

Token-Based Platform Governance*

Joseph Abadi[†] Markus Brunnermeier[‡]

September 18, 2024

Abstract

We develop a model to compare the governance of traditional shareholder-owned platforms to that of platforms that issue tokens. A traditional shareholder governance structure leads a platform to extract rents from its users. A platform that issues tokens for its services can mitigate this rent extraction, as rent extraction lowers the platform owners' token seigniorage revenues. However, this mitigation from issuing "service tokens" is effective only if the platform can commit itself not to dilute the "service token" subsequently. Issuing "hybrid tokens" that bundle claims on the platform's services and its profits enhances efficiency even absent ex-ante commitment power. Finally, giving users the right to vote on platform policies, by contrast, redistributes surplus but does not necessarily enhance efficiency.

Keywords: Utility Tokens, Platforms, DeFi, Corporate Governance

*Toni Whited was the editor for this article. We thank Agostino Capponi, Hanna Halaburda, Sylvain Chassang, and an anonymous referee for helpful suggestions, as well as seminar participants at the Canadian Economic Association meetings, the CMU Secure Blockchain Summit, the FRB Philadelphia Digital Currencies conference, the Tokenomics conference, and Princeton University. **Disclaimer:** The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

[†]Corresponding author. Federal Reserve Bank of Philadelphia, 10 N Independence Mall W, Philadelphia, PA 19106 (joseph.abadi@phil.frb.org).

[‡]Princeton University, 20 Washington Road, Princeton, NJ 08542 (markus@princeton.edu).

1 Introduction

The emergence of platforms that issue their own digital currencies or credits to users – called *tokens* – has offered an alternative to traditional platforms’ model of financing and governance. Traditional platforms are owned and governed by shareholders (i.e., residual cash flow claimants). A platform’s users, by contrast, do not necessarily play a role in financing or decision-making. Shareholders may not govern a traditional platform in users’ best interests: they can exercise the platform’s market power to extract rents from users. Proponents of cryptocurrency and decentralized finance (DeFi) argue that by giving users a stake in the system, token-financed platforms are more likely to be governed in accordance with users’ preferences.

Tokens can offer users several types of claims or rights. Some platforms issue tokens that offer *transaction services* while retaining the traditional shareholder governance model. For example, the (centralized) Binance cryptocurrency exchange issues a token (BNB) that users can redeem to receive a discount on trading fees. Blockchain-based platforms often issue “utility tokens” that can be used to purchase a digital service: e.g., the Golem platform’s tokens (GLM) can be used to rent out computational resources, and Chainlink’s token (LINK) is used to pay network operators to retrieve data for smart contracts. Other tokens function essentially like shares of a traditional platform, granting *cash flow claims* and *voting rights*. The Kyber decentralized cryptocurrency exchange, for instance, issues a token (KNC) that can be “staked” to receive a share of the platform’s revenues and to participate in governance. Still other tokens *bundle* transaction services with cash flow claims and/or voting rights. The archetypal example is a proof-of-stake cryptocurrency with “on-chain” governance, like Tezos or Algorand. On these blockchains, tokens play a dual role: users can hold them to transact with others, and “validators” can set tokens aside as collateral (called “staking”) to verify blockchain transactions, earn monetary rewards, and vote on policies.

The advent of token-financed platforms raises key economic questions. Can token-financed platforms succeed in mitigating rent extraction and aligning policies with users’ interests? How should tokens be designed to promote efficient platform governance? Should tokens grant transaction services, cash flows claims, voting rights, or some combination of features?

To answer these questions, we develop a unified model of a platform economy that is general enough to encompass traditional platforms as well as platforms that issue tokens with various features. The model is set in continuous time and has two groups of agents: users, who enjoy the platform’s transaction services, and investors, who hold cash flow claims on the platform but do not engage in transactions.

In our benchmark model, the platform’s policy consists of a transaction fee charged to

users and (possibly) a rate of token seigniorage, when we consider a platform that issues tokens. These policies dictate the *split of surplus* between users and investors, since fees and seigniorage both represent transfers from users to investors, as well as *total surplus*, since the costs users face to use the platform determine their demand for its transaction services. Governance decisions (i.e., opportunities to vote on the platform’s policies) occur periodically over time. Hence, there is *limited commitment* in governance: instead of committing to a full sequence of policies at $t = 0$, the platform can rewrite its policies in each governance decision.

The platform’s *market power* generates scope for inefficiencies. The platform faces no competition, so it can set fees higher than the marginal cost of processing transactions without losing its user base. Importantly, fees are *distortionary* because the platform cannot fully extract surplus from users: higher fees lead to lower transaction volumes and deadweight losses.

For each of the platform designs we consider, there are two types of assets. There is a *transaction asset* that users must hold to receive the platform’s transaction services: users’ flow payoffs depend on their real balances of transaction assets, as in many models of platforms that issue tokens (Cong, Li, and Wang, 2021; Gryglewicz, Mayer, and Morellec, 2021). There is also a *cash flow asset* held by investors that grants claims on the platform’s profits. Either type of asset can have voting rights, depending on the setting. We consider three different platform designs in this general environment.

- **Traditional platform:** Users transact with an asset that originates outside the platform, such as cash, deposits, or other liquid assets. The platform issues *shares* to investors that confer cash flow and governance rights (so shares are the cash flow asset). Investors choose the platform’s policies to maximize its equity value.
- **Service tokens:** Users transact with *tokens* that are issued by the platform (so tokens are the transaction asset). Investors hold shares (the cash flow asset) that grant claims on profits. At first, we assume the platform maintains the shareholder governance model, but we later extend the model to permit token-holding users to vote as well.
- **Hybrid tokens:** The platform does not issue shares – it issues a *token* that serves as both a transaction asset and a cash flow asset. Users hold tokens in order to transact on the platform, whereas investors “stake” tokens to receive a claim on the platform’s profits. Again, we begin by assuming that only the holders of staked tokens (investors) can vote, and then extend the model to permit users to vote as well.

We study the efficiency of each platform design in terms of total surplus vis-à-vis the first-best allocation.

Traditional platform: The shareholders of a traditional platform simply set fees to maximize the present value of profits without regard for user surplus. Consequently, as is typical of models with imperfect competition, the platform’s fees are set higher than the marginal cost of processing transactions. Transaction volumes are therefore *inefficiently low* from a social perspective.

Service tokens: When the platform issues service tokens, investors choose both fees and the rate of seigniorage. They have a reason to internalize user surplus: the equilibrium price of tokens, and therefore the platform’s seigniorage revenues from its initial token issuance, depends on the *service flow* that users expect to receive from tokens in the future (i.e., the marginal benefit of holding a token). Policies that enhance users’ welfare (e.g., a promise of lower fees or less future seigniorage) increase service flows and thus token prices. If investors’ ability to commit to future policies is strong enough, then the prospect of greater initial seigniorage revenues will incentivize them to choose policies that are more beneficial to users. Therefore, under commitment, investors will set lower fees than a traditional platform, leading to greater transaction demand. The equilibrium outcome is unambiguously more efficient than in the traditional case.

However, if investors’ ability to commit to future policies is weak, then this logic breaks down. Investors no longer have as strong an incentive to pass policies that benefit users: they can frequently rewrite policies and do not internalize any reduction in value of tokens issued *in the past*. Instead, users bear those costs. Investors are tempted to set fees too high and over-issue tokens to boost seigniorage revenues. This temptation is so severe that in the limit of no commitment, there does not exist an equilibrium in which tokens are valued: realizing that investors will attempt to extract high rents in the future, users are unwilling to purchase tokens in the first place. Commitment is thus crucial for service tokens to enhance efficiency on their own.

Our analysis of service tokens sets the stage for us to ask: what type of token design can overcome investors’ time-consistency problem? Of course, there are several mechanisms that could enhance platform owners’ commitment to policies that benefit users, such as token retention or smart contracts that pre-program a specified sequence of policies. Our paper’s main result, however, shows that so long as investors have limited commitment power, a hybrid token that *bundles transaction services with cash flow claims* can serve as an effective substitute for commitment.

Hybrid tokens: A hybrid token is held by users for its transaction services and staked by investors for its cash flows. Equilibrium token prices reflect both their *service value* (the present value of service flows, which is users’ valuation) and their *cash flow value* (the present value of dividends, which is investors’ valuation). Even when investors govern the platform

without strong commitment to future policies, equilibrium outcomes are unambiguously more efficient than in the traditional platform case. Here our main result, in fact, is that hybrid tokens overcome investors’ time-inconsistency problem: equilibrium governance decisions are precisely the same as in the case of a platform that issues service tokens with full commitment. The key intuition is that investors hold an asset whose value reflects users’ future service flows, so investors bear part of the costs if they pass policies that harm users. When a platform issues service tokens to users without commitment, by contrast, investors may seek to pass policies that increase the value of their equity while reducing token prices. This mechanism is quite general: in the Online Appendix, we show that the introduction of hybrid tokens resolves investors’ time-inconsistency problem in several extensions of our benchmark model.¹

User voting: We then extend the model to accommodate token-issuing platforms that permit users to vote, rather than just investors. We consider both a platform that issues service tokens and one that issues hybrid tokens. When users can participate in governance, they receive a greater share of total surplus, but equilibrium outcomes are not necessarily more efficient. If users acquire a majority share of voting power, they pass policies that increase their surplus at the expense of lower profits – transaction quantities may be inefficiently *high*. Intuitively, simply *reallocating voting power* does not cause users to internalize investor welfare. By contrast, *bundling cash flow claims with transaction services* causes investors to internalize user surplus, unlike the owners of a traditional platform.

Organization. In the remainder of this section, we give a review of the related literature. Section 2 gives a brief overview of the types of tokens issued by platforms in practice. Section 3 introduces the economic environment and other preliminary elements of the benchmark model. Section 4 studies the governance of a traditional platform as a benchmark. Section 5 introduces service tokens, outlines how token issuance affects equilibrium governance decisions, and highlights the time-consistency problem that is central to our analysis. Section 6 analyzes the hybrid token scheme and proves our main result: that hybrid tokens achieve the full-commitment outcome. Sections 7 lays out an extension with user voting. Section 8 discusses the model’s main assumptions. Section 9 concludes. All proofs are in the Appendix.

Related literature. Our paper is most closely related to the emerging literature that studies the role of tokens in DeFi platforms’ governance. In the context of a platform with network externalities, Sockin and Xiong (2023) study the introduction of a token that grants platform membership and permits users to vote on platform policies, preventing the platform from exploiting their data. However, users are not able to share the costs of investments in the platform and therefore cannot subsidize the admission of new users to the platform. Bakos

¹We thank an anonymous referee for noting that hybrid tokens can potentially overcome time-inconsistent governance in a broad range of environments.

and Halaburda (2022) study a platform that issues tokens that offer cash flow claims and voting rights. They highlight conditions under which token holdings become concentrated among non-users, leading to rent extraction. Similarly, Han, Lee, and Li (2023) develop and empirically test a model in which concentrated token holdings by a large investor can undermine efficient governance. Relatedly, Bena and Zhang (2022) and Gan, Tsoukalas and Netessine (2023) compare the inefficiencies in governance of a platform that issues service tokens to those of a traditional platform. While our analysis shares some of these themes, it is complementary: we characterize the separate roles of tokens' transaction services, cash flow claims, and voting rights, providing novel insights into the optimal *design* of tokens.

The broader literature on financing through token sales and ICOs is also related to our work. Closest to our paper, Goldstein, Gupta, and Sverchkov (2024) show that by issuing utility tokens, a platform can commit to charge lower prices to users, as in our model. Their mechanism, however, is related to the Coase (1972) conjecture and is quite distinct from ours. Gryglewicz, Mayer, and Morellec (2021) and Cong, Li, and Wang (2022) study the optimal issuance of tokens by a financially-constrained platform, demonstrating how seigniorage policies can be used to reward platform owners for investments. Li and Mann (2017), Chod and Lyandres (2021), and Lee and Parlour (2022) study other reasons why firms might finance themselves through the issuance of utility tokens. Li and Mayer (2020) and d'Avernas, Maurin, and Vandeweyer (2022) present models to study the optimal issuance of stablecoins. Similarly, You and Rogoff (2023) study how the tradability of a platform's utility tokens affects the revenue raised by a token offering. Relative to this literature, our paper differs in that it considers the role of tokens *exclusively* for governance – there are no financial frictions that motivate token issuance.

Of course, our paper connects to the corporate governance literature as well. There is an extensive body of work on control rights, ownership structure, and the theory of the firm stemming from the work of Coase (1937), Williamson (1979), and Grossman and Hart (1986). Our paper contributes to this literature by characterizing the specific governance consequences brought about by different token designs – we show that despite the fact that users can potentially be exploited by the platform, it is not always most efficient to give them control rights (Hansmann, 1988). Our work is complementary to the literature that studies how different control structures aggregate information in governance decisions (see Aghion and Tirole, 1997, among many others). Recent work has extended this literature to the study of DeFi platforms (Tsoukalas and Falk, 2020; Benhaim, Falk, and Tsoukalas, 2023).

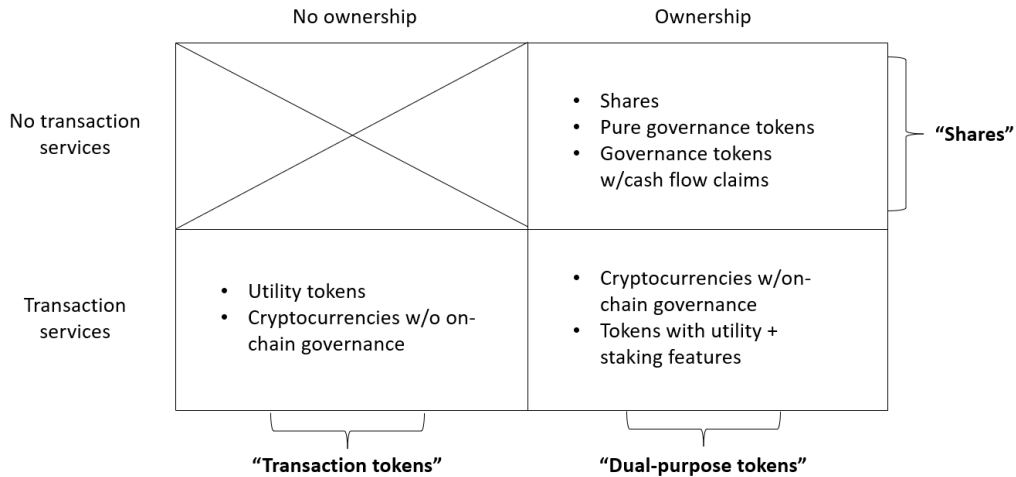


Figure 1: A taxonomy of the types of tokens issued by platforms in practice. Tokens are categorized by whether they confer ownership in the platform (cash flow claims and/or voting rights) and whether they offer transaction services. We also indicate how we refer to each type of token in the model.

2 An Overview of Tokens

Before introducing the model, we briefly outline the different claims or rights that tokens may confer and examples of tokens that are issued in practice. Tokens typically grant at least one of (1) claims on a platform’s transaction services, (2) claims on cash flows, or (3) voting rights. Figure 1 provides a taxonomy of tokens based on whether they provide transaction services and whether they confer platform ownership (cash flow claims and/or voting rights). This section discusses each class of tokens in this taxonomy as well as how the different types of tokens used in practice relate to the assets in the model.

