https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 69, No. 11, November 2023, pp. 6684–6707
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

A Model of Cryptocurrencies

Michael Sockin,[a,*] **Wei Xiong**[b,c]

[a] McCombs School of Business, University of Texas at Austin, Austin, Texas 78712; [b] Bendheim Center for Finance, Princeton University, Princeton, New Jersey 08544; [c] School of Management and Economics, Chinese University of Hong Kong, Shenzhen, Guangdong 518172, China
*Corresponding author

Contact: Michael.Sockin@mccombs.utexas.edu (MS); wxiong@princeton.edu, https://orcid.org/0000-0003-3592-9373 (WX)

Received: February 6, 2022
Revised: August 24, 2022; December 15, 2022
Accepted: March 11, 2023
Published Online in Articles in Advance: April 11, 2023

https://doi.org/10.1287/mnsc.2023.4756

Abstract. We model cryptocurrencies as utility tokens used by a decentralized digital platform to facilitate transactions between users of certain goods or services. The network effect governing user participation, in conjunction with the nonneutrality of the token price, can cause the token market to break down. We show that token retradeability mitigates this risk of breakdown on younger platforms by harnessing user optimism but worsens this fragility when sentiment trading by speculators crowds out users. Elastic token issuance mitigates this fragility, but strategic attacks by miners exacerbate it because users' anticipation of future losses depresses the token's resale value.

History: Accepted by Agostino Capponi, Special Section of *Management Science*: Blockchains and Crypto Economics.

Keywords: cryptocurrency • token price nonneutrality • optimism • platform fragility

1. Introduction

The rapid growth of the cryptocurrency market in the last few years promises a new funding model for innovative digital platforms. Rampant speculation and volatility in the trading of many cryptocurrencies, however, have also raised substantial concerns that associate cryptocurrencies with potential bubbles. The failure of the DAO only a few months after its initial coin offering (ICO) raised $150 million in 2016, together with a number of other similar episodes, particularly highlights the risks and fragility of cryptocurrencies. Understanding the risks and potential benefits of cryptocurrencies requires a systematic framework that incorporates several integral characteristics of cryptocurrencies—their role in funding digital platforms and in serving as investment assets for speculators and their integration of blockchain technology with decentralized consensus protocols to record transactions on the platforms. We develop such a model in this paper.

Our model analyzes the properties of cryptocurrencies on platforms that rely on network effects. Cryptocurrencies cover a wide range of tokens and coins facilitated by crypto technologies. For simplicity, we anchor our analysis on utility tokens, but our model can also be applied to coins and altcoins. Utility tokens are native currencies accepted on decentralized digital platforms that often provide intrinsic benefit to participants.[1] The benefits of utility tokens can range from provision of secure and verifiable peer-to-peer transaction services to the maintenance of smart contracts. Examples of such utility tokens include Ether, which enables participants to write smart contracts with each other; Filecoin, which matches the demand and supply for decentralized computational storage; and GameCredits, which finances the purchase, development, and consumption of online games and gaming content. The development of these platforms is financed by the sale of tokens to investors and potential users through the issuance of utility tokens.

We follow Sockin and Xiong (2023) to model a cryptocurrency as membership in a platform, which has been created by its developer to facilitate decentralized bilateral transactions of certain goods or services among a pool of users by using a blockchain technology. Users face difficulty in making such transactions outside the platform as a result of severe search frictions. The platform fills the users' transaction needs by pooling a large number of users who need to transact with each other. A user's transaction need is determined by its endowment in a consumption good and its preference of consuming its own good together with the goods of other users. As a result of this preference, users need to trade goods with each other, and the platform serves to facilitate such trading. Specifically, when two users are randomly matched, they can trade their goods with each other only if they both belong to the platform. Consequently, there is a key network effect—each user's desire to join the platform grows with the number of other users on the platform and the size of their goods endowments. If more users join the platform, each benefits more from joining the platform and is willing to pay a higher token price. Sockin and Xiong (2023) highlight that tokenization helps to decentralize the control of the platform and makes it possible for the platform to commit to not exploiting its users. This commitment, however, comes at the expense of not having an owner with an equity stake to subsidize user participation and maximize the platform's network effect.[2]

Our analysis builds on two key features. First, a user's benefit from using the token is increasing in the quantity, rather than in value in fiat currency, of tokens that she holds. This assumption is motivated by the nonneutrality of money that underlies modern monetary theory. This can arise, for instance, because of stickiness on the platform in adjusting the number of tokens required for its services in response to token price fluctuation.[3] As a result of such stickiness, a shock to the token price can directly affect user participation. This, in turn, amplifies the price shock through the platform's network effect. Second, in our model, the supply of tokens to users is decentralized in that token issuance follows a predetermined schedule and that token market participants are atomistic and therefore, do not internalize how their trading impacts others.

Our model features infinitely many periods, with users and speculators holding different beliefs about the capital gain from holding the token. In each period, a new generation of users chooses whether to join the platform by purchasing tokens from both existing token holders and from new token issuance by the platform. In deciding whether to join the platform, a user trades off the cost of buying a token with the benefits from both transacting goods on the platform and expected token price appreciation. Each user optimally adopts a cutoff strategy to join the platform by purchasing the token only if its goods endowment is higher than a threshold. This threshold and the token price are jointly determined by users' token demand, which is based on their common goods endowment and optimism about token price appreciation and the net supply of tokens by speculators; this is also determined by their sentiment about token price appreciation. Despite the inherent nonlinearity induced by each user's cutoff strategy, we derive the equilibrium in an analytical form and systematically characterize the platform's performance.

Our analysis highlights the fragility of cryptocurrency platforms induced by the network effect of user participation and decentralized token trading. Because of the network effect, users' demand curve for tokens is hump shaped (rather than downward sloping), whereas the net token supply faced by users is upward sloping. As a result, even though a trivial equilibrium with zero user demand and zero token price always exists, the token price may fail to simultaneously clear the supply and demand for tokens with positive user participation. In this case, the token market breaks down, which occurs when the platform's demand fundamental is sufficiently weak. Such market breakdown represents a severe form of distortion induced by token price on user participation through the network effect. Note that this distortion is specifically relevant to platforms with nonneutral token prices but not platforms that adjust the number of tokens required for their services in response to token price fluctuations.

Users' optimism about token price appreciation can alleviate this instability by inducing users to join the platform even when their transaction needs are low. In contrast, speculators' sentiment exacerbates this fragility by raising the cost for users to participate and crowding them out. Consequently, token retradeability is a powerful tool for improving platform performance when it capitalizes on user optimism. In contrast, it harms performance when it incentivizes outsiders, like speculators, to hold tokens as well, as their enthusiasm acts as a tax on user participation and exacerbates the platform's instability. Furthermore, elastic token issuance mitigates this fragility.

Because the supply of tokens increases deterministically over time, the platform exhibits life-cycle effects that are governed by the substitution of the token's current convenience yield and expected capital gains, which jointly determine the total token return to each user. The inflation of the token base over time lowers expected capital gains by shifting out the token supply curve. As a result, the region of market breakdown and the relative weight of the convenience yield in the total token return increase over time. Both of these effects, in turn, raise the sensitivity of the user base to the current demand fundamental and log token price volatility over time. We illustrate that more mature platforms not only have lower expected log token prices but also, higher log token price volatility and that these life-cycle effects are more pronounced for platforms whose fundamentals have weaker growth rates. Consequently, the ability of retradeability to harness the optimism of users to mitigate platform stability declines as the platform matures.

To further illustrate how outsiders hamper platform performance, we extend the model to incorporate miners who provide accounting and custodial services to record transactions on the platform's blockchain according to the proof of work (PoW) protocol. Each miner incurs a computational cost in providing the service and is compensated by the seigniorage from token inflation, which diminishes deterministically over time, and a transaction fee, which is a fraction of the transaction surplus of the users on the platform. This trade-off determines the number of miners on the platform. When the number of miners falls sufficiently low, some corrupt miners may choose to attack the cryptocurrency so that they can benefit from creating fraudulent seigniorage and stealing other miners' transaction fees. Although such attacks do not directly lead the platform to fail, our analysis shows that users' anticipation of future losses from miner attacks may exacerbate the fragility of the token market, especially when the mining cost is high. Consequently, having outsiders with whom there is a conflict of interest with users exacerbates the instability of cryptocurrency platforms.

Our framework provides a rich set of empirical predictions for token price appreciation. As only part of users' token return, the expected token price appreciation is determined by the marginal user's equilibrium condition—it equals the total cost of capital and participation minus the convenience yield from transaction surplus. Consistent with Liu and Tsyvinski (2021), our model predicts a role for both news and investor sentiment in explaining the time series of cryptocurrency price appreciation, not through risk premia but rather, by predicting the marginal user's convenience yield. In addition, our model can rationalize the momentum patterns that they observe in token price appreciation through the persistence of user participation costs and convenience yields, as well as the size effect that Liu et al. (2022) show in the cross-section of cryptocurrency price appreciation. Nonfundamental shocks to token prices, represented by user optimism and speculator sentiment in our model, can also help explain reversals in cryptocurrency returns, consistent with the evidence of a "value" factor in Cong et al. (2022a). Importantly, our asset pricing predictions are applicable only to tokens on platforms with nonneutral token prices and would not apply, for instance, to (alt-)coins and tokens on platforms with token neutrality, such as stablecoins and nonfungible tokens (NFTs).

Our paper contributes to a literature that studies instability on cryptocurrency platforms. Cong et al. (2021b) show that network effects amplify utility token price but mitigate user base volatility. Biais et al. (2023) develop a structural model of Bitcoin with transaction benefits and costs from hacking and show that Bitcoin is subject to significant extrinsic volatility because of coordination on sunspot equilibria. Pagnotta (2022) shows in an equilibrium model of Bitcoin that the interaction between the network of users and the investment of miners in network security amplifies Bitcoin price volatility. Mei and Sockin (2022) illustrate how speculation as an outside option can slow down learning on token platforms with network effects and lead to participation traps. In contrast to these papers, we highlight how the network effect among users can lead to market breakdown when there is nonneutrality of the token price on users' participation decisions. Furthermore, this instability is mitigated by user optimism and exacerbated by speculator sentiment because the latter shifts upward the supply curve of tokens. We further illustrate how users' anticipation of strategic attacks amplifies this fragility by reducing the token's expected retrade value. Our mechanism also differs from the impact of speculation on other asset classes, such as stocks and commodities, in which speculation can increase price volatility but not lead to a breakdown in which price and demand both collapse to zero.

Our paper is also related to the emerging literature on cryptocurrencies. Our model shares a similar pricing model but differs by deriving a strong network effect in the transaction benefits of the cryptocurrency as well as subtle interactions between strategic attacks by miners and the cryptocurrency's fragility. Cong et al. (2021) also emphasize the strong network effect among platform users. They construct a dynamic model of crypto tokens to study the dynamic feedback between user adoption and the responsiveness of the token price to expectations about future growth on the platform. In contrast to the monetary neutrality assumed in their model, which ensures that the token market is always stable, our model assumes that the token price is nonneutral. This key assumption, together with the network effect in user participation, underlies our mechanism that induces platform fragility. In addition, we show that miner attacks may exacerbate the platform fragility through the users' anticipation of losses from future attacks. Athey et al. (2016) model Bitcoin as a medium of exchange of unknown quality that allows users to avoid bank fees when sending remittances, and they use the model to guide an empirical analysis of Bitcoin users. Schilling and Uhlig (2019) study the role of monetary policy in the presence of a cryptocurrency that acts as a private fiat currency. Mayer (2019) finds that speculators provide or take liquidity from adopters depending on how volatile the platform fundamental is. In contrast to these papers and as a key contribution of our analysis, we examine token prices and platform performance with a realistic information structure that allows us to examine the role of optimism among users and sentiment among speculators. Goldstein et al. (2019) show that when there is token nonneutrality on an online platform, utility tokens that trade in secondary markets can act as a commitment device for an owner to price services competitively.

Our analysis also contributes to the literature on frictions in consensus validation on cryptocurrency platforms. Easley et al. (2019) analyze the rise of transaction fees in Bitcoin through the strategic interaction of users and miners. Chiu and Koeppl (2022) consider the optimal design of a cryptocurrency and emphasize the importance of scale in deterring double spending by buyers. Cong and He (2019) investigate the trade-off of smart contracts in overcoming adverse selection while also facilitating oligopolistic collusion, whereas Biais et al. (2019) consider the strategic interaction among miners. Pagnotta (2022) examines the strategic interaction among miners on the Bitcoin platform. Capponi et al. (2023) illustrate how the nature of mining may lead to a concentration of mining power, whereas Abadi and Brunnermeier (2018) examine disciplining writers to a blockchain technology with static incentives. Saleh (2021) explores how decentralized consensus can be achieved with the proof of stake protocol. Even without strategic attacks, Capponi et al. (2021) demonstrate how miners can impose more subtle costs on users by leaking information about their transactions for front running.

2. The Model

Consider a cryptocurrency that facilitates transactions on a decentralized digital platform. The platform serves to reduce search frictions among a pool of users who share a certain need to transact goods with each other. The benefits of participating on a utility token platform, such as Ether or FileCoin, include securing transactions and writing smart contracts to sharing gaming content and providing secure file storage. As the value of the token may appreciate with the development of the platform over time, the token also serves as an investable asset for users and speculators to speculate about the growth of the platform.

The model is discrete time with infinitely many periods: $t = 1, 2, \ldots$ There are three types of agents on the platform: users, speculators, and validators. The success of the cryptocurrency is ultimately determined by whether the platform can gather a large number of users together. In each period, a new generation of users purchases the cryptocurrency as the membership to the platform, and then, these users are randomly matched with each other to transact their goods endowments. The goods transactions are supported by validators of the decentralized platform who act as service providers and complete all user transactions. They record these transactions in an indelible ledger called the blockchain. A key feature of the blockchain technology underpinning cryptocurrencies is that it is permissionless and verifies transactions through decentralized consensus among an anonymous population of validators. For now, we assume there are no issues of trust on the platform. We will extend the model in Section 4 to incorporate decentralized miners who follow the PoW protocol to record transactions and may collude to strategically attack the platform.

2.1. Users

There are overlapping generations of users who join the platform. In each period t, there is a pool of potential users, indexed by $i \in [0, 1]$. Each of these potential users is endowed with a different consumption good and needs to transact her good with another user so that each user can consume two goods. To complete such a transaction, both users need to participate on the platform by purchasing a unit of the cryptocurrency, which we call a token of the platform. We can divide the unit interval into the partition $\{\mathcal{N}_t, \mathcal{O}_t\}$ in each period, with $\mathcal{N}_t \cap \mathcal{O}_t = \varnothing$ and $\mathcal{N}_t \cup \mathcal{O}_t = [0, 1]$. Let $X_{i,t} = 1$ if user i purchases the token (that is, $i \in \mathcal{N}_t$) and $X_{i,t} = 0$ if he chooses not to purchase. An indivisible unit of currency is commonly employed in search models of money (Kiyotaki and Wright 1993). If user i at $t = 1$ chooses to purchase the token, he purchases one unit at the equilibrium price P_t, denominated in the consumption numeraire. In the next period $t + 1$, each user from period t resells his token to future users and to speculators.

We follow Sockin and Xiong (2023) to model the users' transactions on the platform. In each period, user i is endowed with a certain good and is randomly paired with a potential trading partner, user j, who is endowed with another good. Users i and j can transact with each other only if both have the token. After their transaction, user i has a Cobb–Douglas utility function over consumption of his own good and the good of user j according to

$$U_{i,t}(C_{i,t}, C_{j,t}; \mathcal{N}_t) = \left(\frac{C_{i,t}}{1 - \eta_c}\right)^{1-\eta_c} \left(\frac{C_{j,t}}{\eta_c}\right)^{\eta_c}, \tag{1}$$

where $\eta_c \in (0, 1)$ represents the weight in the Cobb–Douglas utility function on his consumption of his trading partner's good $C_{j,t}$ and $1 - \eta_c$ is the weight on the consumption of his own good $C_{i,t}$. A higher η_c means a stronger complementarity between the consumption of the two goods. Both goods are needed for the user to derive utility from consumption. If one of them is not a member of the platform, there is no transaction, and consequently, each of them gets zero utility. This setting implies that each user cares about the pool of users on the platform, which determines the probability of completing a transaction.

