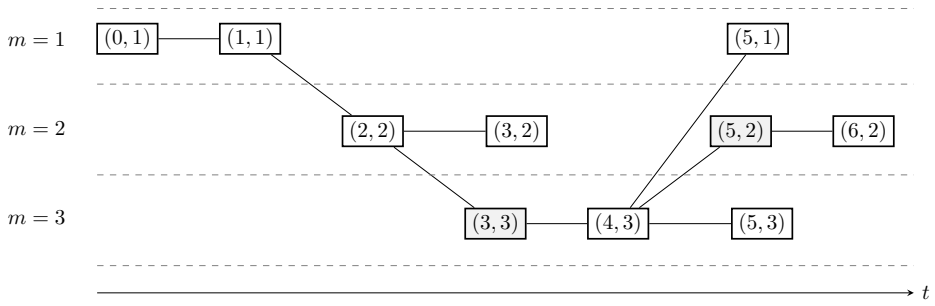

Figure 5: Three miners updating the blockchain

for its new block *and* transmitting the block to the other miners establishes the new canonical longest chain—as described by Section 4.4. In other words, the block of the winning miner has the shortest update time $T^m_{\mathsf{update}}(a)$ defined in Eq. (13).

Fig. 5 illustrates how the blockchain grows over time in my model as a result of mining tournaments played among three miners $m \in \{1, 2, 3\}$, with the t-axis representing the time blocks reach the whole miner network. Blocks are indexed (B, m) with $B \in \{0, 1, 2, 3, 4, 5, 6\}$. In this illustration, forks occurring at $B \in \{3, 5\}$ are resolved by the gray blocks, which are the first in completing transmission among their cohorts.

Remark (Forks): The number of chains originating after a mining tournament with a winner is

$$1 + \left| \left\{ m : T^m_{\mathsf{PoW}} + \Delta^m(a^m) < T^{m^*}_{\mathsf{PoW}} + \Delta^{m^*}(a^{m^*}) < T^m_{\mathsf{PoW}} + \Delta^m(a^m) \right\} \right| \tag{31}$$

The cardinality operator in Eq. (31) counts the number of blocks abandoned in discarded chains. If $T^{m^*}_{\mathsf{PoW}} + \Delta^{m^*}(a^{m^*}) < T^m_{\mathsf{PoW}}$ for all $m \neq m^*$, the blockchain will extend to a single branch, since after receiving the winner's block with a valid PoW for height B, the other miners acknowledge that the blockchain height increased and immediately reset their reference block. Outside this case, more than one miner transmits its blocks in the same round, giving rise to a fork.

A.3 Blockchain security

A malicious miner can attempt to sabotage the blockchain by forking it intentionally attempting to outgrow the longest chain followed by honest miners. The ability to succeed in this attack puts the saboteur in power to arbitrarily modify the ledger. However, as long as honest miners follow LCR, they join their forces to make the honest chain grow faster.

Formally, the honest miners and the attacker engage in a Poisson race where the attacker mines blocks at rate $\mu \frac{A}{M+A}$ while the coalition of honest miners does so at rate

$\mu\frac{M}{M+A}$. The attacker does not suffer from transmission delays since it can be considered as either a single entity or a consortium running a mining farm where block transmission occurs through direct links. Contrariwise, the honest nodes need to communicate updates among each other to account for the growth of their chain; they mine full blocks that require a transmission time Δ defined by Assumption 1 as an ERV with rate θ. Sankagiri et al. (2021) prove the fundamental security condition for the assumptions of my model:[16] Formally, using my notation and letting $\{\mathcal{A} \mid k\}$ denote the event "the attacker grows a longer chain within k blocks",

$$\frac{A}{M+A} < \frac{1}{2}\left(1-\frac{\mu}{\theta}\right) \implies \mathbb{P}\left(\mathcal{A} \mid k\right) \leq e^{-\Omega(k)} \tag{32}$$

The asymptotic lower bound symbol $\Omega(k)$ indicates that there exists a constant, say $h \in \mathbb{R}_+$, such that $kh \leq \Omega(k)$ for k sufficiently large.