Transaction services only: Tokens that grant only claims on transaction services are typically referred to as *utility tokens*. Examples are the Binance token (BNB) or the Golem token (GLM) discussed in the Introduction, which users can redeem for specific services or perks. However, tokens can facilitate transactions on a platform even if they are not redeemable for any particular service. In the context of our model, a *pure cryptocurrency* that has no intrinsic value but is used for transactions among a platform’s users could also be viewed as a token that offers an implicit claim on the platform’s transaction processing services. Even outside of DeFi, some platforms have begun to issue, or have considered issuing, their own currencies (e.g., Alibaba’s “Alipay” or Facebook’s now-defunct Libra/Diem

project).²

Ownership features only: Other tokens grant claims on the platform’s cash flows (e.g., transaction fees or seigniorage revenues).³ Usually, such tokens have voting rights in governance decisions as well. The (previously discussed) KNC token issued by the Kyber decentralized exchange (DEX) is a leading example. Kyber provides automated cryptocurrency market-making services and collects transaction fees from traders. Token holders can “stake” their tokens in order to receive a share of these fees. They may also participate in *on-chain governance*: the community regularly votes on referenda that determine the platform’s policies, including fees and software upgrades, and voting power is allocated proportionally to token holdings. Most of the assets that users transact on the platforms are cryptocurrencies that originate elsewhere.⁴ In the context of our model, a token that offers *only* cash flow claims and voting rights is equivalent to a share of a traditional platform. So, in the case of Kyber, the “investors” would be KNC token holders and “users” would be those who trade various other cryptocurrencies on the platform.

Pure governance tokens offer only voting rights. For example, the Uniswap DEX issues the UNI token, which entitles holders to vote on changes to the market-making protocol. UNI does not currently pay its holders any dividends, but in principle, token holders could vote to pay themselves a dividend at some point in the future.⁵ The COMP token issued by the popular Compound lending platform carries similar voting rights. Tokens that have only voting rights are beyond the scope of our model.

Transaction services and ownership features: Finally, we turn to platforms that issue *native tokens* that bundle transaction services with ownership features. We have already given the example of proof-of-stake cryptocurrency blockchains with on-chain governance, like Algorand or Tezos. On these platforms, tokens are held by users who wish to transact with others. They are staked by “validators” (the analogue of investors in the model) who run the computational hardware needed to verify transactions and collect monetary rewards, which could take the form of transaction fees or newly minted tokens. Some blockchain platforms permit any token holder to stake and vote on proposed policies (e.g., Algorand), whereas others allow users to delegate their votes to validators who they trust to act on their behalf (e.g., Tezos). Typical policies adjust the blockchain’s transaction fees or upgrade transaction

²A key difference, however, is that tokens issued by non-DeFi platforms are typically backed at least partially by existing fiat currencies.

³These are distinct from *security tokens*, which usually represent a claim on another firm’s profits or a claim on a financial asset that exists outside of the blockchain.

⁴Kyber users, for example, may access a “liquidity pool” that permits them to trade Ether for the Tether stablecoin.

⁵See <https://protos.com/to-fee-or-not-to-fee-that-is-the-question-does-uniswap-have-an-answer/>. Whether UNI will eventually pay dividends has been a topic of intense speculation in the community.

verification protocols.

This setup is not restricted to proof-of-stake cryptocurrencies, however. Some DeFi platforms issue tokens that bundle voting rights, cash flow claims, and direct claims on the platform’s services. The Aave lending platform issues a token (AAVE) with voting rights that (1) borrowers can post as collateral to receive discounted interest rates and (2) investors can stake to provide a liquidity backstop and receive a share of Aave’s profits. Similarly, some platforms that enable interoperability across DeFi applications, such as the Cosmos and Polkadot networks, issue tokens with voting rights that users hold to pay network fees and validators stake to provide transaction security.

3 Model

Environment: We consider a continuous-time, infinite-horizon economy in which agents interact on a *platform*. There are two commodities: a numéraire good (referred to as a “dollar”) and transaction services (henceforth “transactions”) produced by the platform at a marginal cost $c > 0$. The economy is populated by a unit mass of two types of agents: *users* $i \in [0, 1]$ and *investors* $j \in [0, 1]$. Users enjoy the platform’s transaction services: a user i who consumes a quantity x_{it} of transactions at time t receives utility $\frac{x_{it}^{1-\gamma}}{1-\gamma}$, where $\gamma \in (0, 1)$.⁶ Investors, on the other hand, hold cash flow claims on the platform but do not enjoy its transaction services. All agents are risk-neutral over consumption of dollars and share a common discount rate $r > 0$.⁷

Assets: In this economy, assets can play two roles. First, there is a *transaction asset* (with endogenous price Q_t^T and supply A_t^T) that can be held by users to receive the platform’s transaction services: users enjoy transaction services equal to their real balances of transaction assets, as is typical in the literature on tokens (as well as models with money in the utility function, e.g., Sidrauski, 1967; Feenstra, 1986).⁸ So a user i who holds a quantity of transaction assets a_{it} at time t receives transaction services

$$x_{it} = Q_t^T a_{it}. \tag{1}$$

Second, there is a *cash flow asset* (with price Q_t^C and supply A_t^C) that is held by investors⁹ and provides a pro-rata claim on the platform’s profits: a cash flow asset pays a dividend

⁶We assume $\gamma < 1$ to ensure users’ transaction demand is sufficiently elastic that a profit-maximizing policy for the platform exists.

⁷Throughout, we make the additional parametric assumption $c/(1-\gamma) > r$ to streamline the presentation of results. However, the main results remain unchanged if we drop this assumption.

⁸Biais et al. (2023) micro-found a similar utility function in an overlapping-generations model.

⁹Section 8 discusses the assumption that users do not hold cash flow assets.

$dD_t = d\Pi_t/A_t^C$ at time t , where $d\Pi_t$ denotes the platform's profits.

We consider three schemes for the design of assets in this economy: a traditional platform, a platform that issues service tokens, and a platform that issues hybrid tokens. The Introduction specifies what plays the role of the cash flow asset and the transaction asset under each scheme. We will mostly study the case in which only cash flow assets have voting rights (so investors govern the platform), but Section 7 extends the model to allow users to vote as well.

Platform governance: The platform's *policy* at time t consists of a *transaction fee* $f_{t+s} \geq 0$ and (possibly) a *token seigniorage* policy at all future dates $t + s$.¹⁰ The platform's policies are determined in *governance decisions* at times $\{\tau_0 = 0, \tau_1, \tau_2, \dots\}$ that arrive according to a Poisson process at rate λ . In a governance decision at time τ_k , any previous policy commitments are torn up, and new policies (for all future dates $\tau_k + s$) are chosen by a vote among the agents who govern the platform. Hence, the parameter λ indexes the degree of *commitment* to future policies: in the limit $\lambda \rightarrow 0$, the platform's policies are fully determined at $t = 0$ (as with an immutable smart contract), whereas in the limit $\lambda \rightarrow \infty$, a new policy is chosen at each instant. Agents are infinitesimal, so no individual agent's vote is ever pivotal. Therefore, all agents take policies as given.

Payoffs: A user i who engages in x_{it} transactions at time t pays a fee f_t per transaction, so the user receives a total flow payoff from transactions

$$dU_{it} = \left(\frac{x_{it}^{1-\gamma}}{1-\gamma} - f_t x_{it} \right) dt.$$

The platform receives the transaction fees paid by users and incurs a marginal cost c per transaction. When we consider token-issuing platforms, the platform will also receive (endogenous) seigniorage revenues from token issuance, which we denote by dS_t for now. Thus, letting $X_t \equiv \int_0^1 x_{it} di$ denote the aggregate quantity of transactions, the platform's flow profits at t are

$$d\Pi_t = (f_t - c)X_t dt + dS_t.$$

In the remainder of this section, we describe the elements of our model that are held constant across the different settings considered. In subsequent sections, we analyze each platform design individually.

Remark. *We illustrate the logic of our main results in a benchmark model that makes several*

¹⁰While we assume that the platform cannot subsidize transactions for simplicity, all that is required for our results is that the platform cannot promise *unboundedly large* subsidies to users. That is, there must exist some lower bound \underline{f} such that $f_t \geq \underline{f}$. The inability to provide large subsidies may be viewed as another dimension of limited commitment.

specific assumptions. However, we extend the model to more general settings to show that the main results continue to hold:

- Online Appendix F considers a platform that must make governance decisions about investment as well as fees and seigniorage;
- Online Appendix G considers a platform in which users' utility exhibits network effects;¹¹
- Online Appendix H considers a setting with monopolistic competition across platforms, rather than a single monopolistic platform;¹²
- Online Appendix I considers a setting in which the platform permits users to redeem transaction assets in exchange for a service (instead of the money-in-the-utility assumption used in the benchmark model).

3.1 Individual optimization problems

We first lay out agents' individual portfolio optimization problems. Users choose their transaction asset holdings to maximize their expected lifetime utility subject to a standard budget constraint, taking the price of transaction assets Q_t^T and fees f_t as given. In Appendix A.1, we show that this problem reduces to a static one:

$$\max_{x_{it}, a_{it}} \left(\frac{x_{it}^{1-\gamma}}{1-\gamma} - f_t x_{it} \right) dt + (\mathbb{E}_t[dQ_t^T] - rQ_t^T dt) a_{it} \quad \text{s.t.} \quad (1).$$

The first term in parentheses is the flow utility of transactions x_{it} , whereas the second term represents the expected return on transaction assets net of holding costs $rQ_t^T a_{it}$. All users i optimally choose the same transaction quantity x_{it} . Users' optimality condition can be used to show that aggregate transaction demand X_t obeys

$$X_t^{-\gamma} dt = f_t dt + (rdt - \mathbb{E}_t \left[\frac{dQ_t^T}{Q_t^T} \right]). \quad (2)$$

Transaction demand is decreasing in the net marginal cost of transacting, which consists of two components: the transaction fee f_t and the opportunity cost $r - \frac{1}{dt} \mathbb{E}_t \left[\frac{dQ_t^T}{Q_t^T} \right]$ of holding transaction assets. Of course, the aggregate demand for transaction assets is equal to

¹¹We assume that a user's utility depends on their own transactions x_{it} as well as other users' aggregate transaction activity X_t via $\frac{(X_t^\nu x_{it}^{1-\nu})^{1-\gamma}}{1-\gamma}$ with $\nu \in (0, 1)$; i.e., ν captures the strength of network effects.

¹²We assume a continuum of platforms k that produce differentiated transaction services. A user's total transaction services x_{it} are a CES aggregate of transactions x_{ikt} on each platform k , $x_{it} = \left(\int_0^1 x_{ikt}^{\frac{\eta-1}{\eta}} dk \right)^{\frac{\eta}{\eta-1}}$.

aggregate transaction demand,

$$X_t = Q_t^T A_t^T. \quad (3)$$

Similarly, investors choose their holdings of cash flow assets to maximize lifetime utility subject to a standard budget constraint. Investors' problem is formally stated in Appendix A.1. Their first-order condition implies that cash flow assets are priced according to the present value of dividends:

$$rQ_t^C dt = dD_t + \mathbb{E}_t[dQ_t^C] \quad \text{where} \quad dD_t = \frac{d\Pi_t}{A_t^C}. \quad (4)$$

Equations (3) and (4) summarize the *demand* for transaction assets and cash flow assets, respectively. The *supply* of each asset is determined in a different way for each of the platform designs we consider.

3.2 Governance decisions

The platform's *status quo policy* at time t consists of a transaction fee and (possibly) a seigniorage rate at all future dates, $\{f_{t+s}, dS_{t+s}\}_{s \geq 0}$. In a governance decision at time τ_k , the platform's status quo policy can be revised. Each agent who has the right to do so votes for a new policy. If some policy attains a majority, then the status quo is abandoned and the new policy is implemented. Otherwise, the status quo is maintained.

In principle, this voting game could be quite complicated since there are infinitely many potential policies. However, our environment offers a useful simplification: within each constituency (investors or users), agents have identical policy preferences. We assume that agents in each constituency vote unanimously for their most-preferred policy.¹³ Throughout most of the analysis, we will focus on the case in which only cash flow assets confer the right to vote. In this case, investors will hold all of the voting power, so they will implement their most-preferred policy. However, we will later extend the model to allow users the opportunity to vote as well.

Each constituency's preferences over policies is determined by the lifetime utility its members expect to obtain after a new policy is passed.¹⁴ Investors' expected lifetime utility V_t^I is equal to the market value of their cash flow assets, whereas users' expected lifetime utility V_t^U

¹³Hence, we ignore self-fulfilling equilibria in which agents do not vote for their most-preferred policy (as is typical in the corporate governance literature, see Levit, Malenko, and Maug, 2024). For instance, if all agents vote for a random policy, then no policy will ever attain a majority, so it is in fact individually rational for each agent to vote randomly.

¹⁴Given an expected path of policies, there may be multiple equilibrium paths of transaction quantities $\{X_t\}$ consistent with those policies. Following convention in the literature, we select the "best-case" equilibrium that generates the greatest total surplus.

is equal to the market value of their transaction assets plus the present value of *infra-marginal rents* R_t that they earn from transactions, defined below.

Proposition 1. *Following the announcement of a new policy at time τ , investors' expected lifetime utility is*

$$V_\tau^I = Q_\tau^C A_{\tau-}^C, \quad (5)$$

where $A_{\tau-}^C$ denotes investors' cash flow asset holdings before the new policy announcement, and users' expected lifetime utility is

$$V_\tau^U = Q_\tau^T A_{\tau-}^T + R_\tau, \quad \text{where } R_\tau \equiv \mathbb{E}_\tau \left[\int_0^\infty e^{-rs} \frac{\gamma}{1-\gamma} X_{\tau+s}^{1-\gamma} ds \right]. \quad (6)$$

3.3 Notation

Throughout the analysis, we will focus on *Markov equilibria*. Suppose that the most recent governance decision occurred at time τ . Then, the only relevant state variable for outcomes at time t is the time $s = t - \tau$ that has elapsed since the most recent decision.

As is typical in models of policy-making with limited commitment, both current policies and anticipated future policies influence agents' behavior in equilibrium. It is therefore necessary to distinguish between *actual* policies chosen at τ and the policies that were *anticipated* before the governance decision at τ , since, at least in principle, the constituency that governs the platform can deviate from the anticipated policy. We index actual policies and outcomes simply by the subscript s : for example, f_s denotes the chosen level of fees at time $\tau + s$, X_s denotes actual aggregate transaction quantities at that time, and so on. We denote anticipated policies and outcomes with a hat as functions of s : so \hat{f}_s denotes the level of fees at time $\tau + s$ that was anticipated *before* the decision at τ , \hat{X}_s denotes anticipated transaction quantities, etc. Of course, in equilibrium, actual outcomes must coincide with anticipated outcomes.

We focus on equilibria in which jumps in variables may occur at the time of a governance decision, but variables evolve smoothly between governance decisions. The drift of a variable is denoted with a dot, so, for example, $dX_s = \dot{X}_s ds$ for $s > 0$.

3.4 The first-best

Before examining specific platform designs, we derive the properties of optimal allocations in this environment. This analysis will facilitate a comparison of the inefficiencies that arise under each platform design studied in subsequent sections.

There is transferable utility in this environment, so an allocation is efficient if and only if it maximizes utilitarian social welfare (i.e., the sum of agents' payoffs). Note that fees will be irrelevant for total welfare, since they are just a transfer from users to investors.

An *allocation* is therefore summarized simply by a sequence of aggregate transaction quantities X_t .¹⁵ Total surplus at time t is just $X_t^{1-\gamma}/(1-\gamma) - cX_t$. To maintain symmetry with our model of governance, we assume a social planner with limited commitment: at the time of a governance decision, the planner chooses an allocation $\{X_s\}_{s \geq 0}$ that is maintained until the next governance decision at time τ (which arrives at rate λ).

Proposition 2. *A first-best allocation solves*

$$\hat{V}_0^P = \max_{X_s} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \left(\frac{X_s^{1-\gamma}}{1-\gamma} - cX_s \right) ds + e^{-r\tau} \hat{V}_0^P \right], \quad (7)$$

where \hat{V}_s^P is the planner's value function. An allocation is first-best if and only if

$$X_s = X^{FB} \equiv c^{-\frac{1}{\gamma}} \quad \forall s. \quad (8)$$

The first-best level of transactions, X^{FB} , is set so that the marginal utility of an additional transaction, $X_s^{-\gamma}$, is equal to the marginal cost c of processing that transaction. Thus, the optimal level of aggregate transactions is constant over time and decreasing in transaction processing costs c .