The goods endowment of user i is $e^{A_{i,t}}$, where $A_{i,t}$ is composed of a component A_t common to all users and an idiosyncratic component $\varepsilon_{i,t}$:

$$A_{i,t} = A_t + \tau_\varepsilon^{-1/2} \varepsilon_{i,t},$$

with $\varepsilon_{i,t} \sim \mathcal{N}(0, 1)$ being normally distributed and independent with each other, across time, and from A_t. We assume that $\int \varepsilon_{i,t} d\Phi(\varepsilon_{i,t}) = 0$ at each date by the strong law of large numbers. The aggregate endowment A_t follows a random walk with a constant drift $\mu \in \mathbb{R}$:

$$A_t = A_{t-1} + \mu + \tau_A^{-1/2} \varepsilon_{t+1}^A,$$

where $\varepsilon_{t+1}^A \sim$ i.i.d $\mathcal{N}(0, 1)$. The aggregate endowment A_t is a key characteristic of the platform. A cleverly designed platform serves to attract users with strong needs to transact with each other. As we will show, a higher A_t leads to more users on the platform, which in turn, implies a higher probability of each user completing a transaction with another user, and furthermore, each transaction gives greater surpluses to both parties. One can, therefore, view A_t as the demand fundamental for the cryptocurrency and τ_ε as a measure of dispersion among users in the platform.[4]

We start with describing each user's problem in period t, conditional on joining the platform and meeting a transaction partner, and then, we go backward to describe his earlier decision on whether to join the platform. At t, when user i is paired with another user j on the platform, we assume that they simply swap their goods, with user i using $\eta_c e^{A_{i,t}}$ units of good i to exchange for $\eta_c e^{A_{j,t}}$ units of good j. Consequently, both users are able to consume both goods, with user i consuming

$$C_{i,t}(i) = (1-\eta_c)e^{A_{i,t}},\ C_{j,t}(i) = \eta_c e^{A_{j,t}}$$

and user j consuming

$$C_{i,t}(j) = \eta_c e^{A_{i,t}},\ C_{j,t}(j) = (1-\eta_c)e^{A_{j,t}}.$$

As formally shown by Sockin and Xiong (2023), these consumption allocations between these two paired users can be microfounded through a trading mechanism between them. Furthermore, we can use Equation (1) to compute the utility surplus $U_{i,t}$ of each user from this transaction.

Before finding a transaction partner on the platform, each user needs to decide whether to join the platform by buying the token. In addition to the utility surplus, $U_{i,t}$, from the transaction, there is also a capital gain from retrading the token, $P_{t+1} - RP_t$, with $R \geq 1$ being the interest rate for the holding period. We assume that users have quasilinear expected utility and incur a linear utility gain equal to this capital gain net of a fixed participation cost $\kappa > 0$ if they choose to join the platform. The participation cost may be either pecuniary or mental and could represent, for instance, the cost of setting up a wallet and installing the software necessary for participating on the platform. Furthermore, we assume that each user needs to give a fraction β of his utility surplus $U_{i,t}$ from the transaction as the service fee to the platform.

In summary, user i makes his purchase decision at t according to

$$\max_{X_{i,t}} \left(\mathbb{E}[(1-\beta)U_{i,t} + P_{t+1}|\mathcal{I}_{i,t}] - RP_t - \kappa\right)X_{i,t}, \tag{2}$$

where $\mathcal{I}_{i,t}$ is the information set of user i at date t. Note that the expectation of the user's utility flow regards the uncertainty associated with matching a transaction partner, whereas the expectation of the capital gain from holding the token regards the uncertainty in the growth of the platform. By adopting a Cobb–Douglas utility function with quasilinearity in wealth, users are risk neutral with respect to the token's capital gain.[5]

By treating the token as a membership to the platform, our model simplifies each user's token demand to a binary choice. In Appendix B, we consider a more general setting in which each user's benefit from holding the token increases in the quantity of tokens that she holds. We show that our key results on platform fragility do not depend on this binary token demand assumption. Instead, what is key to our analysis is the nonneutrality of the token price—the platform does not adjust the number of tokens required for each user to qualify for the platform's matching services. As a result, token price fluctuations directly affect user participation, which may be amplified further by the network effect of user participation.[6]

An important aspect of our analysis is how the weights of the token's convenience yield and capital gain transition over the life of the platform. When the platform is young, there are few tokens in circulation, and users benefit more from the token price appreciation. When the platform matures, there are many tokens in circulation, and users benefit mostly from the convenience yield from transactions on the platform. As we will analyze later, this transition underlies several interesting life-cycle implications; more mature platforms might be more vulnerable to market breakdown, younger platforms might have higher market capitalizations, and token price volatility is increasing over time.

We now describe the information set, $\mathcal{I}_{i,t}$, of each user. In addition to observing the platform fundamental, A_t, each user knows the value of his own goods endowment, $A_{i,t}$. To facilitate our analysis of how users' speculation of the token price may affect their participation in the platform, we also endow all users with a public signal about the next period's innovation to aggregate endowment, ε^A_{t+1}, which by construction, is orthogonal to A_t:

$$Q_t = \varepsilon^A_{t+1} + \tau_Q^{-1/2}\varepsilon^Q_t,$$

where $\varepsilon^Q_t \sim$ i.i.d $\mathcal{N}(0,1)$. This public signal is similar to a "news" shock in the language of Beaudry and Portier (2006). Because the public signal only reveals information about next period's A_{t+1}, it only impacts users' decisions through their beliefs about the next period's token price, $\mathbb{E}[P_{t+1}|\mathcal{I}_{i,t}]$, and therefore, it represents a speculative shock to all the users. Even though we use the term "user optimism" to denote the speculative shock induced by the public signal Q_t, the users are fully rational in information processing in our model. Consequently, $\mathcal{I}_{i,t} = \sigma(\{A_{i,t}, \{P_s, Q_s\}_{s\leq t}\})$ is user i's full information set.

It then follows that user i's purchase decision is given by

$$X_{i,t} = \begin{cases} 1 & \text{if } \mathbb{E}[(1-\beta)U_{i,t} + P_{t+1} - RP_t \mid \mathcal{I}_{i,t}] \geq \kappa \\ 0 & \text{if } \mathbb{E}[(1-\beta)U_{i,t} + P_{t+1} - RP_t \mid \mathcal{I}_{i,t}] < \kappa. \end{cases}$$

As the user's expected utility is monotonically increasing with his own endowment, regardless of other users' strategies, it is optimal for each user to use a cutoff strategy. This, in turn, leads to a cutoff equilibrium, in which only users with endowments above a critical level A^*_t buy the token. This cutoff is eventually solved as a fixed point in the equilibrium to equate the token price, net of the expected resale value and participation cost, with the expected transaction utility of the marginal user from joining the platform. As each user's participation strategy also depends on his expected token resale value $\mathbb{E}[P_{t+1}|\mathcal{I}_{i,t}]$, the common optimism among users induced by Q_t helps to overcome their participation cost κ. Because a user receives all his transaction benefit from holding one token, he will never buy a second and pay the participation cost for just the capital gain. This is because the token price will equal the capital gain plus the transaction benefit of the marginal user.

Given the cutoff strategy for each user who participates if $A_{i,t} \geq A_t^*$, the total token demand of users N_t is given by

$$N_t = \int_{-\infty}^{\infty} X_{i,t}(\mathcal{I}_{i,t})d\Phi(\varepsilon_{i,t}) = \Phi(\sqrt{\tau_\varepsilon}(A_t - A_t^*)). \tag{3}$$

2.2. Token Supply and Speculators

Consistent with the common practice of decentralized crypto platforms, we assume that the token supply, $\Phi(y_t)$, grows over time according to a predetermined schedule:

$$\Phi(y_t) = \Phi(y_{t-1} + \iota), \tag{4}$$

where $\Phi(\cdot)$ is the normal cumulative distribution function. Token platforms may commit to a predetermined inflation schedule, for instance, to mitigate incentives of validators to exploit users through excessive issuance but at the cost of suboptimal platform performance.[7] It also reflects that the inflation rules employed in practice do not, for instance, condition on platform performance; secondary market trading conditions; or the distribution of tokens across users, speculators, and validators. This specification further captures, as in practice, that the increase in supply from token inflation tapers over time. For PoW platforms, such as Bitcoin and Ethereum before the Merge, the number of new coins and tokens created by inflation periodically halves over time, according to a predetermined schedule, so that the total supply asymptotes to a fixed limit.[8] With our specification, at most a unit measure of tokens exists. All of our key qualitative results are unchanged, however, if instead we capped token supply at some maximum $\overline{y} < \infty$.

In addition to the token inflation, we assume that there is a continuum of atomistic and myopic speculators who trade the token for investment and speculation purposes. Speculators provide liquidity by buying tokens, including those from the old generation of users, and then selling them to the new generation of users. In our model, the trading of the token is fully decentralized in the sense that every participant is small and does not internalize the effects of her trading on others.

We consider speculators to be outsiders to the platform who are distinct from the users who actually participate on it. As such, they do not have private information about the platform's fundamental or fully understand how to interpret the implications of the same public information as the users. Instead, similar to Black (1986), we argue that they may trade overconfidently on noisy information or on spurious correlations that give rise to misspecified technical trading strategies. Hackethal et al. (2021), for instance, provide evidence that cryptocurrency investors are prone to investment biases, to following technical analysis heuristics, and to investing in stocks with high media sentiment. Fracassi and Kogan (2022) show that cryptocurrency investors trade based on pure technical analysis.

Specifically, we assume the following token demand curve for speculators:

$$X^S = \Phi(y_t) - \Phi\Big(y_t + \lambda \log(RP_t) - \zeta_t\Big),$$

where ζ_t is the common sentiment shock of speculators about the next-period token price. We separate speculators' sentiment from users' optimism so that we can analyze their distinct effects on the token market equilibrium. This demand curve has the property that when speculators, on average, are more optimistic (i.e., a higher ζ_t), their demand is higher and tightens the supply of tokens for users. In contrast, when the token price is higher, the usual downward-sloping demand effect leads to lower demand from speculators and a higher supply of tokens for users. In Appendix C, we provide a parsimonious microfoundation that derives this demand curve by aggregating the dispersed token demand among a group of atomistic active and passive speculators.

Market clearing in the token market consequently imposes that

$$\Phi(\sqrt{\tau_\varepsilon}(A_t - A_t^*)) + \Phi(y_t) - \Phi(y_t + \lambda \log(RP_t) - \zeta_t) = \Phi(y_t),$$

where we have substituted users' token demand with (3). This condition implies a token price:

$$P_t = \frac{1}{R}\exp\left(\frac{\sqrt{\tau_\varepsilon}}{\lambda}(A_t - A_t^*) - \frac{1}{\lambda}y_t + \frac{1}{\lambda}\zeta_t\right), \tag{5}$$

where the equilibrium token price P_t is a log-linear function of the platform's demand fundamental A_t, the users' participation threshold A_t^*, the token supply y_t, and speculator sentiment ζ_t.[9]

2.3. Validators

The platform requires record keeping of all transactions. For the baseline model, we assume that the decentralized platform has a group of validators who complete all user transactions each period without any frictions and record these transactions on the blockchain.[10] In a later section (Section 4), we expand the model to assume these validators record the transactions for a fee according to the PoW protocol and may also attack the cryptocurrency. In the baseline setting, the payment to validators in period t is both the seignorage from the scheduled inflation of the token base,

$\Phi(y_{t-1}+\iota)-\Phi(y_{t-1})$, and the transaction fees from users,

$$\pi_t = (\Phi(y_{t-1}+\iota)-\Phi(y_{t-1}))P_t + \beta U_t,$$

where U_t is the total transaction surplus on the platform. Validators have no use for tokens and potentially for liquidity reasons, sell them immediately to speculators. Assuming a cutoff strategy for users, we can integrate the expression for the expected utility of a user who joins the platform, as derived in Sockin and Xiong (2023), over $A_{i,t}$ for $A_{i,t} \geq A_t^*$ to arrive at the realized surplus from user transactions:

$$U_t = e^{A_t+\frac{1}{2}((1-\eta_c)^2+\eta_c^2)\tau_\varepsilon^{-1}}\Phi\left((1-\eta_c)\tau_\varepsilon^{-1/2}+\frac{A_t-A_t^*}{\tau_\varepsilon^{-1/2}}\right) \cdot \Phi\left(\eta_c\tau_\varepsilon^{-1/2}+\frac{A_t-A_t^*}{\tau_\varepsilon^{-1/2}}\right).$$

In contrast to Sockin and Xiong (2023), we assume that validators can commit to these policies.[11] As the platform's token base matures from inflation, the compensation to the validators shifts from seigniorage to transaction fees.

2.4. Rational Expectations Equilibrium

Our model features a rational expectations cutoff equilibrium, which requires the rational behavior of each user and the clearing of the token market.

- User optimization. Each user chooses $X_{i,t}$ in each period t to solve his maximization problem in (2) for whether to purchase the token.
- In each period, the token market clears

$$\int_{-\infty}^{\infty} X_{i,t}(A_{i,t},P_t)d\Phi(\varepsilon_{i,t}) = \Phi(y_t-\zeta_t+\lambda\log(RP_t)), \qquad (6)$$

where each user's demand $X_{i,t}$ depends on its information set $\mathcal{I}_{i,t}$. The demand from users is integrated over the idiosyncratic component of their endowments $\{\varepsilon_{i,t}\}_{i\in[0,1]}$, which also serves as the noise in their private information.

3. Equilibrium

We characterize the equilibrium in each period t when A_t and ζ_t are publicly observable. In this case, the token market is characterized by the following state variables: the users' demand fundamental A_t, the token supply y_t, the users' optimism driven by the public signal Q_t, and the speculators' sentiment ζ_t. We use the notation $\mathcal{I}_t = \{A_t, y_t, Q_t, \zeta_t\}$ to represent the state variables at t, which also represent the set of public information to all users. The public signal, Q_t, contains information about A_{t+1}, and thus, it is useful to users for forming their expectations about the token price in period $t+1$, P_{t+1}. Given that all users have a common expectation about P_{t+1}, we drop the i subscript from their information sets. After observing Q_t, users share the same posterior belief about A_{t+1}, which is normal with the following conditional mean:

$$\hat{A}_{t+1} = A_t + \mu + \frac{\tau_Q}{\tau_\varepsilon+\tau_Q}Q_t.$$

As we discussed earlier, the noise in Q_t is a shock to the users' speculative optimism because it has no impact on their current surplus from transacting with other users on the platform.

In each period, users sort into the platform according to a cutoff equilibrium determined by the net benefit of joining the platform, which trades off the opportunity of transacting with other users on the platform and the expected token price appreciation with the cost of participation. Despite the inherent nonlinearity of our framework, we derive a tractable cutoff equilibrium that is characterized by the solution to a fixed-point problem over the endogenous cutoff of the marginal user who purchases the token, A_t^*, as summarized in the following proposition.

Proposition 1. *The rational expectations equilibrium exhibits the following properties.*

1. *Regardless of other users' strategies, it is optimal for each user i to follow a cutoff strategy in purchasing the token:*

$$X_{i,t} = \begin{cases} 1 & \text{if } A_{i,t} \geq A^*(A_t,y_t,Q_t,\zeta_t) \\ 0 & \text{if } A_{i,t} < A^*(A_t,y_t,Q_t,\zeta_t) \end{cases}.$$

2. *In the equilibrium, the cutoff A_t^* solves the following fixed-point condition:*

$$(1-\beta)e^{(1-\eta_c)(A_t^*-A_t)+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2}-\frac{A_t^*-A_t}{\tau_\varepsilon^{-1/2}}\right)\mathbf{1}_{\{\tau>t\}} + \mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa = e^{-\frac{\sqrt{\tau_\varepsilon}}{\lambda}(A_t^*-A_t)-\frac{1}{\lambda}y_t+\frac{1}{\lambda}\zeta_t}, \qquad (7)$$

where τ is the stopping time for the breakdown of the platform because of the failure of the token market clearing

$$\tau = \{\inf t:\, A_t < A^c(y_t,Q_t,\zeta_t)\},$$

with $A^c(y_t,Q_t,\zeta_t)$ as a critical level for A_t, below which Equation (7) has no root.

3. *In each period t, there may be no or multiple equilibria with nontrivial user participation depending on the users' expected token resale value.*
   - *If $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa \leq 0$, Equation (7) has zero or two roots.*
   - *If $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa > 0$, Equation (7) has one or three roots.*
4. *In the dynamic equilibrium, the token price $P(A_t,y_t,Q_t,\zeta_t)$ is determined by Equation (5) according to the users' equilibrium cutoff A_t^* and how users coordinate on their expectations of future equilibria.*

Proposition 1 characterizes the cutoff equilibrium in the platform and confirms the optimality of a cutoff strategy for users in their choice to purchase the token. Users in each period sort into the platform based on their endowments, with those with higher endowments and thus, more gains from trade entering the platform. In this cutoff equilibrium, the token price is a correspondence of the token market state variables (A_t, y_t, Q_t, ζ_t), according to Equation (5), with A_t^* as an implicit function of these state variables.

Equation (7) provides a fixed-point condition to determine the optimal cutoff in each period. The left-hand side (LHS) of Equation (7) reflects the expected benefit to a marginal user with $A_{i,t} = A_t^*$ from acquiring a token to join the platform; the first term is the expected utility flow from transacting with another user on the platform, whereas the other terms $\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa$ represent the user's expected next-period token price net of the user's participation cost κ. The right-hand side (RHS) of Equation (7) reflects the cost of purchasing a token.