In words, Eq. (32) says that the probability that the attack succeeds in a window of k blocks falls exponentially in k as long as honest miners hold a super-majority of the computing power.

Notice that the fraction of malicious CPU in Eq. (32) must be below 1/2 by a factor that depends on the ratio of block rate to transmission rate. So as in the mining game of the main analysis, blockchains that produce blocks too fast relative to the block transmission time are problematic. As long as Assumption 1 holds, i.e. $\frac{\mu}{\theta} < 1$, there exist values of M that satisfy left-side of the implication symbol in Eq. (32). The safety threshold $\underline{M}$ can be though of as a value M that satisfies condition (32) by a small safety margin ϵ:

$$\underline{M} = A\left[2\left(1-\frac{\mu}{\theta}\right)^{-1} - 1\right] + \epsilon$$

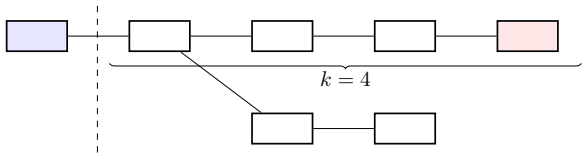

Figure 6: Confirmation rule for $k = 4$

With the security condition in (32) satisfied, sellers can wait that the blocks containing their payments are buried in the longest chain at a sufficiently large depth k before delivering their goods to buyers. In Fig. 6, $k = 4$; the blue block is confirmed as soon as the red block becomes part of the longest chain. This gives them time to react to malicious forks before complying with their contractual agreements.

[16] Their result follows by approximating Poisson processes taking the continuous-time limit of random walks as in Gaži et al. (2020).

The closing link between the security analysis and miner entry incentives is the fact that A is fixed while M is determined by financial gains from updating the blockchain. These are naturally more attractive to honest miners than to malicious ones because honest miners have a vested interest in the blockchain security it keeps positive the value crypto payments they receive from the protocol and merchants. As a consequence, as more honest miners M join the free-entry network in seek for revenues, also their share in computational power relative to a potential attacker grows.

Appendix B Proofs

Before presenting the formal proofs of my main claims, I will state two auxiliary results that provide background technical knowledge for some parts of the main sections and the demonstrations that will follow next.

Auxiliary results

The first auxiliary result relates to exponential random variables. These determine the inter-arrival times in the Poisson processes that I use throughout the paper. They also enter in the calculation of value function driven by such processes.

Lemma B.1. *Let $\{t_j\}_{j\in\{1,2,\dots,J\}}$ denote a collection of ERV's with rates $\{\lambda_j\}_{j\in\{1,2,\dots,J\}}$ such that $\lambda = \sum_{j\in\{1,2,\dots,J\}} \lambda_j$. Let h_t denote a generic function of time and r a constant. Let $V_{(j)}$ denote a value function in state j. V is initialized in state 0, has rate λ_j of switching from state 0 to state $j \in \{1, 2, \dots, J\}$, and discounts time at rate r. The following claims hold:*

- *I (Laplace transform):* $\mathbb{E}\left(e^{-rt_j}\right) = \frac{\lambda_j}{r+\lambda_j}$
- *II (Minimum of ERV's):* $t_{jmin} \triangleq \min\{t_j\}_{j\in\{1,2,\dots,J\}}$ *is an ERV with rate λ. Moreover,* $\mathbb{P}(t_{jmin} = t_j) = \lambda_j/\lambda$.
- *III (Stochastic integration limit):* $\mathbb{E}\left(\int_0^{t_j} e^{-rt} h_t \,\mathrm{d}t\right) = \int_0^{+\infty} e^{-(r+\lambda_j)t} h_t \,\mathrm{d}t = \frac{h}{r+\lambda_j}$ *for $h_t = h \ \forall t$*