4 Traditional platform

In this section, we study the case of a *traditional platform* as a simple benchmark. We demonstrate that as in most models with a monopolistic firm, a traditional platform charges inefficiently high fees, distorting transaction volumes downwards.

4.1 Setup

In the case of a traditional platform, there are two distinct assets: *shares* that are issued by the platform and *transaction assets* that originate outside the platform (such as cash, stablecoins, or another cryptocurrency). Users hold transaction assets for their transaction services but cannot vote in governance decisions. We assume transaction assets are supplied

¹⁵Users have identical concave utility functions. Without loss of generality, then, we can restrict attention to allocations in which each user transacts the same amount, $x_{it} = X_t$ for all i .

elastically at a price $Q_t^T = 1$.¹⁶ Shares serve as the economy's cash flow asset and grant investors the right to vote on the platform's policies. The supply of shares is normalized to $A_t^C = 1$ – we assume, without loss of generality, that the platform does not issue new shares or buy them back from investors.¹⁷

Since the platform does not issue tokens in this environment, it receives no seigniorage revenues. Therefore, it derives profits only from processing transactions:

$$d\Pi_t = (f_t - c)X_t dt. \quad (9)$$

4.2 Equilibrium

We look for a Markov equilibrium in one state variable: the time s elapsed since the most recent governance decision. We begin by solving for users' transaction demand. The price of transaction assets is constant, so (2) implies that users' transaction demand is downward-sloping in the level of fees:

$$X_s = (f_s + r)^{-\frac{1}{\gamma}}. \quad (10)$$

Equations (4) and (5) imply that investors' expected lifetime utility is equal to the present value of profits $\mathbb{E}_0[\int_0^\infty e^{-rt} d\Pi_t]$. In a governance decision, investors unanimously vote in favor of the policy $\{f_s\}_{s \geq 0}$ that maximizes the present value of profits. Investors' governance problem can then be written as

$$\hat{V}_0^I = \max_{f_s, X_s} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} (f_s - c) X_s ds + e^{-r\tau} \hat{V}_0^I \right] \text{ s.t. (10), } f_s \geq 0. \quad (11)$$

where τ denotes the (random) time interval from $s = 0$ until the next governance decision and \hat{V}_s^I denotes investors' value function at time s since the most recent governance decision.

A *Markov equilibrium* consists of a value function \hat{V}_s^I for investors and outcomes $\{\hat{f}_s, \hat{X}_s\}$ that solve (11).

4.3 Welfare and efficiency under the traditional scheme

Under the traditional governance scheme, the platform is just a monopolistic firm that maximizes the present value of its profits. Investors do not internalize how changes in the

¹⁶What matters is not that the price of transaction assets is constant, but rather that (1) their price is exogenous to the platform's policies, and (2) their rate of return is below the discount rate r (that is, there is a *liquidity premium* on transaction assets, as with cash, Treasury bills, or deposits).

¹⁷There are no financial frictions, so the Modigliani-Miller theorem applies in the setting of a traditional platform. The platform's share issuance policy is irrelevant.

platform's policies affect user surplus. As a result, equilibrium policies are inefficient: they maximize investor surplus at the expense of total surplus.

Proposition 3. *Under the traditional scheme, the equilibrium sequence of policies $\{f_t\}$ maximizes expected investor surplus (but not total surplus) over all feasible sequences of policies. That is, regardless of the degree of commitment λ , equilibrium policies solve*

$$\max_{f_t, X_t} \int_0^{\infty} e^{-rt} (f_t - c) X_t dt \quad s.t. \quad X_t = (f_t + r)^{-\frac{1}{\gamma}}. \quad (12)$$

Under the equilibrium policy, fees are

$$f_t = \frac{1}{1-\gamma} c + \frac{\gamma}{1-\gamma} r \quad (13)$$

and transaction quantities are inefficiently low:

$$X_t = X^{trad} \equiv \left(\frac{c+r}{1-\gamma} \right)^{-\frac{1}{\gamma}} < X^{FB}. \quad (14)$$

Note, moreover, that a traditional platform's equilibrium policy is time-consistent, so limited commitment to future policies (parameterized by the frequency λ of governance decisions) is irrelevant. Intuitively, this is the case because time- t transaction demand depends only on time- t policies. When we study a token-issuing platform, this will no longer be true: future token issuance policies will affect the current return on tokens and therefore current transaction demand.

Equilibrium transaction quantities X_t are inefficiently low because users' marginal cost of transacting, fees f_t plus the opportunity cost r of holding transaction assets, is greater than the social marginal cost c of processing transactions. By (13),

$$f_t + r = \frac{c+r}{1-\gamma} > c.$$

This equation for the effective transaction cost faced by users reveals two distortions, each corresponding to a distinct source of user surplus that investors neglect. First, the discount rate r appears in the numerator on the right-hand side. This distortion arises because investors do not take into account the aggregate *service flows*

$$SF_t \equiv X_t^{1-\gamma} - f_t X_t = r X_t, \quad (15)$$

that users earn on their transaction asset holdings, defined as users' marginal utility per

transaction times the aggregate quantity of transactions.¹⁸

Second, an additional factor $1 - \gamma$ appears in the denominator on the right-hand side. This is because investors do not internalize the *infra-marginal rents*

$$IR_t \equiv \left(\frac{X_t^{1-\gamma}}{1-\gamma} - f_t X_t \right) - SF_t = \frac{\gamma}{1-\gamma} X_t^{1-\gamma} \quad (16)$$

that users receive from infra-marginal transactions: since users' utility is concave in transactions, user surplus is greater than their marginal valuation of transaction assets' services.¹⁹

When we analyze a platform that issues tokens for transactions, we will demonstrate that, by contrast, investors take users' service flows into account. Hence, the first distortion will vanish, while the second will remain. Figure 2 illustrates how these distortions cause investors to set fees too high and destroy surplus.

There are three necessary ingredients for the inefficiency in this model. First, the platform has market power, so shareholder value maximization is not equivalent to maximization of social surplus. Put differently, the platform's market power creates a *conflict of interest* between the two constituencies. Second, the platform's owners can extract rents from users only by charging *distortionary* fees that cause deadweight losses: it is not possible for the platform to use a more complex pricing scheme (like a two-part tariff) that fully extracts user surplus. However, the Coase Theorem implies that absent restrictions on contracting, investors and users would nevertheless contract around these inefficiencies and arrive at an efficient outcome. The third necessary ingredient for inefficiency, therefore, is *limited contracting*: users cannot sign a contract in which they commit to compensate investors for choosing a more socially beneficial policy. These limits to contracting could be micro-founded, for instance, by assuming that users are unable to commit to a sequence of payments in response to the policies chosen by investors.

How can agents overcome the problem of limited contracting? In the next section, we outline conditions under which *token issuance* can partially substitute for the missing contracts between users and investors.

5 Service tokens

In this section we consider a platform that issues tokens that provide transaction services *only*. The platform continues to be governed by shareholders (investors) who hold all cash flow

¹⁸That is, SF_t is defined as $X_t \times \frac{\partial}{\partial X_t} \left(\frac{X_t^{1-\gamma}}{1-\gamma} - f_t X_t \right)$. Service flows are equal to rX_t in equilibrium by (10).

¹⁹Note that the discounted infra-marginal rents that enter users' lifetime utility in (6) are just $R_t = \mathbb{E}_t \left[\int_0^\infty e^{-rs} IR_{t+s} ds \right]$.

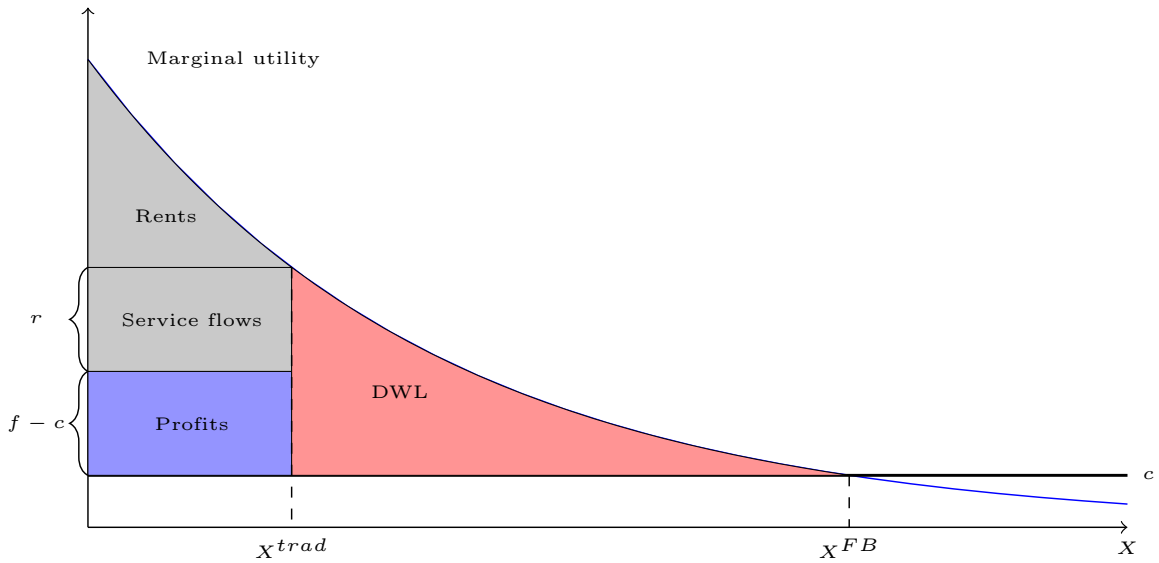


Figure 2: An illustration of the sources of neglected surplus in the case of a traditional platform. Aggregate transaction quantities X are on the horizontal axis. The downwards-sloping curve represents users' marginal utility of transacting (as a function of X). The horizontal line corresponds to the marginal cost c of processing a transaction. The profit-maximizing transaction quantity is X^{trad} , which is below the first-best X^{FB} .

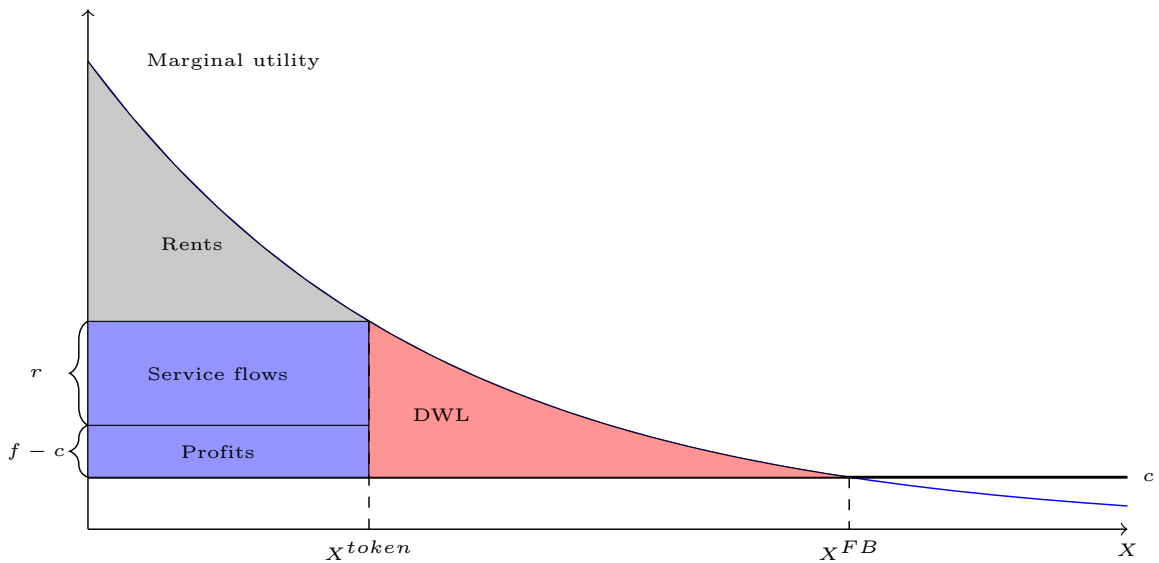


Figure 3: Illustration of the sources of surplus in the case of a platform that issues service tokens. Unlike in the case of a traditional platform (Figure 2), investors internalize users' service flows. Profits are maximized by setting lower fees than in the traditional case, raising transaction quantities to X^{token} (which remains below the first-best X^{FB}).

claims and voting rights. In contrast to the traditional setting, investors' ability to commit to future policies is key. We show that service tokens can substitute for missing contracts between users and investors (and therefore increase welfare) only if investors' commitment power is strong enough. When commitment power is weak, they are tempted to inflate away the value of tokens, depressing transaction quantities and resulting in a less efficient equilibrium.

5.1 Setup

Consider an environment in which the two assets are *shares* held by investors, which serve as the cash flow asset, and *tokens* that the platform issues to users, which serve as the transaction asset. Since the transaction asset is issued by the platform rather than supplied elastically, in this case its price Q_t^T evolves endogenously. We maintain the assumptions that (1) all voting rights are allocated to shareholders, and (2) the platform neither issues nor buys back shares, so the supply of shares is normalized to $A_t^C = 1$. The platform in this setting can be thought of as the issuer of a “utility token” or as a tech platform that issues its own currency.²⁰

Unlike a traditional platform, a token-issuing platform can earn seigniorage revenues by minting new tokens. We denote the growth rate of the token stock by

$$d\mu_t \equiv \frac{dA_t^T}{A_t^T} \geq 0.$$

To ensure that an optimal policy exists, we assume that the rate of token issuance is bounded above by some large positive constant, $\frac{A_t^T - A_{t-}^T}{A_{t-}^T} \leq \Delta$, which implies $d\mu_t \leq \frac{\Delta}{1+\Delta}$.²¹ The token issuance rate $d\mu_t$ is a policy determined in governance decisions. Seigniorage revenues are equal to the current token price times the quantity of tokens issued at time t , $dS_t = Q_t^T dA_t^T = X_t d\mu_t$. The platform's profits at time t are then

$$d\Pi_t = (f_t - c)X_t dt + X_t d\mu_t.$$

5.2 Equilibrium

We again consider a Markov equilibrium in which all outcomes depend only on the time s since the most recent governance decision. We look for an equilibrium in which variables

²⁰In our benchmark model, tokens have no intrinsic value (agents enjoy transaction services only if the price of tokens is positive). Online Appendix I outlines an extension in which tokens have intrinsic value: they can be redeemed for a service at a fixed exchange rate, as is common for utility tokens.

²¹All of our substantive results continue to hold in the limit of unrestricted token issuance, $\Delta \rightarrow \infty$.

may jump at the time of a governance decision ($s = 0$) but evolve smoothly thereafter. In particular, the platform may issue a discrete quantity of tokens when a new policy is implemented, $d\mu_0 > 0$, but afterwards, the supply of tokens evolves smoothly, $d\mu_s = \dot{\mu}_s ds$ for $s > 0$.

When the platform issues service tokens, users' transaction demand is no longer statically pinned down by fees as in (10). Transaction demand also depends on expected changes in token prices: all else equal, higher expected returns on tokens will increase demand. The following lemma characterizes the returns on tokens in this setting.

Lemma 1. *When the platform issues service tokens, the expected return on tokens satisfies*

$$\frac{1}{ds} \mathbb{E}_s \left[\frac{dQ^T}{Q^T} \right] = \frac{\dot{X}_s}{X_s} - \dot{\mu}_s + \lambda \left((1 - \hat{d}\mu_0) \frac{\hat{X}_0}{X_s} - 1 \right).$$

The first two terms represent returns in the absence of a new policy, while the third term represents returns conditional on a new policy (Recall that $\hat{d}\mu_0$ denotes anticipated token issuance at the time of the next governance decision, and \hat{X}_0 is anticipated transaction demand at that time.) Simply put, this lemma implies that the expected return on tokens is equal to the expected growth rate of transaction quantities minus the expected growth rate of the token stock. If the token stock grows without a commensurate increase in transaction demand, there will be inflation (a decrease in the price of tokens and a low return).