Figure 1 illustrates how the intersection of the two sides, each of which is plotted against the number of tokens N_t, determines the equilibrium cutoff A_t^*. Note that $N_t = \Phi(\sqrt{\tau_\epsilon}(A_t^* - A_t))$ is an inverse and monotonic transformation of $A_t^* - A_t$. The dashed hump-shaped line depicts the left-hand side of Equation (7) in a benchmark case when $\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa = 0$. That is, it captures a marginal user's expected utility flow from transacting with another user. This curve goes to zero when N_t goes to either zero or one. If $N_t \nearrow 1$ (i.e., $A_t^* \searrow 0$), the marginal user's own endowment is so low that there is no gain from transacting with the other user. On the other hand, if $N_t \searrow 0$ (i.e., $A_t^* \nearrow \infty$), the equilibrium cutoff is so high that there are no other users on the platform to transact with the marginal user. This network effect makes her expected utility from transaction zero, despite her own high endowment. Once the two end points are determined, it is intuitive that the marginal user's expected utility flow from transacting with another user on the platform has a hump shape. Such a hump-shaped demand curve is ubiquitous in the network effect literature (e.g., Easley and Kleinberg 2010). In the absence of the network effect, this demand curve is monotonically decreasing because the marginal user's expected utility is simply increasing with her own endowment A_t^*.

The right-hand side of Equation (7) is the supply curve of tokens and is represented by the solid upward-sloping curve. It is an exponential function of the inverse of the normal Cumulative Distribution Function of N_t (i.e., $\Phi^{-1}(N_t)$) because the number of users on the platform is decreasing with the equilibrium cutoff A_t^* and because the token price is an increasing function of the number of users as in Equation (5). In the absence of the network effect, this upward-sloping supply curve has a unique intersection with a downward-sloping demand curve. In the presence of the network effect (i.e., the hump-shaped demand curve), however, there may be several possibilities, as is well known in the network effect literature. Figure 1 illustrates that there is always a

Figure 1. (Color online) An Illustration of the Left- and Right-Hand Sides of Equation (7)

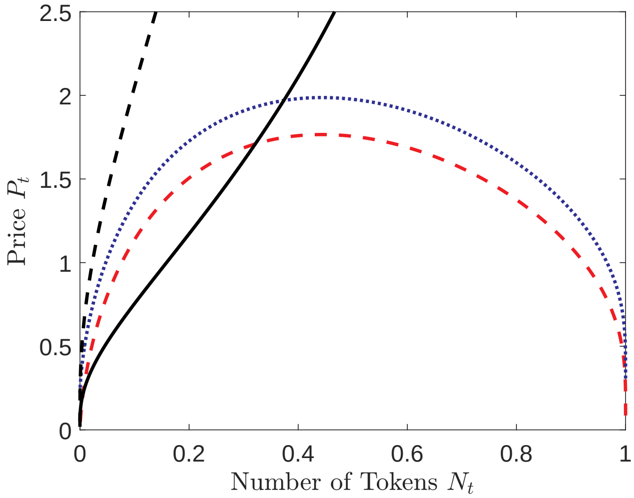

Note. The horizontal axis is the number of tokens $N_t = \Phi(\sqrt{\tau_\epsilon}(A_t^* - A_t))$, which is an inverse and monotonic transformation of $A_t^* - A_t$.

trivial solution in which demand and supply are both zero at a zero token price. Whether there is also a nontrivial solution depends on whether the dashed hump-shaped curve intersects the solid supply curve at some number of tokens greater than zero. As one can see, either this occurs twice or not at all if the solid supply curve lies above the hump-shaped curve for $N_t > 0$. The latter case is particularly important as it represents the breakdown of the token market and consequently, the failure of the platform. This happens when the expected utility from transacting is strictly lower than the cost of acquiring the token, either as a result of the small token supply y_t or as a result of strong speculator sentiment ζ_t. Proposition 1 shows that these two curves do not intersect when A_t falls below a critical level $A_t^c(y_t, Q_t, \zeta_t)$, which is determined by the other three state variables.

The terms $\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa$ may move the hump-shaped curve of the marginal user's expected benefit from participating in the platform up or down relative to the benchmark case. If $\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa > 0$, possibly as a result of the users' optimism about future token price appreciation (i.e., a positive shock to Q_t), the hump-shaped curve moves up relative to the benchmark dashed curve in Figure 1. In this case, the bell curve may intersect with the solid supply curve either once (as illustrated by the dotted curve) or three times.

If $\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa < 0$, either as a result of users' pessimism or a high participation cost κ, the hump-shaped curve moves down relative to the benchmark dashed line in Figure 1, creating the possibility for the token market to break down. That is, an increase in κ may lead to the failure of the platform in which there is only a trivial solution. As each user does not account for his participation decision on other users through the network effect, this externality exacerbates the effect of κ on the equilibrium user participation. Interestingly, users' optimism offsets the effect of their participation cost, which helps to overcome the network externality.

Finally, the dashed upward-sloping curve in Figure 1 illustrates the impact of speculator sentiment on market breakdown. An increase in speculator sentiment (a higher ζ_t) raises the solid black supply curve to the dashed black supply curve. This higher supply curve no longer intersects with the dashed hump-shaped demand curve. In this case, there is only a trivial equilibrium with zero user participation. Consequently, a higher speculator sentiment shifts up the token supply curve and makes it more difficult for there to be an equilibrium with nontrivial participation that clears the token market.

3.1. Market Breakdown

When there is only the equilibrium with zero user participation, the token market breaks down, and the platform fails. Such market breakdown represents a severe form of market dysfunction stemming from the network effect in user demand for tokens. It is important to recognize that this breakdown is a result of two key features of our model. First, token price fluctuations have a real effect on user participation because the platform does not adjust the number of tokens required for users to participate.[12] As discussed earlier, it is common for crypto platforms not to adjust the number of tokens required for their services in response to token price fluctuations. One may still be concerned about the role played by each user's binary token choice in driving the market breakdown. In Appendix B, we analyze a more general setting in which each user's benefit from holding tokens is monotonically increasing in the number of tokens she holds. This assumption maintains the nonneutrality of the token price but allows each user to choose a continuous number of tokens to hold. Interestingly, the token market may still break down because of the same mechanism illustrated by our main setting even in a more general setting, in which users choose tokens in a continuous quantity and face a general token supply curve.

The second key feature for market breakdown is that the token market is decentralized and no participant in the market internalizes the effect of her trading on others. Such externalities are present in both users and speculators. On the user side, no user accounts for the network effect of her participation choice on other users. On the speculator side, each speculator takes the token price as given, which implies that when the token market fails to find a market-clearing price, neither is a single speculator present nor can a group of speculators coordinate with each other to offer a price to clear the users' token demand.[13] It is also important to note that on a decentralized crypto platform, the platform founder by design is unavailable to support the secondary token market either by direct trading or by changing the token issuance protocol.[14]

The following proposition characterizes the conditions for market breakdown to occur.

Proposition 2. *As a result of the network effect, only an equilibrium with zero user participation exists (that is, the token market breaks down) under the following conditions.*

1. *The net speculative motive of users,* $\mathbb{E}[P_{t+1}|\mathcal{I}_t^*] - \kappa$, *is nonpositive.*

2. *The users' demand fundamental is sufficiently low (that is,* $A_t < A^c(y_t, Q_t, \zeta_t)$*), or equivalently, speculator sentiment is sufficiently high (that is,* $\zeta_t > \zeta^c(A_t, y_t, Q_t)$*).*

The critical level $A^c(y_t, Q_t, \zeta_t)$ *is decreasing in user optimism* Q_t *and increasing in speculator sentiment* ζ_t *and the user participation cost* κ.

Proposition 2 characterizes the determinants of the fundamental critical level $A^c(y_t, Q_t, \zeta_t)$ for the token market breakdown to occur. On the demand side, the users' speculative motive, driven by their optimism, helps to overcome the participation externality. On the supply side, speculators' sentiment has the opposite effect.

To further illustrate the properties of the token market equilibrium, we provide a series of numerical examples based on the parameter values given in Table 1. We caution, however, that this exercise is not a calibration but rather, an illustration of our model's behavior. To discipline our numerical examples, we follow Cong et al. (2021b) and choose a growth rate for the platform fundamental of $\mu = 0.02 - \frac{1}{2}\tau_A^{-1}$ (the term $-\frac{1}{2}\tau_A^{-1}$ arises because A_t is equivalent to the log of the fundamental in their model), a risk-free rate of $R = 1.05$, and a degree of complementarity of $\eta_c = 0.3$. We also follow Cong et al. (2022a) and choose a transaction fee rate of $\beta = 0.001$. Finally, we choose a token supply inflation rate of 4% ($\iota = 0.04$) based on Ethereum's average inflation rate before its conversion to proof of stake. We also choose reasonable values for the remaining parameters.

Figure 2 depicts the fundamental critical level A^c across speculator sentiment (the left panel), user optimism (the center panel), and token supply (the right panel). When the platform fundamental A is below A^c, the token market breaks down. The left panel shows that as speculator sentiment increases, the crowding out effect of speculators holding more tokens lowers user participation, shifting up the region of breakdown. In contrast, the center panel shows that an increase in user

Table 1. Baseline Model Parameters

	Model parameters
Demand fundamental	$\tau_A = 10$, $\mu = 0.02 - \frac{1}{2}\tau_A^{-1}$
Platform	$y_0 = -.84$, $\beta = 0.001$, $\iota = 0.04$
Sentiment	$\tau_Q = 5$, $\tau_\zeta = 2$, $\lambda = 1$
Users	$\tau_\theta = 1$, $\eta_c = 0.3$, $\kappa = 0.04$, $R = 1.05$

Figure 2. An Illustration of the Market Breakdown Boundary for the Demand Fundamental A^c with Respect to Speculator Sentiment (Left Panel), User Optimism (Center Panel), and Token Supply (Right Panel)

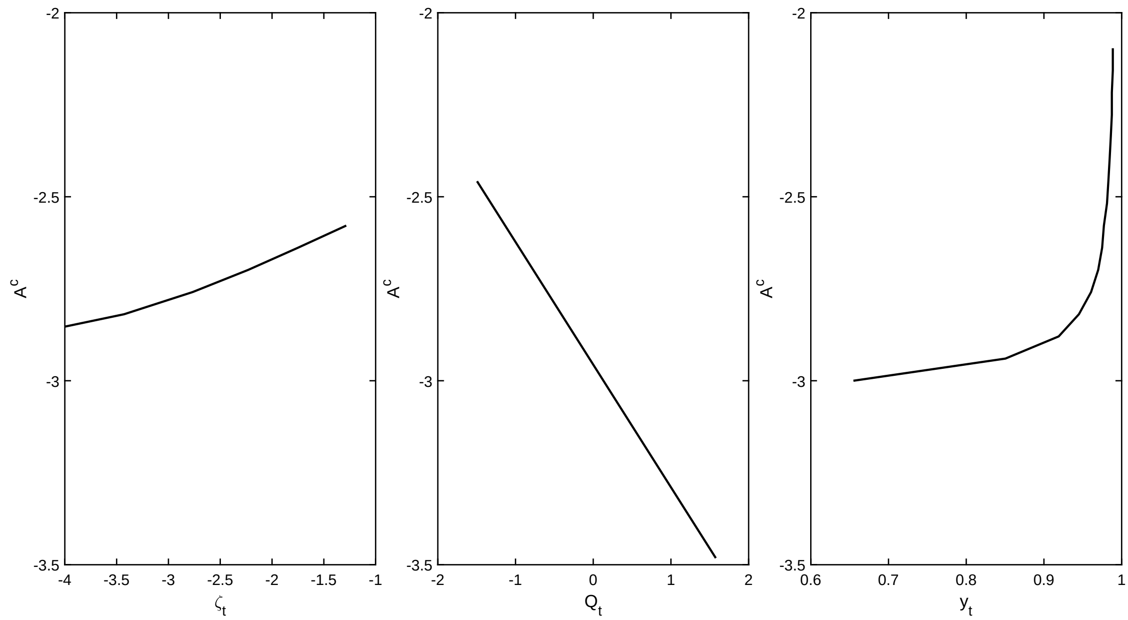

Note. The model parameters are given in Table 1, and the baseline values for the current state are $\zeta_t = 0$, $Q_t = 0$, and $y_t = 0.92$.

optimism, which incentivizes more users to participate, has the opposite effect and shifts down the region of breakdown. Taken together, these two panels illustrate the opposite effects generated by users' optimism and speculators' sentiment on the fragility of the platform, as formally established by Proposition 2.

The right panel of Figure 2 shows that an increase in token supply, by lowering the expected retrade value of the token, increases the breakdown boundary; to the left of the line, there is always an equilibrium. When the token base is small, there are at least two advantages. First, it is easier to clear markets with a small pool of users. Second, the expected growth of the token value is also higher. As the token supply inflates over time, the effects of token supply imply that the platform becomes more fragile over time, as the token's expected retrade value falls and user participation is driven more by the flow of convenience yields from transactions on the platform. This pattern thus suggests that large market capitalization tokens, such as Ethereum, might be more fragile and thus, have more pronounced price volatility than small capitalization tokens. Interestingly, although Cong et al. (2021b) emphasize the role of token resale in facilitating adoption, our model shows that it also helps to stave off the failure of the platform.

3.2. User Participation and Token Price

For the simplicity of our analysis, we assume that all users coordinate on the highest-price (i.e., the lowest-cutoff) equilibrium in each period, regardless of how many equilibria exist. One can motivate this refinement based on the (dynamic) stability of the potential equilibria.[15] Then, the following proposition derives several comparative statistics of the equilibrium user participation and token price.

Proposition 3. *The equilibrium has the following properties.*

1. *Demand fundamental. The token price and the fraction of users who participate in the platform are increasing in the demand fundamental,* A_t.

2. *User optimism. The token price and the fraction of users who participate in the platform are increasing in user optimism,* Q_t.

3. *Speculator sentiment. The fraction of users who participate in the platform is decreasing in speculator sentiment,* ζ_t*, whereas the token price is increasing (decreasing) in* ζ_t *when* $A_t^* - A_t$ *is sufficiently negative (positive).*

Figure 3 illustrates the equilibrium token price across the demand fundamental A for different values of speculator sentiment (the left panel), user optimism (the center panel), and token supply (the right panel). The center panel shows that the token price is increasing with user optimism, as formally established by Proposition 3. The left panel shows that the token price is also increasing with speculator sentiment, which holds, as established by Proposition 3, only when the demand fundamental is high. The difference across user optimism is more pronounced because user optimism increases user participation by raising users' expectations of the token's resale value, which in turn, raises the price today; speculator sentiment, in contrast, raises the token price but also crowds out user participation, which in turn, lowers the

Figure 3. An Illustration of the Token Price Across the Demand Fundamental for Different Values of Speculator Sentiment (Left Panel), User Optimism (Center Panel), and Token Supply (Right Panel)

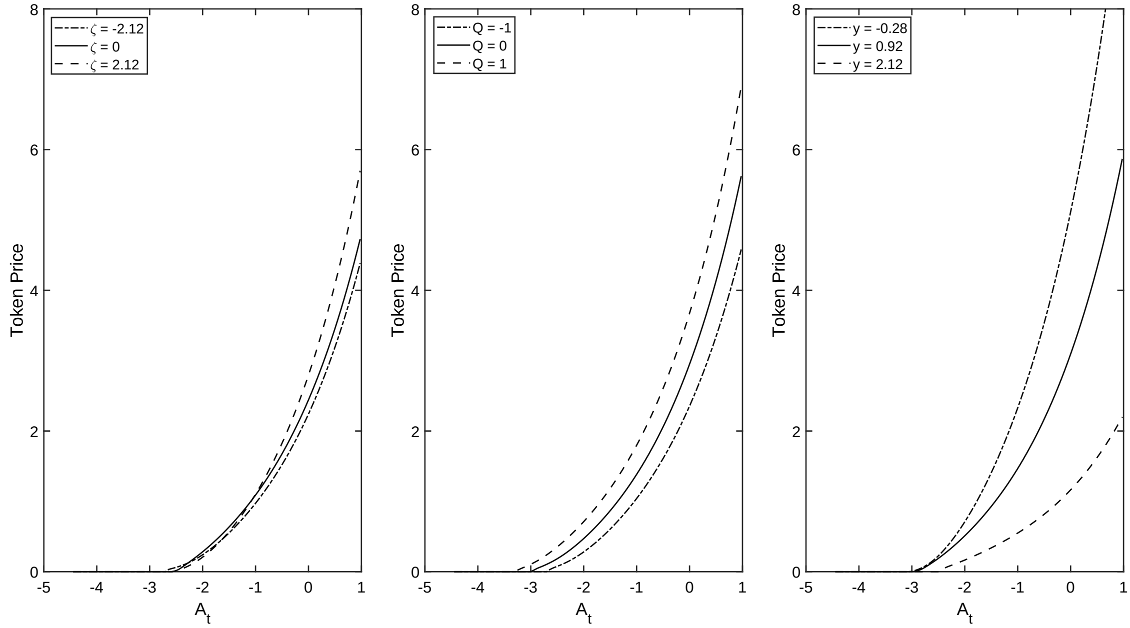

Note. The model parameters are given in Table 1, and the baseline values for the current state are $\zeta_t = 0$, $Q_t = 0$, and $y_t = 0.92$.

price, leading to a more muted overall effect on the token price. Finally, the right panel shows that the token price is decreasing in token supply because it lowers the expected retrade value of the token.

3.3. Life-Cycle Effects

Because our model is nonstationary with the token supply increasing deterministically over time, it has nuanced implications for how platform performance varies over the platform's life cycle. Central to understanding this pattern is the tension between the contemporaneous convenience yield and the capital gains in each user's total return from holding the token. Because users are risk neutral, the sum of the two pieces always equals the cost of carry plus the participation cost, $R + \kappa/P_t$, in equilibrium. Thus, when expected future token price appreciation is high, the current demand fundamental and convenience yield must be low.