Proof. Lemma B.1.I and Lemma B.1.II are standard properties. Their proofs can be found in most graduate probability textbooks. For Lemma B.1.III, proceed integrating by parts.

$$\begin{aligned}\mathbb{E}\left(\int_0^{t_j} e^{-rt} h_t \,\mathrm{d}t\right) &= \int_0^{+\infty} \lambda_j e^{-\lambda_j t_j} \left(\int_0^{t_j} e^{-rt} h_t \,\mathrm{d}t\right) \mathrm{d}t_j = \int_0^{+\infty} e^{-\lambda_j t_j} e^{-rt_j} h_{t_j} \,\mathrm{d}t_j \\ &= \int_0^{+\infty} e^{-(r+\lambda_j)t_j} h_{t_j} \,\mathrm{d}t_j \equiv \int_0^{+\infty} e^{-(r+\lambda_j)t} h_t \,\mathrm{d}t\end{aligned} \tag{33}$$

The last step in the chain of equalities in Eq. (33) uses the fact that $e^{-\lambda_j t_j}\int_0^{t_j} e^{-rt}h_t\,\mathrm{d}t = 0$ both at $t_j \to +\infty$ and $t_j = 0$. ■

The next auxiliary lemma presents the properties of mempool dynamics. Technically these are governed by the simplest queuing system, the M/M/1 queue. The long-run probability distribution of the mempool size $g(Q_t)$ reaches a steady state if the probability inflow and outflow across all states Q_t is balanced. This global balancing principle translates to $\dot{g}(Q) \triangleq \partial g(Q)/\partial t = 0$ for all $Q \geq 0$. This means

$$\begin{aligned} 0 \equiv \dot{g}\,(0) &= \mu g\,(1) - \alpha N g\,(0) \\ 0 \equiv \dot{g}\,(Q) &= \alpha N\; g\,(Q-1) + \mu g\,(Q+1) - (\alpha N + \mu)\, g\,(Q) \qquad \text{for } Q > 0 \end{aligned} \tag{34}$$

The solutions of system (34) are determined by the *load* $\rho = \frac{\alpha N}{\mu}$ (see Eq. (19))

Lemma B.2. *The M/M/1 queue Q_t explodes for $\rho \geq 1$. On the contrary, for $\rho \in [0,1)$, Q_t admits a steady state, stationary distribution. The probabilities $g(Q_t) = g(Q)$ have the geometric mass function*

$$g(Q) = (1-\rho)\rho^Q$$

Proof. See Kleinrock (1975) Chapter 3.2. ■

Notice that $\rho = 1 - g(0)$. In words, it is the fraction of time in which the mempool is non-empty.

Main results

Proof of Proposition 1. To derive Eq. (4), impose stationarity by letting $V_{t_0} = V_{t_0+T_{\mathsf{trade}}} = V$ and express $\mathrm{d}x_t$ using Eq. (3): $\mathrm{d}x_{t_0} = z$, $\mathrm{d}x_t = z\pi\,\mathrm{d}B_t$ for $t \in (t_0, T_{\mathsf{trade}})$. Eq. (2) now gives

$$V^n = \max_z -z + \mathbb{E}\left[-z\pi\int_0^{T_{\mathsf{trade}}} e^{-rt}\,\mathrm{d}B_t + e^{-rT_{\mathsf{trade}}}\left(u\left(\beta z\,(1-\tau)\right) + V^n\right)\right] \tag{35}$$

Since B_t is a Poisson counter, $\mathbb{E}\left(\mathrm{d}B_t\right) = \mu\,\mathrm{d}t + o(\mathrm{d}t)$. Also, T is exponentially distributed with density $\alpha e^{-\alpha r}$, so Lemma B.1 points II and III apply to solve the expectation in Eq. (35),

$$\begin{aligned} V^n &= \max_z -z + \int_0^{+\infty} e^{-(r+\alpha)t}\left[-z\mu\pi + \alpha\left(u\left(\beta z\,(1-\tau)\right) + V^n\right)\right]\mathrm{d}t \\ &= \max_z -z + \frac{1}{r+\alpha}\left[-z\mu\pi + \alpha\left(u\left(\beta z\,(1-\tau)\right) + V^n\right)\right] \end{aligned} \tag{36}$$