Lemma 1 allows us to characterize transaction demand in this setting – (2) implies

$$(r + \lambda)X_s = \underbrace{X_s^{1-\gamma} - (f_s + \dot{\mu}_s)X_s + \dot{X}_s}_{\text{current policy}} + \underbrace{\lambda(1 - \hat{d}\mu_0)\hat{X}_0}_{\text{exp. future policy}}. \quad (17)$$

Users' transaction demand is equal to the present value of service flows net of the costs of dilution from additional token issuance – new seigniorage reduces the price of tokens, representing an implicit transfer from token holders to investors. This transaction demand condition illustrates why lack of commitment matters. Transaction demand depends not only on current policy, but also on anticipated policy after the next governance decision, which arrives at rate λ . If users expect that investors will be tempted to issue a large quantity of tokens in the next governance decision ($\hat{d}\mu_0 > 0$), then they expect a reduction in the value of their tokens, lowering transaction demand.

As before, (4) and (5) imply that investors choose fees f_s and a token issuance policy

$(d\mu_0, \{\dot{\mu}_s\}_{s>0})$ to maximize the present value of profits:

$$\hat{V}_0^I = \max_{f_s, d\mu_0, \dot{\mu}_s, X_s} \mathbb{E}_0 \left[X_0 d\mu_0 + \int_0^\tau e^{-rs} (f_s + \dot{\mu}_s - c) X_s ds + e^{-r\tau} \hat{V}_0^I \right]$$

s.t. (17), $f_s, \dot{\mu}_s \geq 0$, $d\mu_0 \in [0, \frac{\Delta}{1+\Delta}]$, (18)

where τ denotes the time of the next governance decision. We search for a Markov equilibrium consisting of a value function \hat{V}_s^I for investors and outcomes $\{\hat{f}_s, \hat{d}\mu_0, \hat{\mu}_s, \hat{X}_s\}$ that solve (18).

We begin our analysis of equilibrium with some simplifying observations. Note that at the time when the governance decision takes place ($s = 0$), investors can issue tokens without affecting transaction demand for $s \geq 0$. Transactions instead depend on *future* token issuance ($\dot{\mu}_s$ for $s > 0$ and $\hat{d}\mu_0$ in the *next* governance decision), since that is what determines expected returns on tokens by Lemma 1. Therefore, at the time of a governance decision, investors issue a large quantity of new tokens and inflate away the value of existing tokens.

Lemma 2. *When the platform issues service tokens, investors issue the maximum feasible quantity of tokens at the time of a governance decision ($d\mu_0 = \frac{\Delta}{1+\Delta}$).*

In our benchmark model, this is the source of time-inconsistency – ex ante, investors would like to commit not to inflate away the value of tokens, but ex post (at the time of a governance decision), it is optimal to do so. Using this result, we henceforth denote the *expected* rate of dilution from such jumps in token issuance by

$$\hat{\lambda} \equiv \lambda \frac{\Delta}{1+\Delta}.$$

When a governance decision takes place, investors decide to issue a large quantity of tokens that inflates away a fraction $\frac{\Delta}{1+\Delta}$ of the value of existing tokens. The parameter $\hat{\lambda}$ will be a crucial measure of investors' ability to commit to seigniorage policies.²²

Lemma (2) can be used to write the platform's expected profits in a simple form. The platform's expected profits are equal to users' discounted service flows SF_s (given in (15)), plus discounted future profits from transaction processing, $(f_s - c)X_s$, plus a term that depends on anticipated future policies.

Proposition 4. *When the platform issues service tokens, the expected value of platform*

²²Note that even in the limit of unrestricted seigniorage, $\Delta \rightarrow \infty$, expected dilution remains finite, $\lim_{\Delta \rightarrow \infty} \hat{\lambda} = \lambda$.

profits on $[0, \tau)$ is

$$\mathbb{E}_0 \left[\int_0^\tau e^{-rs} d\Pi_s \right] = \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \left(\underbrace{X_s^{1-\gamma} - f_s X_s}_{\text{service flow } SF_s} + \underbrace{(f_s - c)X_s}_{\text{profits}} \right) ds + \underbrace{\frac{1}{1+\Delta} (e^{-r\tau} \hat{X}_0 - X_0)}_{\text{future policy}} \right]. \quad (19)$$

The intuition behind this result is straightforward. The platform makes profits in two ways: by charging transaction fees and by issuing new tokens. Users value tokens according to the present value of their service flows. Hence, when the platform issues new tokens at $s = 0$, it receives revenues equal to the present value of service flows. Thereafter, it earns the transaction processing fees that make up the remainder of its profits.

Unlike in the case of a traditional platform, then, investors take users' service flows into account when choosing policies. By passing policies that raise users' anticipated service flows, they increase token prices and therefore seigniorage revenues. Hence, investors' policy preferences will be more aligned with users' when the platform can issue tokens. Nevertheless, welfare will still fall short of the first-best: despite the fact that investors take users' service flows into account, they still fail to internalize how their policies affect users' infra-marginal rents.

Investors' ability to commit will be key in determining equilibrium outcomes. We will distinguish between two regimes in what follows: the *strong commitment regime* ($\hat{\lambda}$ small enough) and the *weak commitment regime* ($\hat{\lambda}$ large enough). Equilibrium outcomes will differ sharply across these two regimes. We analyze each in turn.

5.3 The strong commitment regime

We begin by analyzing the strong commitment regime. We show that when investors' commitment power is strong enough, their governance problem is *time-consistent* – that is, they choose the same policies that they would have chosen if they could commit to a full sequence of policies at $t = 0$. In this regime, equilibrium policies maximize the present value of service flows plus platform profits.

Proposition 5. *There exists λ^* (defined in (24)) such that if $\hat{\lambda} \leq \lambda^*$, investors' governance problem (18) is time-consistent. Equilibrium policies solve*

$$\begin{aligned} \max_{f_t, \dot{\mu}_t, X_t} \int_0^\infty e^{-rt} \left(X_t^{1-\gamma} - cX_t \right) dt \quad \text{s.t. } f_t, \dot{\mu}_t \geq 0, \\ rX_t = X_t^{1-\gamma} - (f_t + \dot{\mu}_t)X_t + \dot{X}_t. \end{aligned} \quad (20)$$

Under the equilibrium policy, transaction quantities satisfy

$$X_t = X^{token} = \left(\frac{c}{1-\gamma} \right)^{-\frac{1}{\gamma}}. \quad (21)$$

The quantity of transactions X_t and welfare are both higher than in the case of a traditional platform (but below their first-best levels).

In the Appendix, we take a Lagrangian approach to solve for equilibrium policies and transaction quantities, but we summarize the main results here. Transfers from users to investors (fees plus seigniorage revenues) are set statically to maximize service flows plus transaction processing profits,

$$X_t^{1-\gamma} - cX_t = \underbrace{X_t^{1-\gamma} - f_t X_t}_{\text{service flows } SF_t} + \underbrace{(f_t - c)X_t}_{\text{profits}}.$$

The equilibrium level of transactions is given in (21). Equilibrium transaction quantities are greater than in the case of a traditional platform (see (14)) but nevertheless remain below the first-best level (8).

When commitment to future policies is strong enough, then, token issuance enhances efficiency. The key idea is that since investors care about maintaining a high token price to maximize their seigniorage revenues at $s = 0$, they are reluctant to set fees too high. High fees imply low service flows for users, reducing the token price and the platform's seigniorage revenues. Hence, investors commit to lower fees than in the traditional setting. However, investors still fail to internalize how the platform's policies affect infra-marginal rents, so transaction quantities remain distorted downwards (hence the additional factor of $1 - \gamma$ in the denominator of (21) relative to the first-best level (8)). Figure 3 illustrates this point.

Why is the degree of commitment (i.e., the precise value of $\hat{\lambda}$) irrelevant in this setting? The answer lies in how investors choose to raise revenues from users. Conceptually, there are two ways the platform can raise revenues: by charging transaction fees or by issuing additional tokens. Recall that the expected rate of dilution in future governance decisions is $\hat{\lambda}$. Then, between governance decisions, users pay investors transaction fees $f_t X_t dt$ as well as seigniorage revenues $\mu_t X_t dt$. The transaction demand condition (17) then implies that if $f_t + \mu_t$ is constant, transaction demand satisfies

$$X_t = \left(r + f_t + \mu_t + \hat{\lambda} \right)^{-\frac{1}{\gamma}} = \left(\frac{c}{1-\gamma} \right)^{-\frac{1}{\gamma}},$$

where the second equality follows from (21). When $\hat{\lambda}$ is higher (i.e., commitment power is

weaker), investors know that in the future, they will be more tempted to raise revenues by issuing tokens. Then, they choose to charge lower fees and issue fewer tokens in the present to boost transaction demand and keep total transfers from users $(f_t + \dot{\mu}_t + \hat{\lambda})X_t$ constant. The platform's fee policy and its seigniorage policy are therefore substitutes: what matters for the platform's profits and transaction demand is the *sum* $f_t + \dot{\mu}_t + \hat{\lambda}$, not how revenues are split between fees, current seigniorage, and anticipated future seigniorage.

Proposition 6. *Under the optimal policy with commitment, the level of fees f_t and the growth rate of the token stock $\dot{\mu}_t$ are indeterminate. Their sum, $f_t + \dot{\mu}_t$, is uniquely determined in equilibrium:*

$$f_t + \dot{\mu}_t = \frac{c}{1-\gamma} - (r + \hat{\lambda}). \quad (22)$$

5.4 The weak commitment regime

We now turn to the case in which investors' commitment power is weak. In this regime, investors' temptation to inflate away the value of tokens is strong enough that it impairs transaction demand. In the limit of no commitment, expected inflation is so high that transaction quantities go to zero. The following proposition summarizes these results.

Proposition 7. *When $\hat{\lambda} > \lambda^*$ (defined in Proposition 5), investors' policy problem is no longer time-consistent. Equilibrium transaction quantities are*

$$X_t = (r + \hat{\lambda})^{-\frac{1}{\gamma}} < X^{token}. \quad (23)$$

In the no-commitment limit ($\hat{\lambda} \rightarrow \infty$), transaction quantities $X_t \rightarrow 0$.

When investors lack commitment power, their temptation to issue tokens creates inflation and undermines transaction demand. Specifically, Lemma 2 implies that in governance decisions, investors choose to issue a large quantity of new tokens and inflate away the value of existing tokens. Recall that *effective transfers* paid by users to investors are fees f_t plus current seigniorage $\dot{\mu}_t$ plus $\hat{\lambda}$, which determines anticipated future seigniorage. Previously, we argued that investors' desired effective fee satisfies $f_t + \dot{\mu}_t + \hat{\lambda} = c/(1-\gamma) - r$, so that the marginal cost faced by users (transfers to investors plus the opportunity cost r of holding tokens) is a markup over the marginal cost of transaction processing, $c/(1-\gamma)$. When anticipated future seigniorage is greater (higher $\hat{\lambda}$), investors either cut fees f_t or current seigniorage $\dot{\mu}_t$.

Now, however, suppose that

$$\hat{\lambda} > \lambda^* \equiv \frac{c}{1-\gamma} - r. \quad (24)$$

Then, even if investors cut fees and current seigniorage to zero ($f_t = \dot{\mu}_t = 0$), users still face such a high cost $\hat{\lambda}$ from expected future inflation that the total transfer to investors is greater than its optimal level. In this case, the entire transfer from users to investors consists of seigniorage during governance decisions, and the cost faced by users is expected seigniorage $\hat{\lambda}$ plus the opportunity cost r of holding tokens. Users' transaction demand thus falls below investors' desired level X^{token} , as shown by (23).

These results demonstrate that when the commitment problem is severe enough, transaction quantities and welfare may even fall below the levels attained by a traditional platform. Put differently, the introduction of service tokens is socially beneficial only if a platform has sufficient mechanisms to commit to future policies – otherwise, the introduction of tokens *reduces* welfare.

Proposition 7 is reminiscent of results in financing problems without commitment (Admati et al., 2018; DeMarzo and He, 2021). Typically, an issuer that cannot commit to future issuance effectively competes with its future self: it does not internalize the price impact of current issuance decisions and therefore over-issues relative to the commitment outcome. Goldstein, Gupta, and Sverchkov (2024) demonstrate that if a monopolistic platform can commit to a token's redemption value but not an issuance policy, then token issuance mitigates the platform's tendency to under-supply its services.

Our result is similar in spirit to those in the previous literature, but it is distinct. While it is true that the platform's investors are tempted to extract rents from users by over-issuing new tokens, outcomes with weak commitment would exhibit similar inefficiencies even if new seigniorage were prohibited after $t = 0$. Investors have another tool – fees – that they can use to extract rents in the absence of new token issuance. Instead of issuing new tokens at the time of a governance decision, investors would choose to charge very high transaction fees immediately following the implementation of a new policy. Anticipating this rent extraction, users would not place a high value on tokens and would be reluctant to transact on the platform.

A severe enough lack of commitment power is detrimental to the platform's profits, so investors may therefore seek mechanisms that permit them to commit. Of course, in reality, a platform's founders and investors can use smart contracts to commit to future token issuance, or they could use token retention schemes to incentivize them to pass policies that benefit users. To the extent that such mechanisms are imperfect, though, the next section argues that bundling transaction services with cash flow claims can provide an effective substitute for missing commitment mechanisms.

6 Hybrid tokens

We now consider a platform that issues a single *hybrid* token that bundles transaction services with cash flow claims. Users hold the token for transaction services, whereas investors “stake” the token for cash flows. In this setting, we assume tokens must be staked in order to participate in governance – as a result, only investors vote in governance decisions. This assumption is close to the reality for several DeFi platforms: for example, many proof-of-stake cryptocurrency blockchains allow *only* validators who stake their tokens to vote on policy changes. Our main result is that by issuing a hybrid token, a platform governed by investors can achieve the full-commitment outcome of Section 5.3 even if investors lack commitment power.

6.1 Setup

In this economy, there is a single asset called a *token*. Tokens serve as both the economy’s transaction asset and as its cash flow asset: users can hold tokens for their transaction services, whereas investors can hold tokens to receive dividends. Given that there is only one asset, we let Q_t denote the price of tokens (dispensing with our previous notation Q_t^C, Q_t^T). Furthermore, we denote the total supply of tokens at t by $A_t = A_t^C + A_t^T$, where A_t^C (resp. A_t^T) denotes the quantity of tokens held by investors (users) at t . Henceforth, let $M_t = Q_t A_t$ denote the total market capitalization of tokens, and let $\zeta_t = A_t^C / (A_t^C + A_t^T)$ denote the *fraction* of tokens that are held by investors (“staked”).

As before, users’ transaction services are equal to their real balances of token holdings, $x_{it} = Q_t a_{it}^T$, where a_{it}^T is the quantity of tokens held by user i . Aggregate transaction demand is equal to the real value of tokens held by users, so

$$X_t = Q_t A_t^T = (1 - \zeta_t) M_t.$$

The platform can earn revenues both by charging fees and by issuing tokens. Again, $d\mu_t$ will denote the rate of seigniorage,

$$d\mu_t = \frac{dA_t}{A_t} \in \left[0, \frac{\Delta}{1 + \Delta}\right],$$

so the platform’s seigniorage revenues are $dS_t = M_t d\mu_t$, and total profits are

$$d\Pi_t = (f_t - c) X_t dt + M_t d\mu_t.$$

Dividends are distributed pro rata to the holders of staked tokens, so each staked token pays

a dividend $dD_t = d\Pi_t/\zeta_t A_t$. When a greater fraction ζ_t of tokens are staked, the per-token dividend is lower, all else equal. Investors who stake tokens compete over a fixed pool of dividends, as is common in practice.²³ The fraction of staked tokens is a key equilibrium variable: it adjusts until users and investors place an equal valuation on tokens.