The demand fundamental's expected growth rate $\mu - \frac{1}{2}\tau_A^{-1}$ and the token supply y_t are the two key model parameters that determine the expected token price. We illustrate these effects in Figure 4 for two values of μ. A platform with a higher μ will, on average, see A_t trend upward over time, sustaining a high expected token price, whereas a high y_t depresses token prices across all values of A_t from supply saturation. The tension between the convenience yield and the expected future token price also impacts the log token price volatility over time. When the demand fundamental growth rate μ is high, the expected token price remains higher over time. Because more of the token return for high μ platforms is from the capital gains part of the token return, the user base is less sensitive to instantaneous fluctuations in the demand fundamental, which drive the convenience yield. As such, we expect higher μ platforms to have lower token price volatility. In contrast, as the token supply increases, both the region of market breakdown and the importance of the convenience yield in token returns increase, leading to a more volatile token price.

3.4. Implications for Platform Design

Our analysis raises a key issue that the network effect endemic to utility token platforms can lead to fragility when a rigid token supply curve interacts with a demand curve that is subject to a network effect and nonneutrality of the token price. Consequently, policies that make the supply of tokens respond to speculative shocks, such as a state-contingent token issuance schedule, or that subsidize users, such as a state-contingent transaction fee rate (β in the model), can mitigate the risk of market breakdown. We now discuss these two possibilities.

By making the token supply curve more elastic and leaning against speculator sentiment, a state-contingent token issuance policy can potentially help ensure a nontrivial participation equilibrium on the platform. Such an issuance policy would ideally condition not only on the token price but also, on the nonfundamental component of token supply (i.e., speculator sentiment).[16] To see the potential role of a state-contingent token

Figure 4. An Illustration of the Unconditional Expected Log Token Price (Left Panel) and Log Price Volatility (Right Panel) over Time

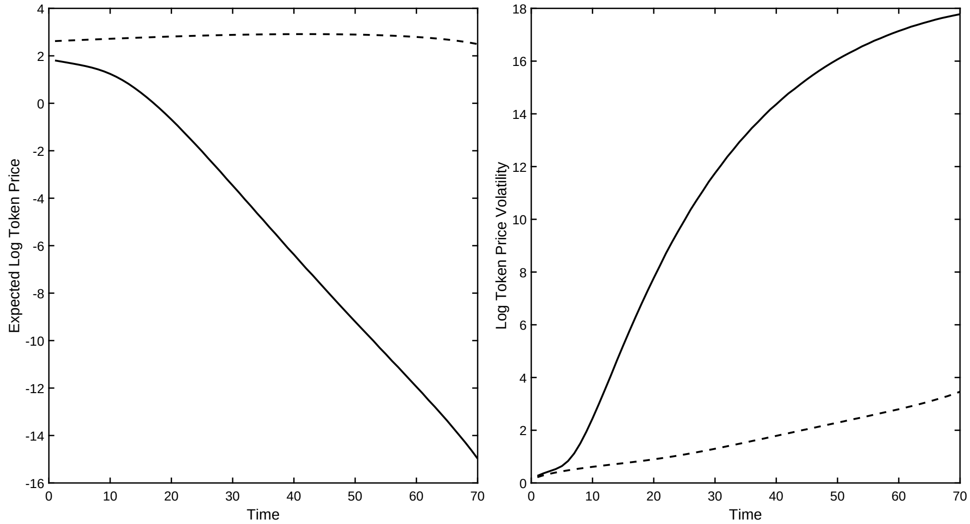

Notes. The model parameters are given in Table 1. The solid lines indicate the case in which the growth rate of the fundamental is $\mu = 0.02 - \frac{1}{2}\tau_A^{-1}$, and the dashed lines indicate the case in which $\mu = 0.10 - \frac{1}{2}\tau_A^{-1}$.

issuance policy, suppose now that the token issuance ι_t, defined in (4), is time varying and contingent on the state of the platform (A_t, y_t, Q_t, ζ_t). The following proposition establishes a condition on ι_t for an equilibrium with nonzero user participation to exist.

Proposition 4. *There exists a state-contingent issuance schedule, ι_t^*, as given by (A.6), such that an equilibrium with nonzero user participation exists at date t if $\iota_t \geq \iota_t^*$, which ensures the minimal supply elasticity with respect to ζ_t.*

We caution, however, that in practice, it may be difficult to implement a token issuance policy ($\iota_t \geq \iota_t^*$) that responds to conditions in the token's secondary market. Realistically, nonfundamental shocks to token prices, such as optimism and sentiment, are not directly observable, and conditioning on the token price and market outcomes, such as trading volume, may not be enough to disentangle the sources of token price fluctuations. Such contingency, if miscalibrated, may make the supply of tokens excessively volatile, which would be at variance with the proper functioning of the platform. It may also introduce unnecessary uncertainty into the revenue of validators, making them reluctant to provide validation services. It may further buffet the platform with nonfundamental fluctuations that impair performance and exacerbate the problem. Note that governments face similar issues in setting monetary policy: for instance, when deciding whether monetary policy should respond to stock or housing market fluctuations.

To illustrate how a miscalibrated token inflation schedule can harm platform performance, we recognize the necessary issuance schedule ι_t^* to avoid breakdown from Proposition 4 loads positively on speculator sentiment ζ_t when the expected token retrade value $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$ is below κ. When $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$ exceeds κ, there will be an equilibrium for any inflation rate. Suppose the platform designer miscalibrates the token inflation schedule and sets

$$\iota_t = -\zeta_t - y_{t-1} - p_t^* \text{ if } \mathbb{E}[P_{t+1}|\mathcal{I}_t] < \kappa,$$

where p_t^* is defined in (A.5) and depends on $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$. In this case, the token supply at date t is $\Phi(y_t + \lambda \log(RP_t) - \zeta_t) = \Phi(\lambda \log(RP_t) - p_t^* - 2\zeta_t)$, instead of $\Phi(\lambda \log(RP_t) - p_t^*)$ under ι_t^*. There are two effects of this miscalibration: one static and one dynamic. The static effect is immediate. As a result of the miscalibration, the token supply schedule now loads more negatively on speculator sentiment ζ_t (by a factor of –2) rather than positively as implied by the minimal inflation schedule. As such, the token supply not only fails to buffer the speculator sentiment shock but also, doubles its impact. More subtle is the dynamic effect. Because a more volatile inflation schedule affects future token prices, the expected retrade value $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$ at time t is also impacted by this miscalibration through p_t^*. More generally, this dynamic effect will depend on how the static miscalibration interacts with the token price and platform breakdown over time.

A state-contingent transaction fee rate may also potentially mitigate market breakdown by adjusting how much transaction surplus users need to give up to compensate the platform's validators. A reduction of β represents an effective subsidy for users to raise their demand curve. Similar to the case of a state-contingent token issuance policy, however, it may be difficult in practice to condition platform policy on unobservable speculator sentiment in secondary markets, and the alternative of conditioning on the token price can even be destabilizing. In addition, decentralized platforms lack an owner who has incentives to subsidize user participation, and consequently, the subsidy through a reduction of transaction fees is bounded from below by a transaction rate of $\beta = 0$, limiting its effectiveness. Because transaction fees are used to compensate validators, validators also have an incentive to maintain high fees on the platform.

4. Mining and Strategic Attacks

The risk of strategic attacks by validators is a central concern for cryptocurrency platforms. Attacks on Bitcoin Gold, ZenCash, Vertcoin, Monacoin, Ethereum Classic, and Verge (twice) have already led to losses of approximately \$18.6 million, \$550,000, \$50,000, \$90,000, \$1.1 million, and \$2.7 million, respectively. Such attacks include, for instance, 51% attacks under the proof of work protocol that lead to "double-spending" fraud and transaction failures through denials of service.[17] In this section, we demonstrate that strategic attacks occur when the platform fundamental is sufficiently weak. More importantly, the risk of such attacks in the future exacerbates the region of market breakdown by reducing the token's retrade value, which feeds back into the likelihood of a strategic attack. This adverse feedback loop is novel to decentralized cryptocurrency platforms.

To illustrate how consensus protocols can impact platform performance and stability, we consider a simple extension of our setting in this section that incorporates proof of work mining. We focus on the most ubiquitous type of attack on PoW blockchains, a 51%, attack, but our general insights will also be valid for other types of attacks, such as a selfish mining attack, and other consensus protocols, such as proof of stake, provided that the interests of validators may conflict with those of users.

We now assume that in each period, a new population of potential miners mines the token by providing accounting and custodial services using its underlying blockchain technology.[18] As in practice, there is free entry of miners onto the platform. All miners provide computing power to facilitate transactions among users, subject to a cost of setting up the required hardware and software to mine the token: $e^{-\xi_t}M_{j,t}$, where $M_{j,t} \in \{0,1\}$ is the miner's decision to mine and ξ_t measures the miner's mining efficiency by inversely parameterizing the miner's cost of mining.[19] This mining efficiency ξ_t is common to all miners and follows an AutoRegressive AR(1) process:

$$\xi_t = \xi_{t-1} + \tau_\xi^{-1/2}\varepsilon_t^\xi,$$

with $\varepsilon_t^\xi \sim$ i.i.d $\mathcal{N}(0,1)$. Validators are now miners who are compensated with the transaction fee βU_t, which is a fraction of the transaction surplus, and the seigniorage from token inflation, $(\Phi(y_{t-1}+\iota) - \Phi(y_{t-1}))P_t$. Consistent with many token platforms with PoW mining, miners also earn transaction fees because over time, the number of tokens created by inflation will diminish. It is thus necessary to shift the compensation toward fees. Miners have no use for tokens and sell them to users and speculators. If $N_{M,t}$ miners join the platform at date t, each miner earns $\frac{\beta U_t + (\Phi(y_{t-1}+\iota) - \Phi(y_{t-1}))P_t}{N_{M,t}} - e^{-\xi_t}$ in expected net gain.[20]

Suppose that when a strategic attack occurs, users lose a fraction $1-\gamma$ of their transaction surplus from failed transactions in the current period as a result of service delays and denials. The interruption of service also reduces transaction fees by a fraction $1-\gamma$. Furthermore, we assume that a strategic attack occurs whenever

$$(\Phi(y_t + \psi\iota) - \Phi(y_t))P_t + \frac{(\Phi(y_{t-1}+\iota) - \Phi(y_{t-1}))P_t + \beta\gamma U_t}{2} \geq \alpha N_{M,t}^2, \quad (8)$$

where $\alpha, \psi > 0$. On the left-hand side of this condition, the first term has the interpretation of fraudulent seignorage created by corrupt miners from double spending, and the second is a fraction γ of the mining fees, in the forms of legitimate seigniorage and transaction fees, earned from mining the attack. The right-hand side is the cost of attack, which is a convex function of the number of miners, reflecting that a larger pool of miners makes it increasingly costly for corrupt miners to acquire the necessary computing power for completing a 51% attack. In Appendix D, we provide a microfoundation for this strategic attack condition, although all that we require is that strategic attacks occur whenever the cost of mining is sufficiently high and the number of miners is sufficiently low.

Consider the incentives of miners to join the platform at date t. With rational expectations, miners choose whether to join, fully anticipating the possibility of a strategic attack. Miner j with the common mining efficiency ξ_t thus maximizes his expected gain:

$$\Pi_j = \max_{M_{j,t}} \left(\frac{(\Phi(y_{t-1}+\iota) - \Phi(y_{t-1}))P_t + (1-(1-\gamma)\chi_t)\beta U_t}{(1+\chi_t)N_{M,t}} - e^{-\xi_t} \right) M_{j,t}, \quad (9)$$

where $\chi_t \in \{0,1\}$ is the indicator for whether there is a strategic attack at date t. The $(1-(1-\gamma)\chi_t)$ factor reflects that the mining pool receives only γ of the total mining revenue from completing less than half of the blocks when a strategic attack occurs.

Note that relative to the equilibrium characterized in Section 3, the miners' common mining efficiency ξ_t becomes an additional state variable. The following proposition shows that strategic attacks occur when either A_t or ξ_t falls below a certain level.

Proposition 5. *The equilibrium has the following properties.*

1. *There exists a critical level* $\xi^a(A_t, y_t, Q_t, \zeta_t)$ *such that strategic attacks occur when* $\xi_t < \xi^a(A_t, y_t, Q_t, \zeta_t)$.
2. *There exists a critical level* $A^a(y_t, Q_t, \zeta_t, \xi_t)$, *which is decreasing in* ξ_t, *such that strategic attacks occur when* $A_t < A^a(y_t, Q_t, \zeta_t, \xi_t)$.
3. *Both an attack equilibrium and a no-attack equilibrium can exist as a result of the positive relationship between the benefits and costs of attacks.*

From Proposition 5, a strategic attack occurs when the mining fundamental and/or the user demand fundamental are sufficiently weak because in these situations, the number of miners is too small to deter a strategic attack. Although the impact of each strategic attack is transitory, the occurrence of strategic attacks is persistent because an attack will occur every period in which the platform is in the attack region. As attacks reduce the token price and thus, the incentives of miners to join the platform, it may be possible for both a no-attack equilibrium and an attack equilibrium to be self-fulfilling.

Figure 5 depicts the strategic attack boundary (left panel) and the platform breakdown boundary with and without mining (center panel) for $\tau_\xi = 10$, $\alpha = 0.8$, and $\psi = 3$. Miners choose to attack the cryptocurrency if the user fundamental A_t falls below the attack boundary A^a. This attack boundary is decreasing with the mining fundamental ξ_t, as formally derived in Proposition 5. Although each strategic attack does not lead to the failure of the platform, the expected losses induced by future attacks lead to a higher-threshold A^c for market breakdown. As such, the possibility of strategic attacks by miners also exacerbates platform fragility. This is reflected by the raised dashed line in the center panel.

As our analysis highlights, the PoW protocol introduces several novel features to cryptocurrency platforms. First, the anticipation of future attacks makes such a strategic attack easier to execute. An attack lowers the revenue each honest miner receives, which reduces the number of miners who join the platform and thus, lowers the cost of an attack. Interestingly, the decentralized consensus protocol exacerbates the problem by dispersing the revenue from mining over the whole population of miners. As a result, an honest miner captures only a fraction of the revenue that is recovered by increasing its own mining power to preempt attacks.[21] In this way, decentralized consensus averts the internalization of incentives to ensure the platform's security.

Second, the feedback effects from mining to the platform token's intrinsic value through service delays and denials are peculiar to the decentralized consensus protocol. Users are effectively also shareholders in the platform through the retradeability of the token. As such, delays and expectations of future delays have an important impact on the token price because they reduce user participation and consequently, demand for the token.

Figure 5. An Illustration of the Strategic Attack Boundary (Left Panel), Market Breakdown Boundary (Center Panel), and Token Price (Right Panel) with Respect to Mining Fundamental ξ_t

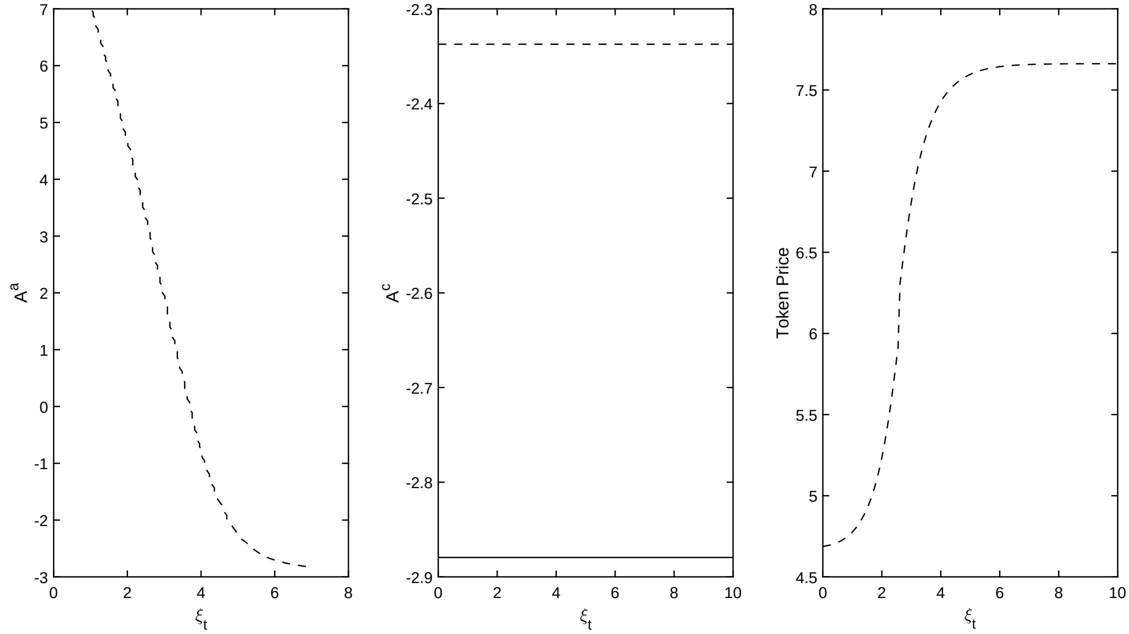

Notes. The market breakdown boundary without mining (solid line) is for comparison. User optimism is turned off ($\tau_Q = 0$) in this illustration, and the lost surplus is set to $\gamma = \frac{1}{2}$. In addition, platform participants always coordinate on the no-attack equilibrium when it exists. Baseline values are $A_t = 0.98$, $\zeta_t = 0$, and $y_t = 0.92$.