Multiplying both sides of Eq. (36) by $r+\alpha$ and rearranging we obtain V^n in Eq. (4). Eq. (5) follows from differentiating (4) with respect to z and relabeling variables. ■

Proof of Lemma 1. Since $c(a)$ has $a = a_{(Q_t,B_t)} \equiv a_{(Q,B)}$ constant during the current round and the game switches from the current round at the ERV T with rate $\alpha N + \mu$, Lemma B.1.III gives $\mathbb{E}\left(\int_0^T e^{-rt} c(a)\,\mathrm{d}t\right) = \frac{c(a)}{r+\alpha N+\mu}$. Now we have

$$\mathbb{E}\left[e^{-rT} R_{(Q_t,B_t)}(a,m^*)\right] = \mathbb{E}\left(e^{-rT}\right)\mathbb{E}\left[R_{(Q_t,B_t)}(a,m^*)\right] = \frac{\alpha N+\mu}{r+\alpha N+\mu}\mathbb{E}\left[R_{(Q_t,B_t)}(a,m^*)\right]$$

where the first equality comes from independence among T and $R_{(Q_t,B_t)}(a,m^*)$ and the last equality from Lemma B.1.I. Solving $\mathbb{E}\left(R_{(Q_t,B_t)}(a,m^*)\right)$ only requires weighting $R_{(Q_t,B_t)}(a,m)$ by $P\left(a,\mathbf{a}^{-m}\right)$ as $R_{(Q_t,B_t)}(a,\varnothing) = R_{(Q_t,B_t)}(a,m') = 0$. Lemma B.1.II gives $\mathbb{P}(T = T_{\mathsf{block}}) = \mathbb{P}(\min(T_{\mathsf{block}}, T_{\mathsf{trade}}) = T_{\mathsf{block}}) = \frac{\mu}{\alpha N+\mu}$, so letting $\mathbb{P}(m = m^* \mid m^* \neq \varnothing) \equiv P(a,\mathbf{a}^{-m})$, we have $\mathbb{E}\left[R_{(Q_t,B_t)}(a,m^*)\right] = \frac{\mu}{\alpha N+\mu} R_{(Q_t,B_t)}(a,m) P(a;\mathbf{a}^{-m})$.

The prior steps of the proof give

$$\begin{aligned}\mathbb{E}\left[e^{-rT} R_{(Q_t,B_t)}(a,m^*) - \int_0^T e^{-rt} c(a)\,\mathrm{d}t\right] &= \frac{1}{r+\alpha N+\mu}\left(\mu R_{(Q_t,B_t)}(a,m) P(a,\mathbf{a}^{-m}) - c(a)\right) \\ &\equiv \frac{1}{r+\alpha N+\mu} v_{(Q_t,B_t)}(a,\mathbf{a}^{-m})\end{aligned}$$

The missing step to complete the proof,

$$\max\left\{\frac{v_{(Q_t,B_t)}(a,\mathbf{a}^{-m})}{r+\alpha N+\mu} + \mathbb{E}\left(e^{-rT}\Pi_{t+T}\right)\right\} = \frac{\max v_{(Q_t,B_t)}(a,\mathbf{a}^{-m})}{r+\alpha N+\mu} + \mathbb{E}\left(e^{-rT}\Pi_{t+T}\right)$$

follows from the restriction to Markovian strategies and the fact that the mempool distribution only depends on $\mathbf{S}$, not on single miner actions. ■

Proof of Proposition 2. (i) is a direct consequence of the monotone-likelihood-ratio property of the densities $f_0(t)$ and $f_1(t)$ from the equation panel (15). $\frac{\mathrm{d}}{\mathrm{d}t}\frac{f_1(t)}{f_0(t)} = \theta\, e^{-t(\theta-\mu/M)} > 0$ implies that $f_1(\cdot)$ first-order-stochastically dominates $f_0(\cdot)$.