For now, we assume that staking is required to vote, so only investors participate in governance decisions. In Section 7, though, we extend the model to the case in which users may vote on policies as well.

6.2 Equilibrium

We look for a Markov equilibrium in the time s since the most recent policy change in which variables can jump at the time of a governance decision and evolve smoothly thereafter, as in previous sections. The expected return on tokens is equal to the expected growth rate in the market capitalization of tokens net of the issuance rate, analogously to the return on tokens in Lemma 1 in the previous section.

Lemma 3. *When the platform issues hybrid tokens, their expected return satisfies*

$$\frac{1}{ds} \mathbb{E}_s \left[\frac{dQ}{Q} \right] = \frac{\dot{M}_s}{M_s} - \dot{\mu}_s + \lambda \left(\frac{\hat{M}_0}{M_s} (1 - \hat{d}\mu_0) - 1 \right) \quad \text{for } s > 0. \quad (25)$$

Recall that in (25), $\hat{d}\mu_0$ denotes the anticipated jump in the token stock at the time of the *next* governance decision, whereas \hat{M}_0 is the anticipated market capitalization of tokens at that time.

When the platform issues a hybrid token, then both constituencies must be willing to hold tokens at the same price. That is, both users and investors are marginal in the market for tokens. This is the key difference between a hybrid token platform and a platform that issues tokens that offer only transaction services. When investors pass policies, they have an additional incentive to consider how those policies affect users: a policy that is detrimental to users will decrease users' token valuation, reducing the price of the tokens held by investors. We therefore begin by studying how each constituency values tokens.

Users price tokens according to their *service value* (i.e., the present value of tokens' service

²³For example, in proof-of-stake blockchains, the rate at which new blocks are produced (and therefore the rate at which block rewards are issued) typically does not scale linearly with the quantity of staked tokens. When an additional token is staked, then the chance that any other staked token is selected to propose a new block is reduced. Several papers in the literature, such as John, Rivera, and Saleh (2021) and Jermann (2023), make a similar assumption.

flows). The transaction demand condition (2) can be combined with (25) for $s > 0$ to obtain

$$(r + \lambda)M_s = \frac{X_s^{1-\gamma} - f_s X_s}{1 - \zeta_s} + \dot{M}_s - \dot{\mu}_s M_s + \lambda(1 - \hat{d}\mu_0)\hat{M}_0. \quad (26)$$

The way to understand this equation is that users earn aggregate service flows $X_s^{1-\gamma} - f_s X_s$ while holding a stock of tokens worth $(1 - \zeta_s)M_s$ (hence the first term). The remaining terms represent the expected return on tokens.

By contrast, investors value tokens for their dividends when staked. Their optimality condition (4) requires that the token price be equal to tokens' *cash flow value* (i.e., the present value of dividends), which can be written for $s > 0$ as

$$(r + \lambda)M_s = \frac{(f_s - c)X_s + \dot{\mu}_s M_s}{\zeta_s} + \dot{M}_s - \dot{\mu}_s M_s + \lambda(1 - \hat{d}\mu_0)\hat{M}_0. \quad (27)$$

The pricing conditions (26)-(27) then imply a relationship between tokens' *service flows* and their *dividend yields* that must hold so that tokens' service value is equal to their cash flow value:

$$\underbrace{\frac{(f_s - c)X_s + \dot{\mu}_s M_s}{\zeta_s}}_{\text{dividend yield}} = \underbrace{\frac{X_s^{1-\gamma} - f_s X_s}{1 - \zeta_s}}_{\text{service flow}}. \quad (28)$$

How are users' and investors' token valuations kept in line with one another? The fraction of staked tokens ζ_s adjusts in equilibrium. Off-equilibrium, if the dividend yield were higher than the service flow, then tokens would be more attractive for their cash flows than for their transaction services. Investors would therefore purchase and stake additional tokens, increasing ζ_s and driving down the per-token dividend $dD_s = d\Pi_s/\zeta_s A_s$.

Under the hybrid token scheme, then, token prices reflect both their usefulness in transactions and staking dividends. In fact, the market capitalization of tokens is equal to the present value of transaction processing profits plus service flows.

Lemma 4. *Under the hybrid token system, for $s > 0$ the market capitalization of tokens satisfies*

$$M_s = \mathbb{E}_s \left[\int_s^\tau e^{-r(u-s)} \left(X_u^{1-\gamma} - cX_u \right) du + e^{-r(\tau-s)} \hat{M}_0 (1 - \hat{d}\mu_0) \right]. \quad (29)$$

where τ is the time of the next governance decision.

This equation conveys a key intuition: when investors pass a policy, they succeed in increasing the cash flows from their staked tokens only if that policy increases current service flows plus profits. This result may seem counter-intuitive: why is it that they cannot always

increase cash flows by raising fees, as in the case of a platform that issues a service token? Indeed, raising fees may increase the platform's *profits* while decreasing the per-token *dividend yield*. If a policy increases the platform's profits at the expense of a larger *decrease* in tokens' service flows, then users will sell tokens to investors, causing investors to stake more tokens (increasing ζ_s) and decreasing the dividend yield below its initial level.

Of course, despite the fact that fees and the seigniorage rate do not appear in the valuation equation (29), they are not irrelevant in equilibrium. Fees and the seigniorage rate matter because they determine transaction demand through (26). All else equal, higher fees or seigniorage increase users' costs of transacting on the platform. Nevertheless, the pricing equation (29) does imply an important *monetary neutrality* result.

Lemma 5 (Monetary neutrality). *For a given path $\{X_s\}$, the token market capitalization $\{M_s\}$ is independent of $\{\dot{\mu}_s\}$ for $s > 0$, and*

$$M_0(1 - d\mu_0) = M_{0+}, \quad (30)$$

where M_{0+} denotes the market capitalization of tokens immediately after the governance decision at $s = 0$.

Holding transaction quantities fixed, the path of seigniorage is irrelevant for the total market capitalization of tokens. If investors immediately choose to issue a large quantity of tokens $d\mu_0$ when they make a governance decision, (30) implies that this issuance is immediately offset by a drop in the token price, so that the market capitalization of tokens before the issuance, $M_0(1 - d\mu_0)$, is equal to the market capitalization after, M_{0+} . That is, when the platform issues a hybrid token, *seigniorage imposes a capital loss on investors*. This is the main difference from the setting with service tokens: in that setting, the value of investors' assets (shares) would not drop upon the issuance of new tokens.

In governance decisions, only investors are permitted to vote on the level of fees f_s and the seigniorage rate $d\mu_s$. Per (5), investors choose policies to maximize the value of their tokens. The value of tokens outstanding at the time of the governance decision is $M_0(1 - d\mu_0)$, with investors holding some fraction of those. Therefore, investors' governance problem is equivalent to maximizing the market capitalization $M_0(1 - d\mu_0)$ of tokens outstanding, which by Lemmas 4-5 satisfies

$$\begin{aligned} \hat{M}_0(1 - \hat{d}\mu_0) &= \max_{\substack{f_s, d\mu_0, \dot{\mu}_s, \\ X_s, M_s, \zeta_s}} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} (X_s^{1-\gamma} - cX_s) ds + e^{-r\tau} \hat{M}_0(1 - \hat{d}\mu_0) \right] \\ \text{s.t. (26), (28), } &1 - \zeta_s = \frac{X_s}{M_s}, f_s, \dot{\mu}_s \geq 0, d\mu_0 \in [0, \frac{\Delta}{1 + \Delta}]. \end{aligned} \quad (31)$$

6.3 Attaining the full-commitment outcome

We now prove our main result: equilibrium allocations are identical to those attained under the full-commitment outcome in Section 5.3. The logic underlying this result has two steps. First, as is plain from (31), the value of investors' staked tokens is proportional to future aggregate service flows plus platform profits. Second, investors' problem is *time-consistent*: the degree of commitment is irrelevant to equilibrium policies.

Proposition 8. *When the platform issues a hybrid token, investors' governance problem (31) is time-consistent: optimal policies and equilibrium allocations are invariant to the frequency λ of governance decisions.*

Time-consistency implies that no matter how strong investors' commitment power, they will choose policies to maximize the present value of service flows plus profits.

How, exactly, does a hybrid token overcome the time-inconsistency problem? After all, even in the case where the platform issues a token for transactions only, investors' profits depend on users' anticipated service flows. This logic is deceptive, though. When the platform issues a service token, share prices reflect the present value of profits and token prices reflect anticipated future service flows net of inflation. Investors can capture future service flows by issuing new tokens, diluting the value of existing tokens held by users. This is precisely the source of time-inconsistency: investors are tempted to issue a large quantity of new tokens after each governance decision. However, this token issuance is socially detrimental: it causes inflation, raising the cost of transacting on the platform and reducing transaction demand. Investors inefficiently choose to issue tokens and boost share prices at the cost of lowering the value of users' tokens.

When the platform issues a hybrid token, on the other hand, there is a single asset (tokens) held by both constituencies whose value reflects the present value of future profits plus service flows. Regardless of whether investors vote to issue new tokens, the value of their assets will depend on users' future service flows, so they have an incentive to keep service flows high. Indeed, the monetary neutrality result demonstrates that investors cannot benefit from issuing new tokens: any seigniorage revenues are offset by a decrease in the value of their tokens. The fact that investors' tokens lose value when they vote to issue new tokens ($d\mu_0 > 0$) eliminates the time-inconsistency problem. Formally, one can see this fact from (31): seigniorage at $s = 0$ does not enter investors' objective function (unlike investors' problem (18) in the case of a platform that issues service tokens). Intuitively, under the hybrid token system, users are protected from devaluation because they hold the same asset as the investors who govern the platform. With service tokens, by contrast, seigniorage reduces the price of *users' tokens* rather than *investors' shares*.

Since investors pass policies to maximize the present value of service flows plus profits, the equilibrium allocation is precisely the same as the full-commitment outcome of Section 5.3.

Proposition 9. *When governed by investors, a platform with a hybrid token achieves the full-commitment outcome of Section 5.3. Equilibrium policies solve*

$$\max_{\substack{f_t, \dot{\mu}_t, X_t, \\ M_t, \zeta_t}} \int_0^{\infty} e^{-rt} \left(X_t^{1-\gamma} - cX_t \right) dt \quad \text{s.t. } f_t, \dot{\mu}_t \geq 0, \quad (32)$$

$$rX_t = X_t^{1-\gamma} - (f_t + \dot{\mu}_t)X_t + (1 - \zeta_t)\dot{M}_t, \quad (28), \quad 1 - \zeta_t = \frac{X_t}{M_t}.$$

Equilibrium transaction quantities are $X_t = X^{token}$, defined in (21).

The equilibrium is unambiguously more efficient than the equilibrium with a traditional platform but less efficient than the first-best.

We should note that in this benchmark model a hybrid token can enhance efficiency by alleviating time-inconsistency in *seigniorage* policies. However, the benefits of hybrid tokens are actually more general. Online Appendix F studies an extension with investment in which a hybrid token can resolve time-inconsistency in *investment* policies as well as seigniorage.

7 Giving users the right to vote

Up until this point, we have assumed that the platform is governed by investors. Our main results demonstrated how various token designs can align investors' policy preferences with users' (or not). An important feature of the DeFi landscape, however, is the prevalence of platforms that give users the power to vote on policies directly. In this section, we discuss the consequences of giving token-holding users the right to vote in our model.

The model's basic ingredients are unchanged. Governance decisions arrive at a Poisson rate λ . Now, at the time τ of a governance decision, a new policy is chosen by whichever constituency (users or investors) has a voting majority at that time. If investors control the platform, as before, they choose policies that maximize the value of their assets. However, if users control the platform, then they choose a sequence of policies that maximizes the combined value of their tokens plus their expected future infra-marginal rents R_τ ,

$$Q_\tau^T A_{\tau-}^T + R_\tau \quad \text{where} \quad R_\tau = \mathbb{E}_\tau \left[\int_0^{\infty} e^{-rs} \frac{\gamma}{1-\gamma} X_{\tau+s}^{1-\gamma} ds \right],$$

as shown by Proposition 1. We analyze both types of platforms that issue tokens in our model:

- **Service token platform:** Users are allocated voting power in proportion to their token holdings. Without loss of generality, we assume that users hold a majority of voting power²⁴ and choose a fee policy and a seigniorage policy $\{f_s, d\mu_0, \dot{\mu}_s\}$ at the time of each governance decision.
- **Hybrid token platform:** *All* agents, both users and investors, are allocated voting power proportional to their token holdings. In reality, this setup is akin to a DeFi platform that does not limit voting power to those who stake their tokens. At the time of a governance decision, users get to choose fee and seigniorage policies $\{f_s, d\mu_0, \dot{\mu}_s\}$ if they hold a majority of the token stock, whereas investors choose the policy otherwise, as in Section 6. (Control may pass from one constituency to the other over time as their token holdings change.)

In both cases, we look for a Markov equilibrium as before. The analysis leads to optimization problems quite similar to (18) or (31), so we present the main results here but defer the formal analysis to Appendix E.

The main difference is that users seek to pass policies that maximize the value of their transaction assets plus future infra-marginal rents, which is different from investors' objective of maximizing the value of cash flow assets. At the time τ of a governance decision, the value of users' tokens satisfies

$$Q_\tau^T A_{\tau-}^T = \begin{cases} \mathbb{E}_\tau \left[\int_0^\infty e^{-rs} (X_{\tau+s}^{1-\gamma} - (f_{\tau+s} + \dot{\mu}_{t+s})X_{\tau+s}) ds \right] & \text{Service token} \\ (1 - \zeta_{\tau-}) \mathbb{E}_\tau \left[\int_0^\infty e^{-rs} (X_{\tau+s}^{1-\gamma} - cX_{\tau+s}) ds \right] & \text{Hybrid token} \end{cases}$$

The value of a service token is equal to the present value of service flows net of future token dilution, whereas the value of a hybrid token is equal to the present value of service flows plus profits.

Our main result is that giving users the right to vote redistributes economic rents away from investors, but it does not necessarily enhance efficiency.

²⁴Consider a platform that issues service tokens and allocates some voting power to users in proportion to their token holdings, with the remaining votes being allocated to shareholders. Each constituency always votes unanimously, so all that matters is whether token holders are allocated a minority or a majority of votes in the aggregate. When token holders have a minority of voting power, we are back in the case of an investor-controlled platform analyzed in Section 5.

Proposition 10. *Giving users the right to vote can increase user surplus, but it does not necessarily increase total surplus. Moreover, when users control the platform, transaction quantities may be above the first-best level X^{FB} . Specifically,*

1. *When the platform issues service tokens, users vote to set fees and seigniorage equal to zero ($f_t = \mu_t = 0$), and transaction quantities satisfy*

$$X_t = X^{user} \equiv r^{-\frac{1}{\gamma}} > X^{FB} \quad \forall t. \quad (33)$$

2. *When the platform issues a hybrid token, for γ sufficiently small, users have a voting majority in the long run ($t \rightarrow \infty$). In this case, there exists $\zeta^* \in (0, \frac{1}{2}]$ such that*

$$X_t \rightarrow \min \left\{ \left(\frac{1 - \zeta^*}{(1 - \gamma)(1 - \zeta^*) + \gamma} c \right)^{-\frac{1}{\gamma}}, X^{user} \right\} > X^{FB} \quad \text{as } t \rightarrow \infty, \quad (34)$$

where $1 - \zeta^*$ is the fraction of tokens held by users as $t \rightarrow \infty$.

Figure 4 shows that a service token-issuing platform may indeed achieve lower surplus when governed by users rather than investors.