These two features contribute to a rich dynamic adverse feedback loop between strategic attacks and market breakdown. The anticipation of future strategic attacks lowers the expected retrade value of the token, which in turn, reduces users' incentives to join the platform and exacerbates both the problem of strategic attacks and market breakdown.

Finally, from Figure 5 (right panel), we see a nonlinear relation between the mining fundamental and token price. When the mining fundamental is far away from the strategic attack boundary, an incremental change in the efficiency of mining has a limited impact on the token price because the probability of an attack is small. When the mining fundamental is close to the strategic attack boundary, however, a small change in the efficiency of mining can have a substantial impact on the token price, which in turn, leads to a substantial impact on the platform's stability.

Our insights about the adverse dynamic feedback loop associated with strategic attacks are also applicable to other types of attacks beyond a 51% attack under the PoW consensus protocol. Another type of PoW attack is a selfish mining attack, in which a miner secretly validates blocks until she can broadcast it as the longest chain. Like a 51% attack, this attack is more difficult to execute when there are more miners, and each miner has less probability of winning a block. Consequently, our analysis would also apply to this type of attack. Similarly, under the proof of stake protocol, the most prevalent type of strategic attack is a Sybil attack. Under a Sybil attack, a rogue validator can acquire 51% of all staked tokens and create false validator nodes to manipulate consensus on the blockchain to engage in a distributed denial of service or "double-spending" attack. Like the 51% attack under PoW, such a strategic attack is more difficult when the revenue from validating transactions is higher and there are a lot of tokens staked to compete for this revenue. Furthermore, a Sybil attack is also harder when the token price is high because acquiring a 51% stake size is more expensive. As in our analysis, an attack is more likely to occur when the platform is weak and user participation is low. By similar logic to our 51% attack analysis, the anticipation of a Sybil attack also reduces user incentives to join the platform, which increases the region of market breakdown and strategic attacks. Consequently, our analysis of strategic attacks applies more generally to vulnerabilities of consensus protocols.

5. Empirical Implications

In this section, we discuss several empirical implications of our conceptual framework for cryptocurrency returns. It is important to note that our model implications are specific to tokenized platforms that do not adjust the number of tokens required for their services and would not apply, for instance, to (alt-)coins, stablecoins, or NFTs. Cryptocurrency returns in our framework have three components: a convenience yield of the marginal user, which acts like a dividend; a capital gain from the token price appreciation; and an embedded discount in the token price to compensate users for their participation cost. By the marginal user's equilibrium condition in (7), these three components satisfy the following relationship:

$$R = \frac{(1-\beta)U_t^*}{P_t} + \frac{\mathbb{E}[P_{t+1}|\mathcal{I}_t]}{P_t} - \frac{\kappa}{P_t}.$$

In contrast to fiat currencies, the expected capital gain can be quite positive, despite token inflation, and substantial, which has attracted many speculators to the nascent asset class. In addition and novel to cryptocurrencies, the convenience yield is created by shareholders acting in their dual capacity as users of the platform, which gives rise to a feedback mechanism from the cryptocurrency return to user participation. As the platform matures and participation increases, the cryptocurrency return transitions from being driven more by the capital gain component to more by the convenience yield.[22]

The empirical literature is mostly focused on the capital gain component of the cryptocurrency return, as it is directly measurable by the econometrician. In equilibrium, the expected excess capital gain can be expressed as

$$\frac{\mathbb{E}[P_{t+1}|\mathcal{I}_t]}{P_t} - R = \frac{\kappa}{P_t} - \frac{(1-\beta)U_t^*}{P_t}. \tag{10}$$

Consistent with the empirical findings of Hu et al. (2019) and Liu and Tsyvinski (2021), the expected excess capital gain in our setting does not exhibit conventional risk premia. The capital gain may still exhibit predictability through the underlying state variables that explain the convenience yield. These state variables are the demand fundamental, user optimism, speculator sentiment, and token supply. Liu and Tsyvinski (2021), for instance, show that investor attention, measured either with Google searches or with Twitter post counts for "Bitcoin," predicts future cryptocurrency returns, with positive (negative) attention, as measured by keywords, positively (negatively) predicting future weekly returns.[23] Liu and Tsyvinski (2021) also find that investor sentiment, measured as either the log ratio between the number of positive and negative phrases of cryptocurrencies in Google searches or the ratio of trading volume to return volatility, predicts future cryptocurrency returns. Such nonfundamental shocks to token prices, represented by user optimism and speculator sentiment in our model, can also explain reversals in cryptocurrency returns, consistent with the evidence of a "value" factor in Cong et al. (2022a).

Our model also suggests that the participation cost borne by users, which is not directly observed by the econometrician, is an additional channel of return predictability. As this cost effect is inversely related to the token price and consequently, to market capitalization, our model predicts a size effect in the capital gain of cryptocurrencies. This prediction is consistent with Liu et al. (2022), who find a size factor in the cross-section of cryptocurrency returns, with size measured as market capitalization, price, or maximum price.

In addition, the persistence of the two return components $\frac{\kappa}{P_t}$ and $\frac{(1-\beta)U_t^*}{P_t}$ in (10) can lead to a positive autocorrelation in the capital gain:

$$Cov\left(\frac{P_{t+2}}{P_{t+1}}, \frac{P_{t+1}}{P_t}\middle|\mathcal{I}_{t-1}\right) = Cov\left(\frac{\kappa}{P_{t+1}} - \frac{(1-\beta)U_{t+1}^*}{P_{t+1}}, \frac{\kappa}{P_t} - \frac{(1-\beta)U_t^*}{P_t}\middle|\mathcal{I}_{t-1}\right) > 0$$

because the innovations $\frac{P_{t+1}-\mathbb{E}[P_{t+1}|\mathcal{I}_t]}{P_t}$ and $\frac{P_{t+2}-\mathbb{E}[P_{t+2}|\mathcal{I}_{t+1}]}{P_{t+1}}$ are uncorrelated with rational expectations. This positive autocorrelation implies momentum, as empirically documented by Liu and Tsyvinski (2021) in the prices of cryptocurrencies. Furthermore, the momentum effect in our model is independent of investor attention and sentiment, which is also consistent with Liu and Tsyvinski (2021), who find time series momentum over one- to eight-week horizons that is not subsumed by their measures of attention or sentiment.

Our model also highlights the importance of network effects in utility token pricing. Shams (2020) shows that return comovement arising from overlapping exposures to demand shocks is significantly stronger among "high community-based" cryptocurrencies, whereas Schwenkler and Zheng (2021) find evidence of comovement among peer cryptocurrencies based on news reactions. Cong et al. (2022a) also provide evidence of a network factor that prices the cross-section of cryptocurrencies.

Finally, our extension with mining suggests that the capital gain from a cryptocurrency has a nonlinear relation with the marginal cost of mining. When the cost of mining is low relative to the strategic attack threshold, small changes in it have a muted impact on the capital gain, as the potential loss from strategic attacks, which can be viewed as an extended form of the participation cost in (10), is small. As the mining cost increases toward the strategic attack boundary, however, incremental changes become more relevant. Our model, therefore, predicts that measures of mining costs should have more predictive power for the capital gain when there is a nontrivial chance of strategic attacks, such as when the hash rate or the number of miners is low.

6. Conclusion

This paper develops a model to analyze the price dynamics and stability of cryptocurrencies. In our model, a cryptocurrency comprises both an asset and a membership in a platform developed to facilitate transactions of certain goods or services. As a result of the strong network effect among users to participate on the platform and the rigidity induced by market clearing with token speculators, the market can break down so that there is only an equilibrium with zero user participation. In such a setting, token retradeability plays an important role in harnessing the optimism of users to mitigate this instability. In contrast, it can exacerbate such fragility if it attracts speculators whose enthusiasm crowds out users. As a result of token inflation, this novel benefit of token retradeability fades as the platform matures and the token price becomes driven more by the current platform fundamental. We further illustrate how consensus validation protocols can exacerbate the platform's instability through strategic attacks on the blockchain. The potential for strategic attacks feeds back into the incentives both of miners to mine and of users to join the platform, which makes such attacks more likely. Our model also provides several implications for cryptocurrency price changes that are broadly consistent with recent empirical evidence.

Acknowledgments

The authors thank An Yan for a comment that led to this paper. The authors also thank Will Cong, Haoxiang Zhu, and Aleh Tsyvinski as well as seminar participants at Instituto Tecnológico Autónomo de México, the National Bureau of Economic Research Asset Pricing Meeting, the National Bureau of Economic Research Summer Institute, Tsinghua, University of British Columbia, University of North Carolina, and Yale for helpful comments. The authors particularly thank Agostino Capponi (editor), an associate editor, and three referees for their highly constructive comments and suggestions.

Appendix A. Proofs of Propositions

Proof of Proposition 1. We first examine the decision of a user to purchase the token. We first recognize that each user's expectation about P_{t+1}, $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$, depends on each user's expectation of A_{t+1}. By the Bayes rule, it is straightforward to conclude that the conditional posterior of users about A_{t+1} after observing A_t and Q_t is Gaussian $A_{t+1}|\mathcal{I}_t \sim \mathcal{N}(\hat{A}_{t+1}, \hat{\tau}_A^{-1})$, where the conditional estimate and precision satisfy

$$\hat{A}_{t+1} = A_t + \mu + \frac{\tau_Q}{\tau_\varepsilon + \tau_Q} Q_t,$$
$$\hat{\tau}_A = \tau_\varepsilon + \tau_Q.$$

We define τ as the stopping time, at which the platform fails as a result of the breakdown of the token market. We shall derive the conditions that determine τ later. Conditional on $t < \tau$, the expected utility of user i, who chooses to purchase the token at t, from transacting with another user is

$$\mathbb{E}[U_{i,t}|\mathcal{I}_t, \tau > t, A_{it}, \text{matching with user } j] = e^{(1-\eta_c)A_{i,t}}\mathbb{E}[e^{\eta_c A_{j,t}}|\mathcal{I}_t],$$

which is monotonically increasing with the user's own endowment $A_{i,t}$. Note that $\mathbb{E}[e^{\eta_c A_{j,t}}|\mathcal{I}_t]$ is independent of $A_{i,t}$

but dependent on the strategies used by other users. It then follows that user i will follow a cutoff strategy that is monotonic in its own type $A_{i,t}$.

Suppose that every user uses a cutoff strategy with a threshold of A_t^*. Then, the expected utility of user i is

$$\mathbb{E}[U_{i,t}|\mathcal{I}_t,\tau>t]=e^{(1-\eta_c)A_{i,t}+\eta_c A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2}+\frac{A_t-A_t^*}{\tau_\varepsilon^{-1/2}}\right)\mathbf{1}_{\{\tau>t\}}$$

because losing a transaction is independent of the identities of the two transacting parties.

To determine the equilibrium threshold, consider a user with the critical endowment $A_{it}=A_t^*$. As this marginal user must be indifferent to his purchase choice, it follows that

$$\mathbb{E}[(1-\beta)U_{i,t}+P_{t+1}|\mathcal{I}_t,A_{it}=A_t^*]=RP_t+\kappa,$$

which is equivalent to

$$(1-\beta)e^{(1-\eta_c)A_{i,t}+\eta_c A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2}+\frac{A_t-A_t^*}{\tau_\varepsilon^{-1/2}}\right)\mathbf{1}_{\{\tau>t\}}+\mathbb{E}[P_{t+1}|\mathcal{I}_t]=RP_t+\kappa, \tag{A.1}$$

with $A_{i,t}=A_t^*$. Fixing the critical value A_t^*, the expected token price $E[P_{t+1}|\mathcal{I}_t]$, and the price P_t, we see that the LHS of Equation (A.1) is monotonically increasing in $A_{i,t}$ because $1-\eta_c>0$. This confirms the optimality of the cutoff strategy that users with $A_{i,t}\ge A_t^*$ acquire the token to join the platform and that users with $A_{i,t}<A_t^*$ do not. Because $A_{i,t}=A_t+\varepsilon_{i,t}$ it then follows that a fraction $\Phi(-\sqrt{\tau_\varepsilon}(A_t^*-A_t))$ of the users enters the platform and that a fraction $\Phi(\sqrt{\tau_\varepsilon}(A_t^*-A_t))$ chooses not to. As one can see, it is the integral over the idiosyncratic endowment of users ε_i that determines the fraction of potential users on the platform.

By substituting P_t from Equation (5) into Equation (A.1), we obtain an equation to determine the equilibrium cutoff $A_t^*=A_t^*(\mathcal{I}_t)$:

$$(1-\beta)e^{A_t+(1-\eta_c)(A_t^*-A_t)+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2}+\frac{A_t-A_t^*}{\tau_\varepsilon^{-1/2}}\right)\mathbf{1}_{\{\tau>t\}}+\mathbb{E}[P_{t+1}\,|\,\mathcal{I}_t]=e^{\frac{\sqrt{\tau_\varepsilon}}{\lambda}(A_t-A_t^*)-\frac{1}{\lambda}y_t+\frac{1}{\lambda}\zeta_t}+\kappa. \tag{A.2}$$

Define $z_t=\sqrt{\tau_\varepsilon}(A_t^*-A_t)$, which determines the population that buys the token. We can rewrite Equation (A.2) as

$$(1-\beta)e^{\left[(1-\eta_c)\tau_\varepsilon^{-1/2}+\frac{1}{\lambda}\right]z_t+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi(\eta_c\tau_\varepsilon^{-1/2}-z_t)\mathbf{1}_{\{\tau>t\}}+e^{\frac{1}{\lambda}z_t}(\mathbb{E}[P_{t+1}\,|\,\mathcal{I}_t]-\kappa)=e^{-\frac{1}{\lambda}y_t+\frac{1}{\lambda}\zeta_t}. \tag{A.3}$$

Note the first term in the LHS of Equation (A.3) has a humped shape with respect to z_t, and the second term is an exponential function of z_t with a coefficient that may be either positive or negative. As the RHS of Equation (A.3) is constant with respect to z_t, this equation may have zero, one, two, or three roots.

- If $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa\le 0$, the LHS has a humped shape with a maximum at $\bar{z}$, and it may intersect with the RHS at zero or two points.
  1. If $LHS(\bar{z})<RHS$, then Equation (A.3) has no root.
  2. If $LHS(\bar{z})>RHS$, then Equation (A.3) has two roots.
- If $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa>0$, the LHS is nonmonotonic with $LHS(-\infty)=0$, $LHS(\infty)=\infty$, and one local maximum $\ddot{z}$ and one local minimum $\dot{z}$ in $(-\infty,\infty)$, and it may intersect the RHS at one or three points.
  3. If $RHS<LHS(\dot{z})$ or if $RHS>LHS(\ddot{z})$, then Equation (A.3) has one root.
  4. If $LHS(\dot{z})<RHS<LHS(\ddot{z})$, then Equation (A.3) has three roots.

In the first scenario outlined, there is only an equilibrium with trivial user participation, and the token market breaks down. Note that A_t shifts up and down the left-hand side of Equation (A.3). Thus, Equation (A.3) has no root when A_t is sufficiently small. For this situation to occur, the speculative motive, $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa$, must be nonpositive; otherwise, Equation (A.3) has one or three roots. This condition is also satisfied when A_t is sufficiently small because $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$ is increasing with A_t. Thus, the token market breaks down when A_t falls below a certain critical level, which we denote as $A^c(y_t,Q_t,\zeta_t)$. Thus, the stopping time τ of the platform's disbandment is

$$\tau=\{\inf t:\ A_t<A^c(y_t,Q_t,\zeta_t)\}.$$

Finally, note that because the only difference among users is the value of their transaction benefit $\mathbb{E}[U_{i,t}|\mathcal{I}_t,\tau>t]$, which is monotonically increasing in $A_{i,t}$ regardless of the mass of users who join the platform, it follows that, regardless of the strategies of other users, it is always optimal for each user i to follow a cutoff strategy. □

Proof of Proposition 2. The first part of the proposition follows from the derivation of Proposition 1 and the definition of A^c. This proof characterizes the determinants of the fundamental critical level A^c.

With regard to speculator sentiment, notice from Equation (A.3) that, when $\mathbb{E}[P_{t+1}|\mathcal{I}_t]-\kappa$ is nonpositive, there is a critical value of speculator sentiment $\zeta^c(A_t,y_t,Q_t)$:

$$\zeta_t^c=\lambda\log\left\{\sup_{z_t}\left\{(1-\beta)e^{\left[(1-\eta_c)\tau_\varepsilon^{-1/2}+\frac{1}{\lambda}\right]z_t+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi(\eta_c\tau_\varepsilon^{-1/2}-z_t)+e^{\frac{1}{\lambda}z_t}(\mathbb{E}[P_{t+1}\,|\,\mathcal{I}_t]-\kappa)\right\}\right\}+y_t,$$

such that nontrivial equilibrium exists if $\zeta_t\ge\zeta^c(A_t,y_t,Q_t)$, with the convention that $\zeta_t^c=-\infty$ if the argument in the log is negative.