Now, $P(m0_o, m1_o|a)$ from Eq. (16) is the expectation of $(1-F_0(t))^{m0_o}(1-F_1(t))^{m1_o}$ over the density $f_a(t)$. Given that $(1-F_0(t))^{m0_o}(1-F_1(t))^{m1_o}$ is **decreasing** in t, $P(m0_o, m1_o|0) \geq P(m0_o, m1_o|1)$ by $f_1(t) \succeq_{\text{FOSD}} f_0(t)$.

For (ii) and (iii), just observe that the integrand in Eq. (16) decreases so the resulting integral is also lower.

(iv) is a standard property of IID random variables. Below I provide a proof that solves Eq. (16) integrating by parts. Let $F_a(t) \equiv F(t)$ and $f_a(t) \equiv f(t)$. As $\left[(1-F(t))^{M-1}F(t)\right]_{t=0}^{t\to\infty}$

$= 0,$

$$\int_0^\infty (1-F(t))^{M-1} f(t)\,\mathrm{d}t = (M-1)\int_0^\infty (1-F(t))^{M-2}F(t)f(t)\,\mathrm{d}t$$
$$= (M-1)\left[\int_0^\infty (1-F(t))^{M-2}f(t)\,\mathrm{d}t - \int_0^\infty (1-F(t))^{M-1}f(t)\,\mathrm{d}t\right]$$

Collecting the integrals, $\displaystyle\int_0^\infty (1-F(t))^{M-1}f(t)\,\mathrm{d}t = \frac{M-1}{M}\left[\int_0^\infty (1-F(t))^{M-2}f(t)\,\mathrm{d}t\right]$

The result is immediate proceeding by induction on M with base step $M=2$. ■

Proof of Proposition 3. (i) requires

$$P(m0_o-1, m1_o+1|0)P(m0_o, m1_o|1) - P(m0_o, m1_o|0)P(m0_o-1, m1_o+1|1) \geq 0 \quad (37)$$

For ease of exposition, let $P_o(t) \triangleq (1-F_0(t))^{m0_o}(1-F_1(t))^{m1_o}$. Define the operator $\psi(f,f') \triangleq \int_0^\infty P_o(t)f(t)\,\mathrm{d}t \cdot \int_0^\infty P_o(t)f'(t)\,\mathrm{d}t$ over two functions f and f' and let $\tilde{f}(t) \triangleq \frac{\mu}{M}e^{-t\theta}$ and $\hat{f}(t) \triangleq \frac{\mu}{M}e^{-t(2\theta-\mu/M)}$. Since $f_0(t) = \left(1-\frac{\mu}{\theta M}\right)f_1(t) + \frac{\mu e^{-t\theta}}{M}$ and $\frac{1-F_1(t)}{1-F_0(t)} = \frac{\theta}{\theta-\mu/M} - \frac{\mu/M}{\theta-\mu/M}e^{-t(\theta-\mu/M)}$,

$$P(m0_o-1, m1_o+1|0)P(m0_o, m1_o|1) = \frac{\theta^2\psi(f_0,f_0) - (\frac{\mu}{M}\theta+\theta^2)\psi(f_0,\tilde{f}) - \frac{\mu}{M}\theta\psi(\tilde{f},\tilde{f})}{(\theta-\mu/M)^2}$$
$$P(m0_o, m1_o|0)P(m0_o-1, m1_o+1|1) = \frac{\theta^2\left(\psi(f_0,f_0)-\psi(\tilde{f},f_0)\right) + \theta\frac{\mu}{M}\left(\psi(\hat{f},f_0)-\psi(\tilde{f},f_0)\right)}{(\theta-\mu/M)^2}$$