Proponents of DeFi sometimes claim that enfranchising users is key to mitigating inefficiencies. This argument contains a kernel of truth, but it is incomplete. When users are given the right to vote, the platform is more likely to be run in accordance with their interests, so there is less rent extraction. Users unambiguously benefit from the right to vote. However, users have incentives to run the platform in their own favor, to investors' detriment. In particular, users do not bear the platform's costs of operation – those are instead borne by investors. Users are therefore willing to pass policies that increase their infra-marginal rents at the expense of platform profits. This desire to increase their own rents is why users pass policies that lead to inefficiently high transaction quantities. A redistribution of voting power can increase welfare only to the extent that it brings the *median* voter's preferences closer in line with maximizing total surplus (Hart and Moore, 1998). Just as there is no reason to presume that consumer cooperatives are more efficient than stockholder corporations, then, there is no reason to presume that a user-governed platform is more efficient than an investor-governed platform.

Note, however, that the policies preferred by users depend on whether the platform issues a service token or a hybrid token. In particular, users vote for policies that result in (weakly) lower transaction quantities under the hybrid token scheme – that is, they vote for higher fees and platform profits. Why would users vote to give some profits to investors? In this case, the value of users' tokens depends not only on their own service flows but also on the expected




Figure 4: A comparison of total welfare (for a platform that issues service tokens) under governance by each constituency. The figure is plotted with parameters $\gamma = 0.5, c = 0.06, r = 0.035$. The blue curve plots total surplus as a function of transaction demand X . Transaction demand under investor governance is X^{token} , whereas transactions under user governance are X^{user} . The red lines mark the level of surplus in each case.

dividends that will be paid out to investors, as explained in Section 6. Consequently, users do not want to reduce the platform’s profits too much.

Notice also that unlike a service token platform, a hybrid token platform that issues tokens with voting rights does not always lead to user control, even in the long run. If investors control the platform, they may pass policies that let them continue to keep a voting majority of tokens. This may involve setting fees *above* the level that maximizes token prices, thereby limiting transaction quantities (and thus users’ token holdings). Investors will be willing to relinquish control to users only if they do not expect users to distort transaction quantities too far above their preferred level X^{token} . Users, in turn, have less of an incentive to distort transaction quantities upwards when their infra-marginal rents are small, i.e., when their utility function is not too concave. This is the case when γ is small, hence the result in Proposition 10 that users control the platform in the long run for γ close to zero.

The key to the hybrid token scheme is that it introduces an asset whose value is determined by the welfare of *both* constituencies. The fact that the value of investors’ token holdings depends on users’ welfare moderates their desire to extract rents from users. Similarly, if users govern the platform, their desired policies are moderated by the fact that they hold tokens whose value depends on the platform’s profits. It is this alignment of preferences, not a redistribution of voting rights, that mitigates inefficiencies in platform policies.

8 Discussion of assumptions

Here we comment on the model’s main assumptions. The model has three key features that shape the results. First, there are distinct groups of agents who interact on the platform, generating scope for conflicts of interest. Second, the agents who run the platform lack the ability to commit to a sequence of policies. Third, we limit the types of contracts that can be written between users and investors.

Conflicts of interest: In our model, there are two distinct groups of agents: users who enjoy the platform’s services and investors who hold all cash flow claims on the platform. This separation of users and investors may seem unnatural, since the users of a platform’s service are typically not prevented from holding the platform’s equity.

We make two comments on this issue. First, while this stark assumption helps to illustrate the logic of our model, it could be substantially relaxed without significantly altering the main results. What really matters for our results is that there must be some heterogeneity in preferences: there must be some agents who are more interested in the platform’s services and others who are more natural investors (even if investors also enjoy the platform’s services to some extent). For instance, investors could be interpreted as individuals who are more

patient or who have deeper pockets. In the context of DeFi, those who stake tokens and earn cash flows often need some technical aptitude or computational resources – for example, blockchain “validators” typically use powerful hardware to certify transactions. As long as there is some heterogeneity in preferences, there will be conflicts of interest: natural users will prefer low fees and greater investment in the platform’s technology, whereas natural investors will prefer to maximize profits at users’ expense (see Hart and Moore, 1998; or Bakos and Halaburda, 2022; who emphasize a similar theme).

Second, even within the context of our model, our main results would mostly go through if we were to permit users to hold cash flow claims on the platform as well. In this case, there would be a continuum of equilibria. There would still be an equilibrium in which investors hold all of the cash flow claims on the platform. However, for each $k \in (0, 1]$, there would also be an equilibrium in which users hold a fraction k of all cash flow claims, with investors holding the remaining fraction $1 - k$. The inefficiencies that we highlight would continue to arise in this setting: investors would vote for policies that maximize profits at users’ expense, whereas users would try to maximize their rents at investors’ expense. The *only* efficient equilibrium would be the case in which users hold *all* cash flow claims on the platform ($k = 1$) – when investors sit out entirely, then conflicts of interest become irrelevant. By focusing on the case in which investors hold all cash flow claims and voting rights, our model addresses concerns about the possible “centralization” of DeFi platforms via a concentration of token holdings among non-users.

Lack of commitment: While a lack of commitment is of course central to our main results, we consider it natural to assume that investors cannot make a fully binding commitment to a sequence of policies. In reality, it is quite common for digital platforms to change their terms of service and fees. Even in the context of DeFi, despite the fact that some policies can be hard-coded at the time a platform is founded, it is possible for the rules of almost any platform to be amended in some way. While some blockchains require a “hard fork” (i.e., a split in the blockchain) to re-write the initially set rules, other platforms allow for rules to be updated in arbitrary ways by a simple vote among stakers. For example, the DeFi platforms MakerDAO, Curve, and Uniswap operate in this way. The possibility of amending the rules, moreover, is typically considered to be a desirable feature rather than a detriment to the platform’s viability. However, our main results of course also encapsulate the case of full commitment as $\lambda \rightarrow 0$ (which could be achieved, e.g., if the platform is governed by a smart contract whose rules can never be superseded).

Limited contracting: The last key assumption in our model is limited contracting between users and investors. In particular, we assume away the Coase-style solution in which users and investors collectively agree on a contract that rewards investors for passing a socially

beneficial sequence of policies. We view this as a reasonable restriction for several reasons. It may be costly for users to coordinate and collectively bargain with investors. Moreover, if a contract that anticipates all possible future contingencies were feasible, then it would not even be necessary to have platform governance in the first place – all decisions could simply be encoded in the initial contract. We take the view of Grossman and Hart (1986): if the platform’s ownership structure is to play an essential role, contracts must be incomplete to some extent.

9 Conclusion

We develop a general model of platform governance that is flexible enough to capture both traditional platforms as well as platforms that issue tokens with some combination of transaction services, cash flow claims, and voting rights.

A traditional platform extracts rents from its users by setting fees above its marginal costs. This discourages users from transacting and distorts volumes downwards, below the first-best level.

The issuance of *service tokens* can partially align shareholders’ policy preferences with those of users, as long as shareholders are able to commit to future policies ex ante. If the platform passes policies that benefit users, they will be willing to pay a greater price to purchase tokens. Hence, if investors commit to pass such policies, they can reap large seigniorage revenues. However, if investors lack the ability to commit, this mechanism no longer aligns preferences: after selling tokens to users, investors will again be tempted to extract rents from them by inflating away their value.

The platform can overcome the commitment problem by issuing a *hybrid token* that bundles transaction services with cash flow claims. The key idea is that by issuing a single asset whose value reflects *both* constituencies’ welfare, investors can be disincentivized from extracting rents. If they pass policies that are detrimental to users, the token price will fall, hurting investors as well.

Giving users the right to vote, on the other hand, redistributes economic rents to users but does not necessarily enhance welfare. Just like investors, users are self-interested and pass policies that increase their rents at investors’ expense.

References

- [1] M. Sockin, W. Xiong, Decentralization through tokenization, *Journal of Finance* 78 (1) (2023) 247–299. doi:10.1111/jofi.13192.

- [2] I. Goldstein, D. Gupta, R. Sverchkov, Utility tokens as a commitment to competition, *Journal of Finance* Available at SSRN: <https://ssrn.com/abstract=3484627> (In Press). doi:10.2139/ssrn.4248492.
- [3] J. Chod, E. Lyandres, A theory of icos: Diversification, agency, and information asymmetry, *Management Science* 67 (10) (2021) 5969–5989. doi:10.1287/mnsc.2020.3754.
- [4] J. Li, W. Mann, Digital tokens and platform-building, available at SSRN: <https://ssrn.com/abstract=3088726> (2017). doi:10.2139/ssrn.3088726.
- [5] L. W. Cong, J. Li, N. Wang, Token-based platform finance, *Journal of Financial Economics* 144 (3) (2022) 972–991. doi:10.1016/j.jfineco.2021.10.002.
- [6] J. Lee, C. Parlour, Consumers as financiers: Consumer surplus, crowdfunding, and initial coin offerings, *Review of Financial Studies* 35 (3) (2022) 1105–1140. doi:10.1093/rfs/hhab058.
- [7] Y. Li, S. Mayer, Money creation in decentralized finance: A dynamic model of stablecoin and crypto shadow banking, available at SSRN: <https://ssrn.com/abstract=3757083> (2020). doi:10.2139/ssrn.3757083.
- [8] A. d’Avernas, V. Maurin, Q. Vandeweyer, Can stablecoins be stable?, *Becker-Friedman Institute Working Paper No. 2022-131* (2022). doi:10.2139/ssrn.4226027.
- [9] Y. You, K. Rogoff, Redeemable platform currencies, *Review of Economic Studies* 90 (2) (2023) 975–1008. doi:10.1093/restud/rdac028.
- [10] N. Stokey, R. Lucas, *Recursive Methods in Economic Dynamics*, Harvard University Press, 1989.
- [11] S. Gryglewicz, S. Mayer, E. Morellec, Optimal financing with tokens, *Journal of Financial Economics* 142 (3) (2021) 1038–1067. doi:10.1016/j.jfineco.2021.05.004.
- [12] S. Grossman, O. Hart, The costs and benefits of ownership: A theory of vertical and lateral integration, *Journal of Political Economy* 94 (4) (1986) 691–719. doi:10.1086/261404.
- [13] Y. Bakos, H. Halaburda, Will blockchains disintermediate platforms? limits to credible decentralization in daos, available at SSRN: <https://ssrn.com/abstract=4221512> (2022). doi:10.2139/ssrn.4221512.

- [14] J. Han, J. Lee, T. Li, Dao governance, available at SSRN: <https://ssrn.com/abstract=4346581> (2023). doi:10.2139/ssrn.4346581.
- [15] J. Gan, G. Tsoukalas, S. Netessine, Decentralized platforms: Governance, tokenomics, and ico design, *Management Science* 69 (11) (2023) 6667–6683. doi:10.1287/mnsc.2021.02076.
- [16] R. Coase, Durability and monopoly, *Journal of Law and Economics* 15 (1) (1972) 143–149. doi:10.1086/466731.
- [17] R. Coase, The nature of the firm, *Economica* 4 (16) (1937) 386–405. doi:10.1017/cbo9780511817410.009.
- [18] O. Williamson, Transaction-cost economics: The governance of contractual relations, *The Journal of Law and Economics* 22 (2) (1979) 233–261. doi:10.1086/466942.
- [19] H. Hansmann, Ownership of the firm, *The Journal of Law, Economics, and Organization* 4 (2) (1988) 267–304. doi:10.1093/oxfordjournals.jleo.a036953.
- [20] P. Aghion, J. Tirole, Formal and real authority in organizations, *Journal of Political Economy* 105 (1) (1997) 1–29. doi:10.1086/262063.
- [21] G. Tsoukalas, B. H. Falk, Token-weighted crowdsourcing, *Management Science* 66 (9) (2020) 3843–3859. doi:10.1287/mnsc.2019.3515.
- [22] A. Benhaim, B. H. Falk, G. Tsoukalas, Balancing power in decentralized governance: Quadratic voting under imperfect information, available at SSRN: <https://ssrn.com/abstract=4416748> (2023). doi:10.2139/ssrn.4416748.
- [23] O. Hart, J. Moore, Cooperatives vs. outside ownership, nBER Working Paper 6421 (1998). doi:10.3386/w6421.
- [24] D. Levit, N. Malenko, E. Maug, Trading and shareholder democracy, *Journal of Finance* 79 (1) (2024) 257–304. doi:10.1111/jofi.13289.
- [25] M. Sidrauski, Rational choice and patterns of growth in a monetary economy, *American Economic Review* 57 (2) (1967) 534–544.
URL <https://www.jstor.org/stable/1821653>
- [26] R. Feenstra, Functional equivalence between liquidity costs and the utility of money, *Journal of Monetary Economics* 17 (2) (1986) 271–291. doi:10.1016/0304-3932(86)90032-2.

- [27] L. W. Cong, Y. Li, N. Wang, Tokenomics: Dynamic adoption and valuation, *Review of Financial Studies* 34 (3) (2021) 1105–1155. doi:10.1093/rfs/hhaa089.
- [28] A. Admati, P. DeMarzo, M. Hellwig, P. Pfleiderer, The leverage ratchet effect, *Journal of Finance* 73 (1) (2018) 145–198. doi:10.1111/jofi.12588.
- [29] P. DeMarzo, Z. He, Leverage dynamics without commitment, *Journal of Finance* 76 (3) (2021) 1195–1250. doi:10.1111/jofi.13001.
- [30] U. Jermann, A macro-finance model for proof-of-stake ethereum, available at SSRN: <https://ssrn.com/abstract=4335835> (2023). doi:10.2139/ssrn.4335835.
- [31] K. John, T. Rivera, F. Saleh, Equilibrium staking levels in a proof-of-stake blockchain, available at SSRN: <https://ssrn.com/abstract=3965599> (2021). doi:10.2139/ssrn.3965599.
- [32] J. Bena, S. Zhang, Token-based decentralized governance, data economy and platform business model, available at SSRN: <https://ssrn.com/abstract=4248492> (2022). doi:10.2139/ssrn.4248492.
- [33] B. Biais, C. Bisière, M. Bouvard, C. Casamatta, A. Menkveld, Equilibrium bitcoin pricing, *Journal of Finance* 78 (2) (2023) 967–1014. doi:10.1111/jofi.13206.

A Model

A.1 Agents' optimization problems

Users: Users' optimization problem is

$$U_i = \max_{x_{it}, a_{it}, c_{it}} \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \left(\frac{x_{it}^{1-\gamma}}{1-\gamma} - f_t x_{it} \right) dt + dc_{it} \right]$$

s.t. $Q_t^T da_{it} + dc_{it} = 0$, $x_{it} \leq Q_t^T a_{it}$, a_{i0} given.

The only individual state variable for user i is asset holdings a_{it} .

Using integration by parts, it is possible to rearrange an agent's lifetime utility from consumption of dollars:

$$\begin{aligned} \int_0^\infty e^{-rt} dc_{it} &= - \int_0^\infty e^{-rt} Q_t^T da_{it} \\ &= -e^{-rt} Q_t^T a_{it} \Big|_0^\infty + \int_0^\infty e^{-rt} (a_{it} dQ_t^T - r a_{it} Q_t^T dt) \\ &= Q_0^T a_{i0} + \int_0^\infty e^{-rt} a_{it} (dQ_t^T - r Q_t^T dt). \end{aligned}$$

Using this equation, agent i 's optimization problem can be reformulated as

$$U_i = \max_{x_{it}, a_{it}} Q_0^T a_{i0} + \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \left(\frac{x_{it}^{1-\gamma}}{1-\gamma} - f_t x_{it} \right) dt - a_{it} (r Q_t^T dt - \mathbb{E}_t[dQ_t^T]) \right] \quad (\text{A.1})$$

s.t. $x_{it} \leq Q_t^T a_{it}$, a_{i0} given.

Users' problem then reduces to a sequence of static optimizations over (x_{it}, a_{it}) .