It is straightforward to see that, in the high-price (low-cutoff) equilibrium, the implicit function theorem implies that $\frac{dz_t}{d\zeta_t}>0$. Because the user participation is $\Phi(-z_t)$, it follows that an increase in ζ_t exacerbates the market breakdown region by lowering user participation. Because ζ_t is i.i.d., there is only this static impact of an increase in speculator sentiment on the equilibrium cutoff. As such, by lowering user participation, it shifts up $A^c(y_t,Q_t,\zeta_t)$ for any given pair of $\{y_t,Q_t\}$.

We next consider how user optimism Q_t impacts the market breakdown region. Because user optimism Q_t raises each user's estimate of the resale value of the token at date $t+1$, it raises user participation and the token price at date t. Because Q_t is i.i.d., this is the only impact of an increase in user optimism. As such, it shifts down the market breakdown threshold, $A^c(y_t,Q_t,\zeta_t)$, for any given pair of $\{y_t,\zeta_t\}$.

Similarly, an increase in the user participation cost, κ, deters user participation at all dates and therefore, exacerbates the market breakdown by both increasing the cost today and lowering the expected retrade value of the token tomorrow through the reduced participation in the future. As such, it also shifts up $A^c(y_t, Q_t, \zeta_t)$. □

Proof of Proposition 3. We first establish that the map from the demand fundamental A_t to the equilibrium user cutoff for joining the platform is monotone when the highest-price equilibrium is always played.[24]

Suppose that the token price at date $t+1$, P_{t+1}, is increasing in A_t for all (y_t, Q_t, ζ_t) triples in the high-price equilibrium. Then, because A_t follows a random walk, its cumulative distribution function satisfies the Feller property, and the conditional expectation operator preserves this relation:

$$\frac{\partial \mathbb{E}[P_{t+1}|\mathcal{I}_t]}{\partial A_t} = \mathbb{E}\left[\frac{\partial P(A_t+\mu+\varepsilon_{t+1}, y_{t+1}, Q_{t+1}, \zeta_{t+1})}{\partial A_t}|\mathcal{I}_t\right] > 0,$$

where the expectation is take over ε_{t+1}. Consequently, $\mathbb{E}[P_{t+1}|\mathcal{I}_t]$ is increasing in A_t. Then, we can rewrite Equation (A.3) as the function G_t:

$$G_t = (1-\beta)e^{\left[(1-\eta_c)\tau_\varepsilon^{-1/2}+\frac{1}{\lambda}\right]z_t + A_t + \frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi(\eta_c\tau_\varepsilon^{-1/2} - z_t)\mathbf{1}_{\{\tau>t\}} + e^{\frac{1}{\lambda}z_t}(\mathbb{E}[P_{t+1}\,|\,\mathcal{I}_t] - \kappa) - e^{-\frac{1}{\lambda}y_t + \frac{1}{\lambda}\zeta_t} \equiv 0. \tag{A.4}$$

Assuming the existence of an equilibrium with nontrivial user participation, applying the implicit function theorem to G_t, one has that

$$\frac{\partial z_t}{\partial A_t} = -\frac{\partial G_t/\partial A_t}{\partial G_t/\partial z_t},$$

where

$$\frac{\partial G_t}{\partial A_t} = (1-\beta)e^{\left[(1-\eta_c)\tau_\varepsilon^{-1/2}+\frac{1}{\lambda}\right]z_t + A_t + \frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2} - z_t\right) + e^{\frac{1}{\lambda}z_t}\frac{\partial \mathbb{E}[P_{t+1}|\mathcal{I}_t]}{\partial A_t} > 0.$$

In the high-price equilibrium, the RHS of Equation (A.3) intersects the hump-shaped curve of the LHS in z_t on the left side of the hump, and consequently, $\frac{\partial G_t}{\partial z_t} \geq 0$.[25] It then follows that, in the high-price equilibrium, $\frac{\partial z_t}{\partial A_t} < 0$. Therefore, user participation $\Phi(-z_t)$ is increasing in A_t.

Furthermore, because $P_t = e^{-\frac{1}{\lambda}z_t - \frac{1}{\lambda}y_t + \frac{1}{\lambda}\zeta_t}$, it follows that

$$\frac{\partial P_t}{\partial A_t} = -\frac{P_t}{\lambda}\frac{\partial z_t}{\partial A_t} > 0.$$

Consequently, P_t is increasing in A_t in the high-price equilibrium. Because the choices of t and $t+1$ are arbitrary, P_t is increasing in A_t generically if the high-price equilibrium is played at each date.

Finally, because user optimism Q_t enters into the user's problem by raising the expected resale token price, it raises user participation and the token price. In contrast, speculator sentiment ζ_t lowers user participation by leading to nonfundamental upward pressure on the token price. Because it also lowers user participation, the overall impact on the token price is ambiguous. To see this, we rewrite Equation (A.4) as

$$H_t \equiv (1-\beta)e^{(1-\eta_c)\tau_\varepsilon^{-1/2}(\tilde{z}_t+\zeta_t) + A_t + \frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \tilde{z}_t - \zeta_t\right)\mathbf{1}_{\{\tau>t\}} - e^{-\frac{1}{\lambda}\tilde{z}_t - \frac{1}{\lambda}y_t} + \mathbb{E}[P_{t+1}\,|\,\mathcal{I}_t] - \kappa = 0,$$

where the change of variables $\tilde{z}$ now absorbs speculator sentiment, so that the price is $P_t = e^{-\frac{1}{\lambda}\tilde{z}_t - \frac{1}{\lambda}y_t}$. Because speculator sentiment is i.i.d. and the equilibrium is Markovian in the state space (A_t, y_t, Q_t, ζ_t), the retrade value of the token is unaffected by changes in sentiment today. It is straightforward by the implicit function theorem to the equation that

$$\frac{\partial \tilde{z}_t}{\partial \zeta_t} = -\frac{dH_t/d\zeta_t}{dH_t/d\tilde{z}_t}.$$

Because $\tilde{z}$ enters H_t symmetrically as z does in Equation (A.3), $dH_t/d\tilde{z}_t > 0$ in the high-price equilibrium. In contrast, $dH_t/d\zeta_t$ is

$$dH_t/d\zeta_t \propto (1-\eta_c)\tau_\varepsilon^{-1/2} - \frac{\phi\left(\eta_c\tau_\varepsilon^{-1/2} - \tilde{z}_t - \zeta_t\right)}{\Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \tilde{z}_t - \zeta_t\right)} = (1-\eta_c)\tau_\varepsilon^{-1/2} - \frac{\phi\left(\eta_c\tau_\varepsilon^{-1/2} - z_t\right)}{\Phi\left(\eta_c\tau_\varepsilon^{-1/2} - z_t\right)}.$$

Consequently, if z_t is sufficiently small, then $dH_t/d\zeta_t > 0$, whereas if z_t is sufficiently large, then $dH_t/d\zeta_t < 0$. Because $\frac{\partial P_t}{\partial \zeta_t} = -\frac{1}{\lambda}P_t\frac{\partial \tilde{z}_t}{\partial \zeta_t}$, it follows that $\frac{\partial P_t}{\partial \zeta_t} > 0$ for z_t sufficiently small and that $\frac{\partial P_t}{\partial \zeta_t} < 0$ for z_t sufficiently large. Because $z_t = A_t^* - A_t$, the result follows. □

Proof of Proposition 4. From the proof of Proposition 1, the left-hand side of Equation (7) is a hump-shaped curve in the number of tokens demanded $N_t = \Phi(\sqrt{\tau_\varepsilon}(A_t^* - A_t))$, whereas the right-hand side (the token price) can be expressed as an exponential function of $\frac{1}{\lambda}(\zeta_t - \Phi^{-1}(N_t) - y_{t-1} - \iota_t)$ for an inflation rate ι_t. Notice that a higher ι_t lowers the exponential curve, whereas the fundamental cause of market breakdown (only a trivial participation solution) is that the exponential curve is always above the hump-shaped curve for $N_t > 0$. Consequently, there exists a minimum i_t^* such that the exponential curve just touches the peak of the log of the hump-shaped curve. Equating the log of the hump-shaped demand curve with the token price, we can define p^* as

$$p^* = \begin{cases} \max\limits_n \lambda \log\Big[(1-\beta)e^{\frac{1-\eta_c}{\sqrt{\tau_\varepsilon}}\Phi^{-1}(n) + A_t + \frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}} & \\ \quad \Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \Phi^{-1}(n)\right)\mathbf{1}_{\{\tau>t\}} & \text{if } \mathbb{E}[P_{t+1}|\mathcal{I}_t] \leq \kappa \\ \quad +\mathbb{E}[P_{t+1}|\mathcal{I}_t] - \kappa\Big] + \Phi^{-1}(n) & \\ \infty & \text{if } \mathbb{E}[P_{t+1}|\mathcal{I}_t] > \kappa. \end{cases} \tag{A.5}$$

This minimal state-contingent inflation rate then takes the form

$$\iota_t^* = \zeta_t - y_{t-1} - p_t^*. \tag{A.6}$$

Consequently, for $\iota_t \geq \iota_t^*$, an equilibrium with nontrivial participation exists. □

Proof of Proposition 5. From the miner optimization Problem (9), it is straightforward to see that, with free entry, miners must be indifferent to participating on the platform. Consequently, the number of potential miners who choose to mine is given by

$$N_{M,t} = \frac{(\Phi(y_{t-1}+\iota) - \Phi(y_{t-1}))P_t + \beta(1-(1-\gamma)\chi_t)U_t}{1+\chi_t}e^{\xi_t}.$$

Substituting the optimal number of miners, $N_{M,t}$, from (9) into the attack condition given in (8) conjecturing an attack,

$\chi_t = 1$, we can define

$$f(y_t, P_t, \mathbb{E}[U_t \mid \mathcal{I}_t]) = \left(\Phi(y_t + \psi\iota) - \frac{1}{2}\Phi(y_t) - \frac{1}{2}\Phi(y_t - \iota)\right)P_t + \frac{1}{2}\beta\gamma\mathbb{E}[U_t \mid \mathcal{I}_t] - \frac{\alpha e^{2\xi_t}}{4}\left(\left(\Phi(y_t) - \Phi(y_t - \iota)\right)P_t + \frac{\beta}{2}U_t\right)^2.$$

There is an attack whenever $f(y_t, P_t, U_t) > 0$.[26] It is clear because ξ enters only through the quadratic term that there exists a threshold $\xi^c(A_t, Q_t, \zeta_t)$ such that

$$\{\chi_t = 1 : \xi_t < \xi^a(A_t, Q_t, \zeta_t)\},$$

where

$$\xi^c(A_t, y_t, Q_t, \zeta_t) = \frac{1}{2}\log\frac{\left(\Phi(y_t + \psi\iota) - \frac{1}{2}\Phi(y_t) - \frac{1}{2}\Phi(y_t - \iota)\right)P_t + \frac{1}{2}\beta\gamma U_t}{\frac{\alpha}{4}((\Phi(y_t) - \Phi(y_t - \iota))P_t + \frac{\beta}{2}\mathbb{E}[U_t|\mathcal{I}_t])^2}.$$

Assume now that $\mathbb{E}[U_t|\mathcal{I}_t]$ and P_t are (weakly) increasing in A_t whenever P_t is positive, and we define $P_t = 0$ whenever a market equilibrium does not exist. Define

$$x_t = \frac{(\Phi(y_t) - \Phi(y_t - \iota))P_t + \frac{\beta}{2}U_t}{2},$$

and rewrite $f(y_t, P_t, \mathbb{E}[U_t|\mathcal{I}_t])$ as

$$f(y_t, P_t, x_t) = \left(\Phi(y_t + \psi\iota) - \Phi(y_t)\right)P_t + x_t - \alpha e^{2\xi_t}x_t^2.$$

Notice that $f(y_t, P_t, x_t)$ is concave in x_t, increasing for $x_t < \frac{1}{2\alpha e^{2\xi_t}}$ from zero to $\frac{1}{4\alpha e^{2\xi_t}}$, and then, decreasing to $-\infty$ for $x_t > \frac{1}{2\alpha e^{2\xi_t}}$. It has two roots at $x_t \in \left\{0, \frac{1}{\alpha e^{2\xi_t}}\right\}$.

It then follows that a strategic attack occurs whenever $x_t \le \frac{1}{\alpha e^{2\xi_t}}$, or when A_t is sufficiently small. This occurs because U_t and P_t are (weakly) increasing in A_t and because U_t and P_t converge to zero as $A_t \to -\infty$, as there is no benefit to any (positive measure of) users joining the platform. Consequently, because P_t and U_t are (weakly) increasing in A_t, it follows there is a connected set $\underline{\mathcal{A}}_t = \{A_t : A_t < A^a(y_t, Q_t, \zeta_t; \xi_t)\}$, where $A^a(y_t, Q_t, \zeta_t; \xi_t) = \inf_{A_t}\{f(y_t, P_t, x_t) = 0\}$, such that $\chi_t = 1$ when $A_t < \underline{\mathcal{A}}_t$.

In contrast, when A_t is sufficiently large, it must be the case that $\lim_{A_t\to\infty} f(y_t, P_t, x_t) < 0$ because the highest-order terms in P_t and U_t are quadratic through $-x_t^2$. Consequently, there is a connected set $\overline{\mathcal{A}}_t = \{A_t : A_t > \overline{A}^a(y_t, Q_t, \zeta_t; \xi_t)\}$, where $\underline{A}^a(y_t, Q_t, \zeta_t; \xi_t) = \sup_{A_t}\{f(y_t, P_t, x_t) = 0\}$, such that $\chi_t = 0$ when $A_t > \overline{\mathcal{A}}_t$.

Consequently, it follows that there is a strategic attack when $A_t \in \underline{\mathcal{A}}_t$ and no attack when $A_t \in \overline{\mathcal{A}}_t$. What remains is to determine if $\underline{\mathcal{A}}_t \cup \overline{\mathcal{A}}_t = \mathbb{R}$ or if there are more strategic attack regions for some $A_t > \underline{\mathcal{A}}_t$. Notice now that $f(y_t, P_t, x_t)$ is a quadratic function of x_t and by Descartes' rule of signs, has at most one positive root, which we know must exist by these arguments. Consequently, $f(y_t, P_t, x_t)$ has one zero when, substituting for x_t,

$$\frac{\beta}{2}\mathbb{E}[U_t|\mathcal{I}_t] = \frac{1}{\alpha e^{2\xi_t}} + \sqrt{\left(\frac{1}{\alpha e^{2\xi_t}}\right)^2 + 4\frac{\Phi(y_{t-1} + \psi\iota) - \Phi(y_t)}{\alpha e^{2\xi_t}}P_t} - (\Phi(y_t) - \Phi(y_t - \iota))P_t. \tag{A.7}$$

Therefore, it must be the case that $A^a(y_t, Q_t, \zeta_t; \xi_t) = \underline{A}^a(y_t, Q_t, \zeta_t; \xi_t)$, and therefore, the strategic attack region can be characterized as

$$\chi_t = \begin{cases} 1 & \xi_t < \xi^a(A_t, y_t, Q_t, \zeta_t) \\ 0 & \xi_t \ge \xi^a(A_t, y_t, Q_t, \zeta_t) \end{cases}$$

or alternatively,

$$\chi_t = \begin{cases} 1 & A_t < A^a(y_t, Q_t, \zeta_t; \xi_t) \\ 0 & A_t \ge A^a(y_t, Q_t, \zeta_t; \xi_t). \end{cases}$$

In addition, we recognize from (A.7) that because a higher ξ_t lowers the critical $\frac{\beta}{2}U_t$, all else equal, it follows that $A^a(y_t, Q_t, \zeta_t; \xi_t)$ is decreasing in ξ_t.

One may be concerned that no mining equilibrium may exist if, conditional on no attack, miners want to attack the blockchain, whereas conditional on an attack, no miner ex post wants to attack the blockchain. This does not occur because the (convex) cost of attacks from fewer miners falls faster than the benefit from the attack from lower revenue. To see this, notice that the only endogenous object determined by users is A_t^*, and a strategic attack raises A_t^*, lowering prices and transaction fees, by reducing the benefit of joining the platform for all users. This is equivalent to a fall in A_t to some $\tilde{A}_t$. Because if an attack that would occur at A_t would also occur at $A_t' < A_t$, by these arguments, it follows that if a strategic attack would occur when users and miners do not anticipate an attack, it would also occur if it is anticipated. Consequently, such a strategic attack inconsistency issue does not arise.

Furthermore, although there cannot be an inconsistency in the attack decision on the platform, there can be self-fulfilling prophecies, in which both the no-attack and attack equilibria can be sustained. This arises because both the benefit $(\Phi(y_t + \psi\iota) - \Phi(y_t))P_t$ and the cost $x_t - \alpha e^{2\xi_t}x_t^2$ of an attack are positively correlated.