Using the previous expressions, inequality (37) simplifies to $\psi(\tilde{f},\tilde{f}) - \psi(\hat{f},f_0) \leq 0$. Now I will show that

$$\psi(\tilde{f},\tilde{f}) = \min_q \psi\left(\theta\frac{\mu}{M}e^{-t(2\theta-q)}, \theta\frac{\mu}{M}e^{-tq}\right)$$

Claim (i) is thus satisfied since $\psi(\hat{f},f_0)$ uses $q = \mu/M \neq \theta$. To prove that $\psi(\tilde{f},\tilde{f})$ is a minimum, compute the first and second-order minimization conditions.

$$D_q\psi\left(\theta\frac{\mu}{M}e^{-t(2\theta-q)}, \theta\frac{\mu}{M}e^{-tq}\right) = 0$$
$$\iff \int P_o e^{-t(2\theta-q)}t\,\mathrm{d}t \int P_o e^{-tq}\,\mathrm{d}t = \int P_o e^{-t(2\theta-q)}\,\mathrm{d}t \int P_o e^{-tq}t\,\mathrm{d}t \quad (38)$$

Eq. (38) holds iff $e^{-t(2\theta-q)} = e^{-tq}$. That is, for $q=\theta$. For sufficiency we need

$$D_q^2\psi\left(\theta\frac{\mu}{M}e^{-t(2\theta-q)},\theta\frac{\mu}{M}e^{-tq}\right)\Big|_{q=\theta}>0$$

$$\iff \int_0^\infty P_o(t)e^{-t\theta}\,\mathrm{d}t\int_0^\infty P_o(t)e^{-t\theta}t^2\,\mathrm{d}t>\left(\int_0^\infty P_o(t)e^{-t\theta}t\,\mathrm{d}t\right)^2$$

$$\iff \int_0^\infty\left(\sqrt{P_o(t)e^{-t\theta}}\right)^2\mathrm{d}t\int_0^\infty\left(\sqrt{P_o(t)e^{-t\theta}}t\right)^2\mathrm{d}t>\left(\int_0^\infty\sqrt{P_o(t)e^{-t\theta}}\sqrt{P_o(t)e^{-t\theta}}t\,\mathrm{d}t\right)^2$$

The last step holds by the Cauchy-Schwartz inequality. Moving to (ii),

$$L(M-1,0)=\frac{1/M}{\int_0^\infty e^{-t\frac{\mu(M-1)}{M}}\frac{\mu\theta/M}{\theta-\mu/M}\left(e^{-t\frac{\mu}{M}}-e^{-t\theta}\right)\mathrm{d}t} \tag{39}$$

The numerator of Eq. (39) follows from Proposition 2.ii. Its denominator simplifies to

$$\frac{\theta\frac{\mu}{M}}{\theta-\frac{\mu}{M}}\int_0^\infty e^{-t\mu}-e^{-t\left(\theta+\frac{\mu(M-1)}{M}\right)}\,\mathrm{d}t=\frac{\theta\frac{\mu}{M}}{\theta-\frac{\mu}{M}}\left[\frac{1}{\mu}-\frac{1}{\theta+\mu-\frac{\mu}{M}}\right]=\frac{\theta}{M(\theta+\mu)-\mu}$$

Hence Eq. (39) simplifies as $L(M-1,0)=\frac{M(\theta+\mu)-\mu}{M\theta}=1+\frac{\mu}{\theta}\left(1-\frac{1}{M}\right)$. The comparative statics results are immediate from the latest expression. ■

Proof of Lemma 2. The probabilities in panel (17) come from the general formulae (41) and (42) which give explicit solutions for any $(a,m0_o,m1_o)$, $M=1+m0_o+m1_o$. The simple procedure that follows derives them.[17] The only non-obvious step is the binomial expansion of $(1-F_1(t))^{m1_o}$ that facilitates integration.