The optimal x_{it} solves

$$\max_{x_{it}, a_{it}} \left(\frac{x_{it}^{1-\gamma}}{1-\gamma} - f_t x_{it} \right) dt - (r Q_t^T dt - \mathbb{E}_t[dQ_t^T]) a_{it} \quad \text{s.t.} \quad x_{it} \leq Q_t^T a_{it},$$

so

$$x_{it}^{-\gamma} dt = f_t dt + \left(r dt - \mathbb{E}_t \left[\frac{dQ_t^T}{Q_t^T} \right] \right), \quad (\text{A.2})$$

as claimed in (2).

We also derive users' value functions, since those will be important in determining users' preferences over sequences of policies. Users' first-order condition (2) implies that the flow utility they receive at time t is

$$\left(\frac{X_t^{1-\gamma}}{1-\gamma} - f_t X_t\right) dt - X_t \left(r dt - \mathbb{E}_t \left[\frac{dQ_t^T}{Q_t^T} \right] \right) = \frac{\gamma}{1-\gamma} X_t^{1-\gamma} dt.$$

Then, plugging this expression into the right-hand side of (A.1),

$$U_i = Q_0^T A_{0-}^T + R_0, \quad \text{where } R_0 \equiv \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \frac{\gamma}{1-\gamma} X_t^{1-\gamma} dt \right], \quad (\text{A.3})$$

as claimed in Proposition 1.

Investors' lifetime utility can be derived with analogous calculations. Investor j 's problem is

$$U_j = \max_{a_{jt}, c_{jt}} \mathbb{E}_0 \left[\int_0^\infty e^{-rt} dc_{jt} \right] \quad \text{s.t.} \quad Q_t^C (da_{jt} - a_{jt} dD_t) + dc_{jt} = 0, \quad a_{j0} \text{ given.}$$

Integration by parts yields

$$U_j = \max_{a_{jt}} Q_0^C a_{j0} + \mathbb{E}_0 \left[\int_0^\infty e^{-rt} a_{jt} \left(dD_t + dQ_t^C - rQ_t^C dt \right) \right] \quad \text{s.t.} \quad a_{j0} \text{ given.}$$

The first-order condition implies that at an optimum (4) holds, and

$$U_j = Q_0^C A_{0-}^C. \quad (\text{A.4})$$

Proof of Proposition 1. The result is immediate from (A.3) and (A.4). \square

A.2 The first-best

In this section, we formally lay out and solve the planner's problem:

$$\hat{V}_0^P = \max_{X_s} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \left(\frac{X_s^{1-\gamma}}{1-\gamma} - cX_s \right) ds + e^{-r\tau} \hat{V}_0^P \right]. \quad (\text{A.5})$$

We now prove Proposition (2), which characterizes the first-best allocation.

Proof of Proposition 2. First, note that the planner's problem is time-consistent by standard arguments (see Stokey and Lucas, 1989): it is a standard dynamic programming problem with control X_s and no state variables. Therefore, the optimal policy solves the corresponding sequence problem. The sequence problem is

$$\max_{X_t} \int_0^{\infty} e^{-rt} \left(\frac{X_t^{1-\gamma}}{1-\gamma} - cX_t \right) dt \quad (\text{A.6})$$

with first-order condition $X_t^{-\gamma} = c$, which immediately implies (8). \square

B Traditional platform

In this section, we prove our main results about the traditional scheme in Proposition 3.

Proof of Proposition 3. Standard arguments (Stokey and Lucas, 1989) imply that investors' governance problem (11) is time-consistent: it is a standard dynamic programming problem with control f_s and no state variables. Therefore, investors' problem can be written as the sequence problem (12). We substitute the transaction demand constraint (10) into investors' profits to obtain

$$(f_t - c)X_t = X_t^{1-\gamma} - (c + r)X_t$$

and use this result to write the sequence problem as.

$$\max_{X_t} \int_0^{\infty} e^{-rt} \left(X_t^{1-\gamma} - (c + r)X_t \right) dt.$$

The first-order condition is

$$(1 - \gamma)X_t^{-\gamma} = c + r.$$

Therefore, the optimal transaction quantity is X^{trad} given by (14), which is clearly less than X^{FB} in (8). Since X^{FB} is the first-best quantity of transactions, it is immediate that welfare is lower than the first-best with a traditional platform. \square

C Service tokens

In this section, we prove the results in Section 5.

Proof of Lemma 1. In equilibrium, conditional on no governance decision, transactions and the stock of tokens grow at rates $\frac{\dot{X}_s}{X_s}$ and $\dot{\mu}_s$, respectively. When a governance decision takes place, transactions jump to \hat{X}_0 and the stock of tokens immediately grows by $\hat{d}\mu_0$. Then, using the fact that $Q_t^T = \frac{X_t}{A_t^T}$, we have

$$\frac{dQ_t^T}{Q_t^T} = \left(\frac{\dot{X}_s}{X_s} - \dot{\mu}_s \right) ds + \left(\frac{\hat{X}_0}{X_s} (1 - \hat{d}\mu_0) - 1 \right) dJ_s.$$

where dJ_s denotes a Poisson process that arrives at rate λ (representing the next governance decision). Taking expectations of both sides, we obtain the equation in Lemma 1. \square

Proof of Lemma 2. This result is immediate from investors' governance problem (18). The rate of seigniorage $d\mu_0$ at $s = 0$ can be increased up to the maximum limit $d\mu_0 = 1$ without affecting profits or the constraint (17) at any future date. \square

Proof of Proposition 4. Using the result of Lemma 2, we have

$$\begin{aligned} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} d\Pi_s \right] &= \mathbb{E}_0 \left[\frac{\Delta}{1+\Delta} X_0 + \int_0^\tau e^{-rs} (f_s + \dot{\mu}_s - c) X_s ds \right] \\ &= \frac{\Delta}{1+\Delta} X_0 + \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \left(X_s^{1-\gamma} - (c+r+\lambda) X_s + \dot{X}_s + \lambda \hat{X}_0 \left(1 - \frac{\Delta}{1+\Delta} \right) \right) ds \right] \\ &= \frac{\Delta}{1+\Delta} X_0 + \int_0^\infty \lambda e^{-\lambda\tau} \int_0^\tau e^{-rs} \left(X_s^{1-\gamma} - (c+r+\lambda) X_s + \dot{X}_s + \lambda \frac{1}{1+\Delta} \hat{X}_0 \right) ds d\tau \\ &= \frac{\Delta}{1+\Delta} X_0 + \int_0^\infty e^{-(r+\lambda)s} \left(X_s^{1-\gamma} - (c+r+\lambda) X_s + \dot{X}_s + \lambda \frac{1}{1+\Delta} \hat{X}_0 \right) ds \\ &= \frac{\Delta}{1+\Delta} X_0 + \int_0^\infty e^{-(r+\lambda)s} \left(X_s^{1-\gamma} - (c+r+\lambda) X_s + \lambda \frac{1}{1+\Delta} \hat{X}_0 \right) ds \\ &\quad + e^{-(r+\lambda)s} X_s \Big|_0^\infty + \int_0^\infty e^{-rs} (r+\lambda) X_s ds \\ &= -\frac{1}{1+\Delta} X_0 + \int_0^\infty e^{-(r+\lambda)s} \left(X_s^{1-\gamma} - c X_s + \lambda \frac{1}{1+\Delta} \hat{X}_0 \right) ds, \end{aligned}$$

as desired. The second equality uses (17). The fifth equality integrates $\int_0^\infty e^{-(r+\lambda)s} \dot{X}_s ds$ by parts. \square

Proof of Propositions 5 - 7. We characterize the properties of equilibrium in both the strong commitment and weak commitment regime. We begin by taking a Lagrangian approach to investors' governance problem. Following Proposition (4), investors' problem can be written as

$$\hat{V}_0^I = \max_{f_s, \dot{\mu}_s, \dot{X}_s} - \frac{\Delta}{1 + \Delta} X_0 + \int_0^\infty e^{-(r+\lambda)s} \left(X_s^{1-\gamma} - cX_s + \lambda \left(\frac{1}{1 + \Delta} \hat{X}_0 + \hat{V}_0^I \right) - \psi_s \left((r + \lambda)X_s - X_s^{1-\gamma} + (f_s + \dot{\mu}_s)X_s - \dot{X}_s - \lambda \frac{1}{1 + \Delta} \hat{X}_0 \right) - \chi_s^f f_s - \chi_s^\mu \dot{\mu}_s \right) ds$$

The Euler-Lagrange conditions for $s > 0$ are

$$\begin{aligned} (f_s) : 0 &= \psi_s X_s + \chi_s^f; \\ (\dot{\mu}_s) : 0 &= \psi_s X_s + \chi_s^\mu; \\ (X_s) : \dot{\psi}_s &= (1 + \psi_s)(1 - \gamma)X_s^{-\gamma} - c - \psi_s(f_s + \dot{\mu}_s) \end{aligned}$$

We will conjecture an equilibrium in which fees $f_s = f$, seigniorage $\dot{\mu}_s = \dot{\mu}$, transaction quantities $X_s = X$, and the Lagrange multipliers are constant as well. Define λ^* as

$$\lambda^* = \frac{c}{1 - \gamma} - r.$$

We proceed by analyzing two distinct cases.

Case 1: Strong commitment ($\hat{\lambda} \leq \lambda^*$). We conjecture that in this case, the Lagrange multiplier ψ is equal to zero. The Euler-Lagrange condition for X_s then implies that transaction quantities must satisfy

$$(1 - \gamma)X_s^{-\gamma} = c \Rightarrow X_s = \left(\frac{c}{1 - \gamma} \right)^{-\frac{1}{\gamma}} \equiv X^{token} \text{ for } s > 0.$$

This quantity of transactions is greater than X^{trad} but below the first-best X^{FB} .

To confirm that this is indeed an equilibrium allocation, we must show that there exist non-negative levels of fees f and seigniorage $\dot{\mu}$ that support it. The transaction demand condition (i.e., the condition corresponding to the multiplier ψ_s) yields

$$f + \dot{\mu} = \frac{c}{1 - \gamma} - (r + \hat{\lambda}). \quad (\text{C.1})$$

Then it is possible that $f, \dot{\mu} \geq 0$ only if $\hat{\lambda} \leq \lambda^*$. Our conjectured equilibrium is indeed an

equilibrium, then, if and only if $\hat{\lambda} \leq \lambda^*$. This concludes the proof of Proposition 5.

Note, moreover, that this argument proves Proposition 6. Any combination $(f, \dot{\mu})$ satisfying (C.1) implements the same equilibrium.

Case 2: Weak commitment ($\hat{\lambda} > \lambda^*$). In this case, we conjecture and verify that fees and seigniorage are equal to zero, $f = \dot{\mu} = 0$. With a constant Lagrange multiplier ψ , we have

$$(1 + \psi)(1 - \gamma)X^{-\gamma} = c \Rightarrow X = \left(\frac{c}{(1 + \psi)(1 - \gamma)} \right)^{-\frac{1}{\gamma}}.$$

The transaction demand condition (17) pins down X :

$$X = \left(\frac{r + \hat{\lambda}}{1 - \gamma} \right)^{-\frac{1}{\gamma}}.$$

When $\hat{\lambda} > \lambda^*$, clearly, the Lagrange multiplier ψ is indeed non-zero, justifying the fact that both fees and seigniorage are at their constraints. This concludes the proof of Proposition 7.

□

D Hybrid tokens

D.1 Equilibrium

In this section, we prove general results about the equilibrium with a hybrid token.

Proof of Lemma 4. Multiply (26) by $1 - \zeta_s$, multiply (27) by ζ_s , and add the two resulting equations together to obtain

$$(r + \lambda)M_s = X_s^{1-\gamma} - cX_s + \dot{M}_s + \lambda\hat{M}_0(1 - \hat{d}\mu_0).$$

Imposing the transversality condition $\lim_{s \rightarrow \infty} e^{-(r+\lambda)s}M_s = 0$, this differential equation can be solved forward to obtain

$$\begin{aligned} M_s &= \int_0^\infty e^{-(r+\lambda)u} \left(X_{s+u}^{1-\gamma} - cX_{s+u} + \lambda\hat{M}_0(1 - \hat{d}\mu_0) \right) du \\ &= \mathbb{E}_s \left[\int_0^\tau e^{-r(u-s)} \left(X_u^{1-\gamma} - cX_u \right) + e^{-r(\tau-s)} \hat{M}_0(1 - \hat{d}\mu_0) \right], \end{aligned}$$

as desired. The second line uses the fact that for any function $h(u)$,

$$\int_0^{\infty} e^{-(r+\lambda)u} h(u) du = \mathbb{E} \left[\int_0^{\tau} e^{-ru} h(u) du \right],$$

where τ denotes the time of a Poisson event that arrives at rate λ . □

Proof of Lemma 5. That $\{M_s\}$ is independent of $\dot{\mu}_s$ for $s > 0$ is immediate from Lemma 4. It remains to prove that $M_0(1 - d\mu_0) = M_{0+}$. If the token quantity does not jump at $s = 0$, this result is trivial, so we focus on the case in which a positive quantity of tokens is issued at the time of the governance decision. In this case, the following sequence of events takes place:

- Investors announce the quantity of new tokens $A_{0+} - A_0$ that will be issued (where A_0 denotes the quantity of tokens outstanding before the governance decision);
- The new tokens are sold in competitive markets at price Q_0 ;
- Token holders decide whether to stake or hold tokens for transactions;
- Seigniorage profits $(A_{0+} - A_0)Q_0$ are paid out to the holders of staked tokens;
- The token price adjusts to a new ex-dividend level Q_{0+} .

Using this notation, we have $d\mu_0 = (A_{0+} - A_0)/A_{0+}$, $M_0 = Q_0 A_{0+}$, $M_{0+} = Q_{0+} A_{0+}$, and $dS_0 = M_0 d\mu_0$.

Note that if $d\mu_0 > 0$, then (28) implies that investors hold the entire stock of tokens at $s = 0$, i.e., $\zeta_0 = 1$. Then $dD_0 = dS_0/A_{0+}$, and investors' pricing equation for tokens (4) yields

$$\begin{aligned} 0 &= \frac{dS_0}{A_{0+}} + \mathbb{E}_0[dQ_0] \\ &= M_0 d\mu_0 + Q_{0+} A_{0+} - Q_0 A_{0+} \\ &= M_0 d\mu_0 + M_{0+} - M_0 \\ \Rightarrow M_0(1 - d\mu_0) &= M_{0+}. \end{aligned}$$

□

D.2 Attaining the commitment outcome

Next, we prove our main results about the hybrid token system with investor governance.

Proof of Proposition 8. Problem (31) is time-consistent by standard arguments. When written in Lagrangian form, it can be viewed as a typical Bellman equation with controls $(f_s, \dot{\mu}_s, X_s, M_s, \zeta_s)$. Therefore, optimal policies are independent of the times at which governance decisions took place: the policy chosen at time τ for $s \geq \tau$ is precisely the same as the policy chosen for $s \geq \tau$ at $t = 0$. Hence, optimal policies and equilibrium allocations are independent of the frequency λ of governance decisions. \square

Proof of Proposition 9. Given that Problem (31) is a time-consistent Bellman equation, the optimal policy must solve the corresponding sequence problem (32). Note that the objective function is precisely the same as in (20). To show that all of the results carry over, we merely need to demonstrate that there exists some set of controls that implements the full-commitment allocation. Investors can simply choose to charge a constant fee $f_t = c/(1-\gamma) - r$ (as in the case with commitment) and issue no new tokens, $\dot{\mu}_t = 0$ for all t . Since policies are constant over time, token prices will be as well, so users' optimality condition (2) implies

$$X_t = X^{token} = \left(\frac{c}{1-\gamma} \right)^{-\gamma} \quad \forall t.$$

Hence, we have shown that equilibrium transaction quantities are precisely the same as in Proposition 5, and just as in that proposition, the equilibrium allocation is more efficient than that with a traditional platform but less efficient than the first-best. \square

E Giving users the right to vote

In this section, we prove the results in Proposition 10. The Proposition deals with two cases – a setting with a platform that issues service tokens and one with a platform that issues hybrid tokens. The following two subsections address each in turn.