Finally, we verify that the token price and transaction fees are indeed (weakly) increasing in A_t. Let us conjecture that the token price, P_t, and transaction fees are (weakly) increasing in A_t. We further define $P_t = 0$ whenever there is market breakdown. Under this assumption, strategic attacks occur when A_t is sufficiently small by these arguments. It then follows that strategic attacks preserve the monotonicity of P_t in A_t from Proposition 3, confirming the conjecture. Similarly, because a higher token price is associated with a higher user population and consequently, higher transaction fees, this confirms our second conjecture. Further, because the strategic attacks occur when the mining fundamental, ξ_t, is sufficiently small and mining has no direct impact on platform performance when there is no strategic attack, it follows that the token price and user participation are (weakly) increasing in ξ_t. □

Appendix B. A More General Setting

In this appendix, we illustrate the robustness of our key insight about the fragility of token platforms when there are network effects and token nonneutrality. To do this, we first consider a static version of our model and then, discuss the role of token retrading and a more general endowment distribution.

Suppose, as in the main model, that there is a continuum of users who choose whether to join the platform to exchange goods with each other. If a user joins the platform and matches with a trading partner, she derives utility over her own good C_i and that of her trading partner C_j:

$$U_i(C_i, C_j; \mathcal{N}) = \left(\frac{C_i}{1-\eta_c}\right)^{1-\eta_c}\left(\frac{C_j}{\eta_c}\right)^{\eta_c},$$

where $\mathcal{N}$ is the set of users on the platform. User i receives an endowment e^{A_i} and pays a fraction β of his trade surplus in transaction fees. To join the platform, the user has to pay a participation cost κ. In contrast to the main model, however, a user also receives utility from his token holdings X_i such that she receives a Cobb–Douglas convenience yield over his holdings and her expected trading benefit:

$$u_i = X_i^\alpha((1-\beta)\mathbb{E}[U_{i,t}(C_i,C_j;\mathcal{N}) \mid A_i])^{1-\alpha},$$

with share weights $\alpha \in (0,1)$ and $1-\alpha$, respectively. This preference for token holding could reflect an unmodeled convenience from holding tokens. In contrast to Cong et al. (2021b, 2022a), we do not impose token neutrality. As such, we assume users have a preference for the token balance rather than its nominal value in the numeraire good. We could also allow for the convenience yield to be increasing in the token price, P, provided this benefit does not impose token neutrality and is not sufficiently convex that it leads to a trivial corner solution in which all users participate regardless of the price.

If the tokens are sold at a uniform price P and users have quasilinear preferences, then user i solves the following optimization program:

$$\max_{X_{i,t}} (X_{i,t}^\alpha((1-\beta)\mathbb{E}[U_{i,t}(C_{i,t},C_{j,t};\mathcal{N}_t) \mid A_{i,t}])^{1-\alpha} - P_t X_{i,t} - \kappa)\mathbf{1}_{\{X_{i,t}>0\}}.$$

If a user does not join the platform, she receives an outside option normalized to zero.

Furthermore, from our main analysis, recall that if users follow a cutoff strategy and join the platform if $A_{i,t} \ge A_t^*$, then

$$\mathbb{E}[U_i(C_{i,t}, C_{j,t}; \mathcal{N}_t) \mid A_{i,t}] = (1-\beta)e^{(1-\eta_c)(A_{i,t}-A_t)+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}} \Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \frac{A_t^*-A_t}{\tau_\varepsilon^{-1/2}}\right).$$

It is then immediate that if a user intends to join the platform, her optimal choice of tokens is

$$X_{i,t} = \left(\frac{\alpha}{P_t}\right)^{\frac{1}{1-\alpha}}(1-\beta)e^{(1-\eta_c)(A_{i,t}-A_t)+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}} \Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \frac{A_t^*-A_t}{\tau_\varepsilon^{-1/2}}\right).$$

As such, the maximized utility of user i is

$$\max\left\{\left(\frac{\alpha}{P}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)(1-\beta)e^{(1-\eta_c)(A_{i,t}-A_t)+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}} \Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \frac{A_t^*-A_t}{\tau_\varepsilon^{-1/2}}\right) - \kappa, 0\right\}.$$

The marginal user with endowment $A_{i,t} = A_t^*$ is indifferent to joining the platform, which imposes dropping t subscripts

$$\alpha^{\frac{\alpha}{1-\alpha}}(1-\alpha)(1-\beta)e^{(1-\eta_c)(A_t^*-A_t)+A_t+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}} \Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \frac{A^*-A}{\tau_\varepsilon^{-1/2}}\right) = \kappa P^{\frac{\alpha}{1-\alpha}}. \tag{B.1}$$

As in the main model, the token price reflects the marginal user's convenience yield.

Let the supply of tokens be $x(P, y, \xi)$, where y is the amount of outstanding tokens and ξ is a negative supply shock, such as a shock to speculator sentiment. We assume that $x(\cdot, y, \xi)$ is a strictly increasing function for all (y, ξ) and that $x(0, y, \xi) = 0$. This more general token supply curve can reflect an arbitrary predetermined token issuance schedule.

Market clearing in the token market then implies that

$$\int_{-\infty}^{\infty} X_i(A_i)d\Phi(\log A_i) = \left(\frac{\alpha}{P}\right)^{\frac{1}{1-\alpha}}(1-\beta)U = x(P, y, \xi), \tag{B.2}$$

where U is the total transaction surplus:

$$U = e^{A+\frac{1}{2}\left((1-\eta_c)^2+\eta_c^2\right)\tau_\varepsilon^{-1}}\Phi\left((1-\eta_c)\tau_\varepsilon^{-1/2} + \frac{A-A^*}{\tau_\varepsilon^{-1/2}}\right) \Phi\left(\eta_c\tau_\varepsilon^{-1/2} + \frac{A-A^*}{\tau_\varepsilon^{-1/2}}\right).$$

Substituting (B.1) into (B.2), we arrive at the following condition:

$$\frac{\alpha}{1-\alpha}\kappa e^{-(1-\eta_c)(A^*-A)+\frac{1}{2}(1-\eta_c)^2\tau_\varepsilon^{-1}}\Phi\left((1-\eta_c)\tau_\varepsilon^{-1/2} - \frac{A^*-A}{\tau_\varepsilon^{-1/2}}\right) = X(P, y, \xi)P. \tag{B.3}$$

The left-hand side of (B.3) is strictly decreasing in $A^* - A \in (-\infty, \infty)$ from ∞ to zero, whereas the right-hand side is a horizontal line with value $X(P, y)P$ for all values of $A^* - A$. Consequently, there always exists a solution for the token price P from (B.3). Because $x(P, y, \xi)P$ is a strictly increasing function of P, we can invert it to express the price as

$$P = f^{-1}\left(\frac{\alpha}{1-\alpha}\kappa e^{A-A^*+\frac{1}{2}(1-\eta_c)^2\tau_\varepsilon^{-1}}\Phi\left((1-\eta_c)\tau_\varepsilon^{-1/2} + \frac{A-A^*}{\tau_\varepsilon^{-1/2}}\right); y, \xi\right) = p(A-A^*; y, \xi),$$

where $f^{-1}(\cdot, y, \xi)$ is the inverse of $x(P, y, \xi)P$ for a given (y, ξ) pair. Because the left-hand side of (B.3) is strictly decreasing in $A^* - A$, $p(A - A^*; y, \xi)$ is a strictly decreasing function of A^*. If all users join, then $X(P, y)P = \infty$, which suggests a price of $p(\infty; y, \xi) = \infty$.

Although we can always find a unique price for a given participation cutoff A^*, we must now find A^* from (B.1) by rewriting the condition as

$$\frac{\alpha^{\frac{\alpha}{1-\alpha}}(1-\alpha)(1-\beta)}{\kappa}e^{A^*+\frac{1}{2}\eta_c^2\tau_\varepsilon^{-1}}\Phi\left(\eta_c\tau_\varepsilon^{-1/2} - \frac{A^*-A}{\tau_\varepsilon^{-1/2}}\right) = p(A-A^*; y, \xi)^{\frac{\alpha}{1-\alpha}}. \tag{B.4}$$

Notice that the left-hand side of (B.4) is hump shaped in A^*, tending to 0 at $A^* \in \{-\infty, \infty\}$, whereas the right-hand side is a strictly decreasing function of A^* from ∞ at $A^* = -\infty$ to $\underline{0}$ for $A^* = \infty$. Consequently, we have a situation similar to that in the main model, in which the inverse demand curve $p(A - A^*; y, \xi)$ may remain above the hump-shaped curve for $A^* > -\infty$. In this case, only a trivial solution may exist. If $p(A - A^*; y, \xi)$ shifts upward, for instance, because of a more negative supply shock, ξ, then the set of A for which there is a nontrivial solution shrinks. Consequently, the token market fragility illustrated in our main model is still present in this more general model.

B.1. Token Retradeability

Until now, we have abstracted from token retradeability that comes with a dynamic setting. We can incorporate this easily into our analysis by assuming that tokens can be resold after users transact for a final value δ/y per token, where the division by y reflects the dilution of final value based on the amount of outstanding token supply. In this situation, the user's token demand is

$$X_i = \left(\frac{\alpha}{P - \delta/y}\right)^{\frac{1}{1-\alpha}} (1-\beta) e^{(1-\eta_c)(A^* - A_i) + A_t + \frac{1}{2}\eta_c^2 \tau_\varepsilon^{-1}} \Phi\left(\eta_c \tau_\varepsilon^{-1/2} - \frac{A^* - A}{\tau_\varepsilon^{-1/2}}\right),$$

and (B.1) becomes

$$\alpha^{\frac{\alpha}{1-\alpha}} (1-\alpha)(1-\beta) e^{A^* + \frac{1}{2}\eta_c^2 \tau_\varepsilon^{-1}} \Phi\left(\eta_c \tau_\varepsilon^{-1/2} - \frac{A^* - A}{\tau_\varepsilon^{-1/2}}\right) = \kappa (P - \delta/y)^{\frac{\alpha}{1-\alpha}}. \tag{B.5}$$

Note that the issues remain the same as in the static model, except now the term δ/y shifts down the right-hand side of (B.5). Consequently, having a high retrade value increases the region of existence of a nontrivial solution, and this effect is dampened by more tokens that have been issued, y. These again echo the results of our main model.

B.2. More General Endowment Distribution

We now relax the assumption that the endowments of users at each date are normally distributed. Instead of assuming a normal distribution, we let the endowment of agent i follow a general distribution: $A_{i,t} \sim \mathcal{G}(A_{i,t} | A_t)$ with support $A \in [-\infty, \overline{A}]$. Assuming users follow a cutoff strategy,

$$X_{i,t} = \left(\frac{\alpha}{P_t}\right)^{\frac{1}{1-\alpha}} (1-\beta) e^{(1-\eta_c) A_{i,t}} \mathbb{E}[e^{\eta_c A_{j,t}} \mathbf{1}_{\{A \in \mathcal{N}_t\}}].$$

As such, the maximized utility of user i is

$$\max\left\{\left(\frac{\alpha}{P}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha)(1-\beta) e^{(1-\eta_c) A_{i,t}} \mathbb{E}[e^{\eta_c A_{j,t}} \mathbf{1}_{\{A_{j,t} \in \mathcal{N}_t\}}] - \kappa, 0\right\}.$$

It is immediate that it is again optimal for each user to follow a cutoff strategy and join the platform if $A_{i,t} \geq A_t^*$. The marginal user with endowment $A_{i,t} = A_t^*$ is indifferent to joining the platform, which imposes (after dropping t subscripts)

$$\alpha^{\frac{\alpha}{1-\alpha}} (1-\alpha)(1-\beta) e^{(1-\eta_c) A^*} \mathbb{E}[e^{\eta_c A_j} \mathbf{1}_{\{A_j \in \mathcal{N}_t\}}] = \kappa P^{\frac{\alpha}{1-\alpha}}. \tag{B.6}$$

As in the main model, the token price still reflects the marginal user's convenience yield. If the marginal user has $A^* = \overline{A}$, then the left-hand side of (B.6) is zero because the probability of matching with another user, $\mathbb{E}[\mathbf{1}_{\{A_j \in \mathcal{N}\}}]$, is zero. Similarly, if all users join the platform, then $A^* = -\infty$, and the left-hand side is again zero. Consequently, the left-hand side of (B.6) is hump shaped in A^*, as in Figure 1.

It is immediate then that the assumption of a normally distributed endowment process is not essential for our market breakdown analysis.

Appendix C. Microfoundation of Speculator Demand

In this appendix, we provide a parsimonious model of speculators to aggregate their trading. We assume that there are overlapping generations of speculators. At each date, two types of speculators participate in the token market. The first is a group of passive speculators who enter at the beginning of each date and acquire all of the $\Phi(y_t)$ tokens from the previous generation of users, validators, and speculators. These passive speculators provide liquidity to exiting token holders and do not engage in any additional token trading.

The second is a group of active speculators of unit mass who choose whether to short sell tokens based on their expectations of the next-period token price. We assume that active speculator k has a noisy expectation of the next-period token price:

$$\mathbb{E}^{S,k}[P_{t+1} | \mathcal{I}_t] = e^{\zeta_{k,t}} RP_t, \tag{C.1}$$

where $\mathcal{I}_t$ is the public information set, RP_t is the required risk-neutral return for holding the token to the next period, and $\zeta_{k,t} \sim$ i.i.d $\mathcal{N}(\zeta_t - y_t, 1)$. $\zeta_{k,t}$ represents speculator k's sentiment, and $\zeta_t \sim$ i.i.d $\mathcal{N}(0, \sigma_\zeta^2)$ represents speculators' common sentiment shock in period t. That speculator sentiment is decreasing in the token supply, y_t, represents a time trend that speculators' enthusiasm for a new platform declines as the platform matures. Each active speculator can short either zero or one token. In deciding whether to short sell a token, speculator k faces an opportunity cost for her position of $(RP_t)^{1+\lambda}$ for $\lambda > 0$. As such, she chooses $X_{k,t}^S \in \{-1, 0\}$ to maximize

$$U_{k,t}^S = \max_{X_{k,t}^S} [\mathbb{E}^{S,k}[P_{t+1} | \mathcal{I}_t] - (RP_t)^{1+\lambda}] X_{k,t}^S.$$

Choosing $X_{k,t}^S = -1$ indicates that the speculator is short selling a token. Substituting with Equation (C.1), it is straightforward to see that speculator k follows a cutoff policy of short selling a token with a cutoff at the sentiment level $\lambda \log(RP_t)$:

$$X_{k,t}^S = \begin{cases} 0 & \text{if } \zeta_{k,t} \geq \lambda \log(RP_t) \\ -1 & \text{if } \zeta_{k,t} < \lambda \log(RP_t). \end{cases}$$

As a result, the aggregate demand of the speculators X^S is the sum of their passive and active positions

$$X^S = \Phi(y_t) - \int_{-\infty}^{\lambda \log(RP_t)} d\Phi(\zeta_{k,t}) = \Phi(y_t) - \Phi\left(y_t + \lambda \log(RP_t) - \zeta_t\right).$$

It should be clear that we use these two types of speculators to separately capture their collective buying and selling activities. Although this structure is somewhat mechanical, the aggregate demand curve for speculators has sensible economic properties; it increases with speculator sentiment ζ_t and decreases with the token price P_t. This particular functional form facilitates tractability of our model without necessarily driving any of our key results. This microfoundation also makes it clear that each speculator is atomistic and therefore, cannot internalize the impact of his trading on others.

Appendix D. Microfoundation of Strategic Attack

In this appendix, we provide a microfoundation for the strategic attack condition in the main paper. Specifically, we examine whether rogue miners wish to collude to engage in a 51% "double-spending" attack. This requires that a group of miners amasses enough computational power, compared with the rest of the mining community, to be able to verify, on average, the majority of transactions on the blockchain. Conceptually, by winning enough blocks to add to the blockchain, these corrupt miners will be able to eventually validate their own blocks on the longest chain or to mine secretly a second chain longer than the current blockchain and broadcast it to the mining community as the legitimate chain. When this occurs, these miners can reverse their own transactions to undo their expenditures, returning their spent tokens to their wallet to be spent again. This is the so-called "double-spending" problem. By creating duplicate tokens, the strategic attack temporarily increases the token supply through fraudulent inflation.[27]

The benefits and costs of a 51% attack are linked to participation by both users and miners. As more miners join the mining pool, the probability of completing any transaction and adding it to the blockchain falls, increasing the effective computational cost of attacking the currency. In addition, user and miner participation also increases the computational cost of an attack through the difficulty of mining each transaction or the hash rate. Many PoW protocols, such as those of Bitcoin and Ethereum, set the hash rate to maintain a fixed average time for new blocks to be added to the blockchain, and the hash rate increases in the number of users and miners to prevent blocks from being added too quickly. As a consequence, having more subscribers and a more diverse mining pool can make the platform more secure.