$$(1-F_1(t))^{m1_o}=\left(\theta-\frac{\mu}{M}\right)^{-m1_o}\sum_{j=0}^{m1_o}\binom{m1_o}{j}(-1)^j\left(\frac{\mu}{M}\right)^j\theta^{m1_o-j}e^{-t\left[\theta j+\frac{\mu}{M}(m1_o-j)\right]} \tag{40}$$

$$P(m0_o,m1_o|0)=\frac{\mu}{M}\left(\theta-\frac{\mu}{M}\right)^{-m1_o}\sum_{j=0}^{m1_o}\binom{m1_o}{j}(-1)^j\left(\frac{\mu}{M}\right)^j\theta^{m1_o-j}\times\int_0^{+\infty}e^{-t\left[\theta j+\frac{\mu}{M}(m1_o+m0_o-j+1)\right]}\,\mathrm{d}t \tag{41}$$

$$P(m0_o,m1_o|1)=\theta\frac{\mu}{M}\left(\theta-\frac{\mu}{M}\right)^{-(m1_o+1)}\sum_{j=0}^{m1_o}\binom{m1_o}{j}(-1)^j\left(\frac{\mu}{M}\right)^j\theta^{m1_o-j}\times\int_0^{+\infty}\left(e^{-t\left[\theta j+\frac{\mu}{M}(m1_o+m0_o-j+1)\right]}-e^{-t\left[\theta(j+1)+\frac{\mu}{M}(m1_o+m0_o-j)\right]}\right)\mathrm{d}t \tag{42}$$

The solutions to the integrals in Eqs. (41) and (42) are respectively given by Eqs. (41')

[17] Mathematica code to verify the expressions available on request.

and (42'):

$$\left[\mu + j\left(\theta - \frac{\mu}{M}\right)\right]^{-1} \tag{41'}$$

$$\left[\mu + j\left(\theta - \frac{\mu}{M}\right)\right]^{-1} - \left[\mu + (j+1)\left(\theta - \frac{\mu}{M}\right)\right]^{-1} \blacksquare \tag{42'}$$

Proof of Lemma 4. Suppose a seller s is is notified of receiving a pending transaction. Let Q_t^{-s} denote the number of *other* transactions in the mempool by the time the blockchain grows by an additional block. The pending transaction is recorded in that block with probability $1/(1+Q_t^{-s})$. Since $Q_t^{-s} \equiv Q^{-s}$ is geometrically distributed with mass $g(Q^{-s})$ for all t (from Eq. (19)), $\sum_{Q^{-s}=0}^{\infty} g(Q^{-s})\frac{1}{1+Q^{-s}} = (1-\rho)\sum_{Q^{-s}=0}^{\infty}\frac{\rho^{Q^{-s}}}{1+Q^{-s}} = \frac{1-\rho}{\rho}\left(\sum_{Q=1}^{\infty}\frac{\rho^Q}{Q}\right) = \frac{1-\rho}{\rho}\ln(1-\rho)^{-1} \triangleq \nu$. The last equality follows from the Maclaurin expansion of $\ln(1-\rho)^{-1}$. Since $g(Q)$ is stationary and every block is formed independently, $d(b)$ is geometric with parameter ν. $\blacksquare$

Proof of Lemma 6. Since $a(Q_t, B_t)$ is independent of B_t and the distribution of Q_t is stationary, also the profit function is stationary. Letting $\Pi_t = \Pi$ for all t,

$$\Pi = \mathbb{E}\left[e^{-rT}\left(\mathbb{P}\left(T = T_{\mathsf{update}}\right)R + \Pi\right)\right] - \mathbb{E}\left[\int_0^T e^{-rt}c\,\mathrm{d}t\right] \tag{43}$$

$$= \frac{\alpha N + \mu}{r + \alpha N + \mu}\left(\mathbb{E}(R)\frac{\mu}{\alpha N + \mu} + \Pi\right) - \frac{c}{r + \alpha N + \mu} \tag{44}$$

Collecting Π on both sides,

$$\Pi = \frac{1}{r}\left[\mu E(R) - c\right] = \frac{1}{r}\left[\frac{\mu}{M}z\left(N\pi + \rho\tau\right) - c\right] = \frac{1}{r}\left[\frac{Nz(\mu\pi + \alpha\tau)}{M} - c\right] \quad \blacksquare$$

Proof of Lemma 7. W is the sum of

$$MV^m = \frac{\alpha}{r}z(\pi,\tau)\tau + \frac{\mu}{r}z(\pi,\tau)\pi - M\left(\frac{c}{r} + F\right) \tag{45}$$

$$V^n + V^s = \frac{\alpha}{r}\left[u\left(y(\pi,\tau)\right) - y(\pi,\tau) - z(\pi,\tau)\left(1 - \beta(1-\tau)\right)\right] - \frac{\mu}{r}z(\pi,\tau)\pi - z(\pi,\tau) \tag{46}$$

$\frac{\mu}{r}z(\pi,\tau)\pi$ cancels out completely in the sum of Eqs. (45) and (46). Contrarily, $\frac{\alpha}{r}z(\pi,\tau)\tau - \frac{\alpha}{r}z(\pi,\tau)(1-\beta(1-\tau)) = -\frac{\alpha}{r}z(\pi,\tau)(1-\tau)(1-\beta) = \frac{\alpha}{r}y(\pi,\tau)(\frac{1}{\beta}-1)$. The last equality uses $z(\pi,\tau) = \frac{y(\pi,\tau)}{\beta(1-\tau)}$ from Eq. (5) for $y(z(\pi,\tau)) \equiv y(\pi,\tau)$. Eq. (24) follows instantaneously from the resulting sum after expressing $z(\pi,\tau)$ as a function of $y(\pi,\tau)$ in the other terms. $\blacksquare$

Proof of Proposition 7. The proof proceeds in four steps . To begin with, set $\tau' < \tau$ and $\pi' > \pi$ as to satisfy Eq. (26). Step 1. W increases when $(\pi, \tau, y(\pi, \tau))$ shifts to $(\pi', \tau', y(\pi, \tau))$ while preserving PC′; Step 2. $y(\pi', \tau') > y(\pi, \tau)$; Step 3. W increases and PC′ becomes slack moving $(\pi', \tau', y(\pi, \tau))$ to $(\pi', \tau', y(\pi', \tau'))$; Step 4. $\tau' = 0$ maximizes W so $(\pi, \tau) = (\pi^*, 0)$, with π^* satisfying PC′ at $(\pi^*, 0, y(\pi^*, 0))$, is optimal.

Step 1. Immediate as W in Eq. (24) is clearly decreasing in τ and Eq. (26) holds miner revenues constant.

Step 2. Due to diminishing marginal utility, $u'(y(\pi', \tau')) < u'(y(\pi, \tau))$ implies $y(\pi', \tau') > y(\pi, \tau)$. Using the left-hand side of Eq. (5), $u'(y(\pi', \tau')) < u'(y(\pi, \tau))$ if and only if $\frac{1}{r\beta}\left(\frac{r+\alpha+\mu\pi'}{1-\tau'}\right) = \Phi_y(\pi', \tau', y) < \Phi_y(\pi, \tau, y)$. This inequality is satisfied by the last step of (27).

Step 3. Since $W = \max_{\pi,\tau} V^n = \max_{\pi,\tau}\left(\max_y \; -y\Phi_y(\pi, \tau, y) + \frac{\alpha}{r}u(y)\right)$. $\Phi_y(\pi', \tau', y) < \Phi_y(\pi, \tau, y)$ so the objective of the inner maximization problem increases point-wise in y. Hence the outer maximization necessarily achieves a higher maximum, with $y(\pi', \tau') > y(\pi, \tau)$ from Step 2.

Step 4. The previous steps show that reducing τ to $\tau' < \tau$ and increasing π up to the value $\pi'' \in (\pi, \pi')$ that satisfies PC′ at $(\pi'', \tau', y(\pi'', \tau'))$ is welfare-improving and feasible at any $\tau > 0$. The planner can raise welfare with this policy change unless $\tau = \tau^* = 0$. ■