We assume that miners lack commitment, which is consistent with the static incentives miners face because of free entry (e.g., Abadi and Brunnermeier 2018). Any miner can attack the blockchain by engaging in a 51% attack to "double spend" the coins received from seigniorage. If corrupt miners attack the blockchain, the strategic attack artificially inflates the token base by $\Phi(y_t + \psi\iota) - \Phi(y_t)$, for $\psi > 0$, and the miner sells these additional tokens to earn $(\Phi(y_t + \psi\iota) - \Phi(y_t))P_t$ in additional revenue. These additional tokens have to be absorbed by users and speculators by increasing the effective token supply to $\Phi(y_t + \psi\iota)$. In addition, because the corrupt miners add over half the blocks to the blockchain, they earn 50% of the transaction fees from users and seigniorage. As a result of increased waiting times and service denials, users also experience a loss in expectation a fraction $1 - \gamma$ of their trade surplus.[28]

To acquire 51% of the computing power, corrupt miners must replicate the mining power of the existing $N_{M,t}$ miners by expending a convex technological cost $\alpha N_{M,t}^2$, where $\alpha > 0$. That the cost is convexly increasing in the number of miners $N_{M,t}$ reflects that it is increasingly difficult to acquire more mining power because of additional hardware and electricity costs.[29] To join the strategic attack, a potential attacker has to pay a participation cost, which can be viewed as the cost or disutility of coordinating with the other attackers. We normalize this cost to one in the numeraire good.

Suppose that $N_{M,t}$ miners provide mining services at date t and that a fraction p_t of miners attacks and splits the proceeds from the attack equally. They then need to acquire half of the total mining power, and consequently, they must acquire $N_{M,t}$ in additional mining power. An attack will occur when the benefit—the fraudulent seignorage and additional half of the seignorage and transaction fees—is greater than the cost of doubling the existing computing power of the mining community:

$$\left(\Phi(y_t + \psi\iota) - \Phi(y_t) + \frac{1}{2}(\Phi(y_t) - \Phi(y_t - \iota))\right)P_t + \frac{1}{2}\beta\gamma U_t - \alpha N_{M,t}^2 \geq 0.$$

When this happens, a strategic attack occurs. When this condition is satisfied, however, all miners will want to attack the platform, which will dilute the mining power and undermine a strategic attack. As this cannot be an equilibrium, the miners must play a mixed strategy when a strategic attack is possible. The probability of a miner attacking, p_t, is the date t probability then ensures that every miner is indifferent to attacking based on the outcome of an i.i.d. draw of a Bernoulli random variable with $\Pr(Attack) = p_t$. By the weak law of large numbers, exactly a fraction p_t of the existing mining pool will attack. This probability satisfies that the fraction $\frac{1}{p_t}$ of the revenue from attacking is offset by the disutility of participation

$$\frac{\left(\Phi(y_t + \psi\iota) - \frac{1}{2}\Phi(y_t) - \frac{1}{2}\Phi(y_t - \iota)\right)P_t + \frac{1}{2}\beta\gamma U_t - \alpha N_{M,t}^2}{p_t N_{M,t}} - 1 = 0,$$

from which follows, when $p_t > 0$, that

$$p_t = \frac{\left(\Phi(y_t + \psi\iota) - \frac{1}{2}\Phi(y_t) - \frac{1}{2}\Phi(y_t - \iota)\right)P_t + \frac{1}{2}\beta\gamma U_t - \alpha N_{M,t}^2}{N_{M,t}};$$

otherwise, there is no attack. Consequently, we can interpret the strategic attack condition (8) as arising from a 51% attack on the currency, and the possibility of attack leads to a stability boundary in the state space of the platform.

Endnotes

[1] In contrast, coins (and altcoins), such as Bitcoin and Litecoin, are fiat currencies that are maintained on a public blockchain ledger by a decentralized population of record keepers, whereas security tokens are financial assets that trade in secondary markets on exchanges and whose initial sale is recorded on the blockchain of the currency that the issuer accepts as payment. Coins are typically created through "forks" from existing currencies, such as Bitcoin Gold from Bitcoin, and by airdrops, in which the developer sends coins to wallets in an existing currency to profit from the price appreciation of its retained stake if the new currency becomes widely adopted. Security tokens are typically sold through ICOs structured as "smart contracts" on existing blockchains, such as that of Ethereum.

[2] We depart from Sockin and Xiong (2023) along several substantive dimensions. First, in terms of emphasis, their focus is on platform governance, whereas our focus is on token price dynamics and platform stability. Second, in terms of information structure, in their setting the fundamental underpinning the aggregate transaction surplus of users is the only fundamental; in this paper, we extend their analysis to include not only a time-varying token supply but also, nonfundamental fluctuations in the token price from the optimism of users and the sentiment of speculators. Third, although the token price is endogenous in both settings, in ours it is determined by market clearing in a secondary token market; however, in Sockin and Xiong (2023), it is set by the developer that controls the supply of tokens. Finally, in terms of empirical implications, Sockin and Xiong (2023) use their model to explain cross-sectional patterns in ICOs; this paper instead focuses on time series patterns (e.g., momentum, reversal, life-cycle effects, relation to investor attention) and cross-sectional patterns (e.g., size effect) in cryptocurrency returns.

[3] Although transaction fees paid on decentralized crypto platforms can adjust to token price fluctuations, such as gas fees under the proof of stake protocol, the services on these platforms are more rigid in the number of tokens required for their services. To claim a username on Decentraland, for instance, requires 100 MANA tokens regardless of their value in US Dollars. Such costs are necessary to engage with other users on the platform and consequently, contribute to a user's decision to participate.

[4] In an earlier version of the paper, we considered an extended setting in which the platform fundamental is unobservable. In that setting, users use their endowments as a private signal about this fundamental; the token price and the transaction history on the blockchain act as public signals that aggregate their dispersed information. This second public signal reflects that the blockchain technology supporting cryptocurrencies acts as an indelible and verifiable ledger that records the decentralized transactions that take place on the platform. In this extended setting, we show that informational frictions attenuate the risk of breakdown by dampening price volatility and platform performance.

[5] As Liu and Tsyvinski (2021) find little evidence that cryptocurrencies load on conventional sources of systematic risk, such as market or style factors, such an assumption for the token market is realistic.

[6] The nonneutrality of the token price is highly realistic on many crypto platforms. On Axie Infinity, for example, users can breed an axie up to seven times using Small Love Potion (SLP) tokens according to a rising scale (i.e., currently 900 SLP for the first breed up to 15,300 for the seventh). On Socios, users buy fan tokens associated with specific sports teams that convey certain benefits and voting rights on team decisions. These features are independent of the token price. On Friends with Benefits (FWB), users can purchase a local membership with 5 FWB tokens and a global membership with 75 FWB tokens. In addition, users currently pay one FWB token to get access to the newsletter and five for access to the global network of token-gated parties.

[7] Online platforms often face severe commitment issues. Facebook, for instance, changed its data policies over time (for example, Beacon in 2007 and the 2008 Terms of Service update) and settled with the Federal Trade Commission in 2011 for violating privacy promises. Amazon engages in "copycat" practices on two-sided platforms that harm sellers. A rigid token issuance schedule consequently represents one safeguard that decentralization token platforms have in place to protect users. See Sockin and Xiong (2023) for an analysis of trade-offs associated with decentralization.

[8] Such rigid, predetermined inflation schedules are ubiquitous in practice. Solana, for instance, currently has an annual inflation rate of 8% that is scheduled to decrease by 15% per year to a long-term inflation rate of 1.5% (https://blockdaemon.com/products/white-label-validator/how-solana-staking-works/). Polkadot instead sets its inflation rate as a function of the proportion of DOT tokens that are staked but tries to maintain about 10% per year (https://wiki.polkadot.network/docs/learn-staking).

[9] We implicitly assume a frictionless secondary market for tokens. See, for instance, Capponi and Jia (2021) for liquidity issues associated with cryptocurrency exchanges.

[10] In contrast to traditional multisided platforms, such as in Evans (2003) and Rochet and Tirole (2003), the owner issues a native token to users that has a floating exchange rate with other tokens and currencies instead of collecting discriminating participation fees. This potentially buffers the pricing of the platform's services from external shocks, such as monetary policy shocks to fiat currencies, by denominating them in the native token and disciplines their valuation through price discovery in financial markets.

[11] We assume the owner completes all transactions without censorship or charging monopoly markups. See Huberman et al. (2021) for how proof of work-decentralized consensus can overcome these issues at the cost of transaction delays. We also assume that the owner can commit to a token inflation schedule. See Cong et al. (2022a) for a setting in which the owner cannot commit.

[12] This feature also contrasts the neutrality of the token price adopted by Cong et al. (2021b). In their model, each user's benefit from holding a token is determined by the market value of her token holdings in the numeraire rather than the number of tokens.

[13] Even though the market breakdown is a severe form of market dysfunction, it may not present an arbitrage opportunity to speculators for several reasons. First, a platform's tokens derive all their value from the convenience yield that users receive from transacting on the platform. Thus, whether a platform's token price can recover from zero is ultimately determined by users rather than speculators. Second, it is difficult for a platform that relies on network effects to recover once users have lost interest in it. The case of Myspace after the rise of Facebook is a particularly salient example of the fickleness of network effects.

[14] Although one may argue that validators can alter the token supply schedule to mitigate market breakdown, achieving consensus among stakeholders to alter token inflation on decentralized platforms is extremely difficult in practice. For instance, as EthHub describes of the Ethereum platform, "[a]s Ethereum is a decentralized network, the Monetary Policy cannot be successfully modified unless there is overwhelming consensus from the aforementioned stakeholders" (https://cryptobriefing.com/ethereum-sound-money-like-bitcoin/).

[15] The second (high-cutoff) and third (highest-cutoff) equilibria may or may not exist at any given date depending on the expected retrade value of the token. As such, they are dynamically unstable, and we can eliminate them as predictions for the equilibrium outcome. In addition, the second (high-cutoff) equilibrium is unstable, even when fixing the token's expected retrade value. Introducing a small amount of noise into users' participation decisions, for instance, and letting this noise become arbitrarily small would ensure convergence away from this second equilibrium to the highest-price equilibrium.

[16] Although such an analysis of optimal platform policy is beyond the scope of this paper, Mei and Sockin (2022) show that a platform owner may find it optimal to inflate the token base to ensure nontrivial user participation when incentives to speculate are particularly severe.

[17] This issue has also received significant attention in the literature. See, for instance, Budish (2018), Pagnotta (2022), and Chiu and Koeppl (2023).

[18] In practice, several miners are randomly drawn from a queue to compete to complete each transaction, and miners often pool their revenue to insure each other against the risk of not being selected.

See Cong et al. (2021a) for an extensive analysis of this issue. Our modeling of mining as a static problem when there is free entry is consistent with that in Abadi and Brunnermeier (2018).

[19] In an earlier version of this paper, we endow each miner with a heterogeneous but correlated fixed cost of mining, $\xi_{j,t}$. Because this heterogeneity across miners does not impact our qualitative insights, we abstract from it to minimize notation and simplify exposition.

[20] To focus on the broader implications of the cryptocurrency for users, we abstract from the strategic considerations that miners face in adding blocks to the blockchain to collect fees, such as consensus protocols and on which chain to add a block. See, for instance, Biais et al. (2019) and Easley et al. (2019) for game-theoretic investigations into these issues.

[21] Although in principle, mining pools could coordinate to preempt a strategic attack, their primary function is risk sharing. Further, such coordination would undermine the spirit of the decentralized consensus protocol. In May 2019, the BTC.top and BTC.com mining pools, with a combined 44% mining power, were criticized for coordinating an "attack" on the BTC Cash blockchain to reverse a hacker's transactions.

[22] A subtle issue is how to measure the marginal user's convenience yield in practice. If users were all identical, then the average transaction fee would be this yield. With selection onto the platform, however, a reasonable, noisy proxy is the minimum transaction size on the blockchain.

[23] Although the measure is constructed with searches for "Bitcoin" specifically, we view this measure as a noisy proxy for interest in cryptocurrencies more generally.

[24] Our proof is based on a modified argument of Milgrom and Roberts (1994) for comparative statics in the presence of strategic complementarity.

[25] We recognize that $\partial G_t/\partial z_t = 0$ at the critical value of z_t, at which breakdown occurs if the fundamentals deteriorate.

[26] Because there is no profit when $f(y_{t-1}, P_t, E[U_t|\mathcal{I}_t]) = 0$ and only a loss in revenue from honest mining, it follows that miners would rather not attack at the indifference threshold.

[27] To date, the major attacks on blockchains have been 51%. In 2015, the Bitcoin mining pool ghash.io voluntarily committed to reducing its share of mining power from over 50% to less than 40% to assuage fears of it coordinating a potential 51% attack among its miners on the currency. There is even a website, Crypto51, that tracks the computational cost of a 51% attack in real time.

[28] Hackers have also engaged in 51% attacks to disrupt the blockchain to undermine confidence in cryptocurrency. Although hackers can double spend, they cannot steal tokens from user wallets.

[29] Implicitly, we assume that to avoid detection by the mining pool, these rogue miners must acquire additional computing power to compete with their own honest mining.

References

Abadi J, Brunnermeier M (2018) Blockchain economics. Working paper, Princeton University, Princeton, NJ.

Athey S, Parashkevov I, Sarukkai V, Xia J (2016) Bitcoin pricing, adoption, and usage: Theory and evidence. Working paper, Stanford University Graduate School of Business, Stanford, CA.

Beaudry P, Portier F (2006) Stock prices, news, and economic fluctuations. *Amer. Econom. Rev.* 96:1293–1307.

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

Biais B, Bisiere C, Bouvard M, Casamatta C, Menkveld A (2023) Equilibrium Bitcoin pricing. *J. Finance* 78:967–1014.

Black F (1986) Noise. *J. Finance* 41:528–543.

Budish E (2018) The economic limits of Bitcoin and the blockchain. Working paper, University of Chicago, Chicago.

Capponi A, Jia R (2021) The adoption of blockchain-based decentralized exchanges. Working paper, Columbia University, New York.

Capponi A, Jia R, Wang Y (2021) The evolution of blockchain: From lit to dark. Working paper, Columbia University, New York.

Capponi A, Ólafsson S, Alsabah H (2023) Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Management Sci.* 69(11):6455–6481.

Chiu J, Koeppl TV (2022) The economics of cryptocurrencies: Bitcoin and beyond. *Canadian J. Econom.* 55(4):1762–1798.

Cong LW, He Z (2019) Blockchain disruption and smart contracts. *Rev. Financial Stud.* 32:1754–1797.

Cong LW, He Z, Li J (2021a) Decentralized mining in centralized pools. *Rev. Financial Stud.* 34:1191–1235.

Cong LW, Li Y, Wang N (2021b) Tokenomics: Dynamic adoption and valuation. *Rev. Financial Stud.* 34:1105–1155.

Cong LW, Li Y, Wang N (2022a) Token-based platform finance. *J. Financial Econom.* 144:972–991.

Diba BT, Grossman HI (1988) The theory of rational bubbles in stock prices. *Econom. J.* 98:746–754.

Easley D, Kleinberg J (2010) *Networks, Crowds, and Markets: Reasoning About a Highly Connected World* (Cambridge University Press, Cambridge, UK).

Easley D, O'Hara M, Basu S (2019) From mining to markets: The evolution of Bitcoin transaction fees. *J. Financial Econom.* 134:91–109.

Evans D (2003) The antitrust of multi-sided platform markets. *Yale J. Regulation* 20:325–381.

Fracassi C, Kogan S (2022) Pure momentum in cryptocurrency markets. Working paper, University of Texas at Austin, Austin.

Goldstein I, Gupta D, Sverchkov R (2019) Utility tokens as a commitment to competition. Working paper, Wharton School, University of Pennsylvania, Philadelphia.

Hackethal A, Hanspal T, Lammer DM, Rink K (2021) The characteristics and portfolio behavior of Bitcoin investors: Evidence from indirect cryptocurrency investments. *Rev. Finance* 26:855–898.

Hu A, Parlour C, Rajan U (2019) Cryptocurrencies: Stylized facts on a new investible instrument. *Financial Management* 48:1049–1068.

Huberman G, Leshno J, Moallemi CC (2021) An economic analysis of the Bitcoin payment system. *Rev. Econom. Stud.* 88:3011–3040.

Kiyotaki N, Wright R (1993) A search-theoretic approach to monetary economics. *Amer. Econom. Rev.* 83:63–77.

Liu Y, Tsyvinski A (2021) Risks and returns of cryptocurrency. *Rev. Financial Stud.* 34:2689–2727.

Liu Y, Tsyvinski A, Wu X (2022) Common risk factors in cryptocurrency. *J. Finance* 77:1133–1177.

Mayer S (2019) Token-based platforms and speculators. Working paper, HEC Paris, Paris, France.

Mei K, Sockin M (2022) A theory of speculation in community assets. Working paper, University of Texas at Austin, Austin, TX.

Milgrom P, Roberts J (1994) Comparing equilibria. *Amer. Econom. Rev.* 84:441–459.

Pagnotta E (2022) Decentralized money: Bitcoin prices and blockchain security. *Rev. Financial Stud.* 35:866–907.

Rochet J-C, Tirole J (2003) Platform competition in two-sided markets. *J. Eur. Econom. Assoc.* 1:990–1029.

Saleh F (2021) Blockchain without waste: Proof-of-stake. *Rev. Financial Stud.* 34:1156–1190.

Schilling L, Uhlig H (2019) Some simple Bitcoin economics. *J. Monetary Econom.* 106:16–26.

Schwenkler G, Zheng H (2021) News-driven peer co-movement in crypto markets. Working paper, Santa Clara University, Santa Clara, CA.

Shams A (2020) The structure of cryptocurrency returns. Working paper, Ohio State University, Columbus, OH.

Sockin M, Xiong W (2023) Decentralization through tokenization. *J. Finance* 78:247–299.