E.1 A service token platform

As discussed in Section 7, when the platform permits service token holders to vote, we can restrict attention to the case in which users control a majority of voting power for all t . Hence, users choose their most-preferred policy in each governance decision.

We look for a recursive equilibrium in which quantities may jump at the time of a governance decision but evolve smoothly thereafter, with the same notation as in the benchmark model. Users maximize their lifetime utility, which per (6) is equal to the value of their tokens

plus the present value of their infra-marginal rents, which will be denoted by

$$\hat{R}_0 = \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \frac{\gamma}{1-\gamma} X_s^{1-\gamma} ds + e^{-r\tau} \hat{R}_0 \right].$$

The value of users' tokens at the time $s = 0$ of a governance decision satisfies

$$Q_0^T A_{0-}^T = X_0(1 - d\mu_0),$$

since users' initial token holdings are only a fraction $1 - d\mu_0$ of the initial token stock if new token issuance at $s = 0$ is $d\mu_0$. Then, users' governance problem can be written recursively in Lagrangian form as

$$(1 - d\mu_0)\hat{X}_0 + \hat{R}_0 = \max_{f_s, \dot{\mu}_s, d\mu_0, X_s} \int_0^\infty e^{-(r+\lambda)s} \left((1 - d\mu_0)(X_s^{1-\gamma} - (f_s + \dot{\mu}_s)X_s) + \frac{\gamma}{1-\gamma} X_s^{1-\gamma} \right. \\ \left. + \lambda((1 - \hat{d}\mu_0)\hat{X}_0 + \hat{R}_0) + \chi_s^f f_s + \chi_s^\mu \dot{\mu}_s + \psi_s((r + \lambda)X_s - X_s^{1-\gamma} + (f_s + \dot{\mu}_s)X_s - \dot{X}_s - \lambda(1 - \hat{d}\mu_0)\hat{X}_0) \right) ds, \quad (\text{E.1})$$

subject to the additional constraint $d\mu_0 \in [0, \frac{\Delta}{1+\Delta}]$.

Users will always choose to set initial seigniorage $d\mu_0$ to zero – unanticipated seigniorage is just a transfer to investors. The Euler-Lagrange conditions are

$$(f_s) : (\psi_s - 1)X_s = \chi_s^f; \\ (\dot{\mu}_s) : (\psi_s - 1)X_s = \chi_s^\mu; \\ (X_s) : (r + \lambda)\psi_s = X_s^{-\gamma} - (f_s + \dot{\mu}_s) + \psi_s((r + \lambda) - (1 - \gamma)X_s^{-\gamma} + f_s + \dot{\mu}_s) + \dot{\psi}_s.$$

It is simple to see from this problem that users' utility is strictly increasing in transaction quantities X_s and decreasing in fees and seigniorage. Hence, there is no advantage to users of setting fees or seigniorage above their minimum allowable levels. Users thus set fees $f_s = 0$ for all s and choose zero seigniorage, $\dot{\mu}_s = 0$ for all s . By setting $f_s = \dot{\mu}_s = 0$ and conjecturing a constant multiplier $\psi_s = \psi$, we obtain

$$0 = (1 - (1 - \gamma)\psi)X_s^{-\gamma} \Rightarrow \psi = \frac{1}{1 - \gamma}.$$

From there, one can find suitable values for the Lagrange multipliers χ_s^f, χ_s^μ using the Euler-

Lagrange conditions above. Transaction quantities satisfy (17), so

$$X_s^{1-\gamma} = rX_s \Rightarrow X_s = X^{user} \equiv r^{-\frac{1}{\gamma}},$$

as desired.

E.2 A hybrid token platform

Before proving the specific results in this section, we characterize equilibrium in a slightly more general setting and then show how to derive the results from this general characterization.

In the case where users can vote, there are two state variables in a Markov equilibrium: the time s since the most recent governance decision and the fraction of tokens ζ held by investors just prior to the decision. The fraction of tokens held by investors is relevant because it determines which constituency gets to choose the platform's policies.

We consider an equilibrium in which token quantities never jump. Lemma 5 guarantees that this is without loss of generality. The equilibrium quantities are then $\{\hat{f}_s(\zeta), \hat{\mu}_s(\zeta), \hat{X}_s(\zeta), \hat{M}_s(\zeta), \hat{\zeta}_s(\zeta), \hat{R}_s(\zeta)\}$. As in the main body of the paper, several conditions must be satisfied in equilibrium. There is the token demand condition

$$(r + \lambda)X_s = X_s^{1-\gamma} - (f_s + \dot{\mu}_s)X_s + \zeta_s(\dot{M}_s + \lambda\hat{M}_0(\zeta_s)), \quad (\text{E.2})$$

the token pricing condition (29), the condition determining ζ_s , which can be re-written as

$$\frac{X_s^{1-\gamma} - (f_s + \dot{\mu}_s)X_s}{1 - \zeta_s} = \frac{(f_s + \dot{\mu}_s - c)X_s}{\zeta_s}, \quad (\text{E.3})$$

and the Bellman equation determining the present value of infra-marginal rents,

$$\hat{R}_0(\zeta) = \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \frac{\gamma}{1-\gamma} X_s^{1-\gamma} ds + e^{-r\tau} \hat{R}_0(\zeta_\tau) \right]. \quad (\text{E.4})$$

We consider a decision-maker who, starting in state ζ , chooses policies to maximize

$$\begin{aligned} \max_{\substack{f_s, \dot{\mu}_s \\ X_s, \dot{M}_s, \zeta_s}} \mathbb{E}_0 \left[\int_0^\tau e^{-rs} \left(\left(1 + \frac{\gamma}{1-\gamma} b(\zeta)\right) X_s^{1-\gamma} - cX_s \right) ds + e^{-r\tau} (\hat{M}_0(\zeta_\tau) + b(\zeta)\hat{R}_0(\zeta_\tau)) \right] \\ \text{s.t. (29), (E.2), (E.3), } f_s, \dot{\mu}_s \geq 0, \end{aligned} \quad (\text{E.5})$$

where the function b is decreasing, continuous, and differentiable (except for perhaps at

$\zeta = \frac{1}{2}$).

We conjecture an equilibrium in which quantities are constant between governance decisions, e.g., $X_s = X$, $M_s = M$, and so on. To avoid confusion, we denote the fraction of tokens held by investors at the time of the governance decision by ζ , and we let $\zeta' = \zeta_s$ denote the fraction held by investors *after* the decision takes place.

Under the conjecture that all quantities are constant between governance decisions, the decision-maker's problem reduces to

$$\begin{aligned} \max_{f, \dot{\mu}, X, M, \zeta'} & \frac{(1 + \frac{\gamma}{1-\gamma}b(\zeta))X^{1-\gamma} - cX}{r + \lambda} + \frac{\lambda}{r + \lambda}(\hat{M}_0(\zeta') + b(\zeta)\hat{R}_0(\zeta')) \\ \text{s.t. (29), (E.2), (E.3), } & f_s, \dot{\mu}_s \geq 0. \end{aligned} \quad (\text{E.6})$$

We proceed in several steps.

Step 1: In equilibrium, $\hat{M}_0(\zeta)$ ($\hat{R}_0(\zeta)$) is weakly monotonically increasing (decreasing) in ζ . Notice that in the type of equilibria we consider, the policy problem faced by the decision-maker can be written as

$$\max_{f, \dot{\mu}, X, M, \zeta'} M_0 + b(\zeta)R_0 \quad \text{s.t. (29), (E.2), (E.3), } f_s, \dot{\mu}_s \geq 0,$$

where $b(\zeta)$ is a decreasing function of ζ and

$$\begin{aligned} M_0 &= \frac{X^{1-\gamma} - cX + \lambda\hat{M}_0(\zeta')}{r + \lambda}, \\ R_0 &= \frac{\frac{\gamma}{1-\gamma}X^{1-\gamma} + \lambda\hat{R}_0(\zeta')}{r + \lambda}. \end{aligned}$$

The decision-maker chooses policies to maximize a weighted average of token values and users' infra-marginal rents, where the weight on rents is decreasing in ζ . Clearly, then, higher ζ implies that at an optimum, M_0 will be larger and R_0 will be smaller.

Step 2: In equilibrium, the fraction of tokens ζ' held by investors is an increasing function of the transaction quantity X .

Observe that the constraints (E.2) and (E.3) depend on $(f, \dot{\mu})$ only through their sum $f + \dot{\mu}$. Given a transaction quantity X , we can therefore view these constraints as two equations in the two unknowns $f + \dot{\mu}$ and ζ' . The transaction demand condition (E.2) implies an inverse relationship between X and total transfers $f + \dot{\mu}$ from users to investors. The condition (E.3) determining ζ' implies an increasing relationship between transfers $f + \dot{\mu}$ and ζ' . Therefore, the fraction of tokens ζ' held by investors is a decreasing function of the transaction quantity X .

Step 3: Let ζ_k be the fraction of tokens held by investors at the time of the k -th governance decision. Then the sequence $\{\zeta_k\}$ converges to some ζ^ .*

We begin by showing that there exists a monotonically increasing mapping $G(\cdot) : [0, 1] \rightarrow [0, 1]$ such that $\zeta_{k+1} = G(\zeta_k)$. There are two terms in the objective of Problem (E.6) whose weights are decreasing in ζ_k : $\frac{\gamma}{1-\gamma} \frac{X^{1-\gamma}}{r+\lambda}$ and $\hat{R}_0(\zeta_{k+1})$. The former is clearly increasing in X , but the latter is increasing in X as well because $\hat{R}_0(\cdot)$ is increasing in ζ_{k+1} by Step 1 of this argument and ζ_{k+1} is decreasing in X by Step 2. The weights on the other two terms, $-\frac{cX}{r+\lambda}$ and $\hat{M}_0(\zeta_{k+1})$, are constant in ζ_k . Both terms are decreasing in the current transaction quantity X . Thus, the optimal X from the decision-maker's perspective is decreasing in ζ_k , the fraction of tokens held by investors at the time of the most recent governance decision. Applying Step 2 again, the fraction of tokens ζ_{k+1} held by investors at the *next* governance decision is an increasing function of ζ_k .

Convergence follows from the fact that G is a monotonically increasing function that is continuous on the intervals $[0, \frac{1}{2}]$ and $(\frac{1}{2}, 1]$. There are two cases to consider.

Case 1: $G(\frac{1}{2}) \leq \frac{1}{2}$. If $\zeta_0 \in [0, \frac{1}{2}]$, Brouwer's fixed point theorem immediately implies that the sequence ζ_k converges, since the interval $[0, \frac{1}{2}]$ is mapped to itself. If $\zeta_0 > \frac{1}{2}$, then there are again two cases: either $\lim_{\zeta \rightarrow +\frac{1}{2}} G(\zeta) > \frac{1}{2}$ or $\lim_{\zeta \rightarrow +\frac{1}{2}} G(\zeta) < \frac{1}{2}$. In the first case, G maps some interval $[\frac{1}{2} + \epsilon, \frac{1}{2}]$ to itself, so ζ_k converges somewhere in that interval. In the second case, if ζ_k does not converge to a steady state in $(\frac{1}{2}, 1]$, then it eventually enters the interval $[0, \frac{1}{2}]$, after which point it converges in that interval.

Case 2: $G(\frac{1}{2}) > \frac{1}{2}$. The proof in this case is analogous.

We are now ready to move on to the more specific case in our model, where

$$b(\zeta) = \begin{cases} \frac{1}{1-\zeta} & \zeta \leq \frac{1}{2} \\ 0 & \zeta > \frac{1}{2} \end{cases}$$

We proceed with the proof of Proposition 10.

Proof of Proposition 10. Section E.1 proves the result in the case of a platform that issues service tokens. For a platform that issues hybrid tokens, we first characterize transaction quantities in the long run when users control the platform. Then, we show that γ being close to zero is a sufficient condition to ensure users will control the platform in the long run.

When users control the platform in the long run, ζ converges to some limit $\zeta^* \leq \frac{1}{2}$. In this limit, the governance problem (E.6) reduces to maximizing the present value of token

values plus a multiple $\frac{1}{1-\zeta}$ of rents,

$$\max_{X, f, \dot{\mu}} \frac{(1 + \frac{1}{1-\zeta^*} \frac{\gamma}{1-\gamma}) X^{1-\gamma} - cX}{r} \quad \text{s.t.} \quad X^{-\gamma} = r + f + \dot{\mu}, \quad f, \dot{\mu} \geq 0.$$

The first-order condition yields

$$X = \min \left\{ \left(\frac{1 - \zeta^*}{(1 - \gamma)(1 - \zeta^*) + \gamma} c \right)^{-\frac{1}{\gamma}}, r^{-\frac{1}{\gamma}} \right\},$$

as desired.

Next, we show that users control the platform as $t \rightarrow \infty$ whenever γ is close enough to zero. Our argument above demonstrates that equilibrium outcomes with a hybrid token are equivalent to those in the following game:

- If a governance decision occurs at t_k , then whichever constituency controls the platform (users or investors) chooses a level of transactions X_{t_k} that will be maintained for the interval $[t_k, t_{k+1}]$, where t_{k+1} denotes the time of the next governance decision.
- There exists a threshold level of transactions \hat{X} such that users control the platform at t_{k+1} (i.e., $\zeta_{t_{k+1}} \geq \frac{1}{2}$) if and only if $X_{t_k} \geq \hat{X}$.²⁵
- Going forward from time t_k , investors' payoff is

$$V_I(t_k) = \zeta_k \mathbb{E}_{t_k} \left[\sum_{n=0}^{\infty} (e^{-rt_{k+n}} - e^{-rt_{k+n+1}}) \frac{X_{t_{k+n}}^{1-\gamma} - cX_{t_{k+n}}}{r} \right],$$

whereas users' payoff is

$$V_U(t_k) = \mathbb{E}_{t_k} \left[\sum_{n=0}^{\infty} (e^{-rt_{k+n}} - e^{-rt_{k+n+1}}) \left((1 - \zeta_k) \frac{X_{t_{k+n}}^{1-\gamma} - cX_{t_{k+n}}}{r} + \frac{\gamma}{1 - \gamma} \frac{X_{t_{k+n}}^{1-\gamma}}{r} \right) \right].$$

Note that investors' optimal outcome is for the level of transactions to be

$$X^{token} = \left(\frac{c}{1 - \gamma} \right)^{-\frac{1}{\gamma}}$$

at all times. Indeed, if investors control the platform at $t = 0$, and

$$\zeta^{token} = 1 - (1 - \gamma) \frac{r}{c} \geq \frac{1}{2},$$

²⁵This is because ζ_{t_k} is a decreasing function of X_{t_k} , as shown previously.

then investors choose transaction level X^{token} at $t = 0$, giving them control of the platform at t_1 , at which point they choose transaction level X^{token} again (and at all future dates of governance decisions).

However, this equilibrium is not possible if

$$2r \geq \frac{c}{1-\gamma}.$$

If investors choose transaction level X^{token} at $t = 0$, then $\zeta_0 \leq \frac{1}{2}$, so users will control the platform at t_1 . In this case, one of two equilibrium outcomes is possible. Either

1. Investors choose to maintain control of the platform at all dates, choosing a transaction quantity

$$\hat{X} = (2r + c)^{-\frac{1}{\gamma}},$$

leading to $\zeta_k = \frac{1}{2}$ for all k , or

2. Investors choose to let users control the platform and choose some $X_{t_0} \in [\hat{X}, X^{token}]$. In this latter case, users will control the platform at t_1 . Moreover, since they always prefer higher transaction quantities than investors and ζ_k is decreasing in X_k , users will continue to choose transaction quantities larger than \hat{X} , leading to $\zeta_k < \frac{1}{2}$ for all k , so users control the platform starting from t_1 .

Hence, a sufficient condition to ensure that investors will relinquish control to users is:

$$\log\left(\frac{\hat{X}^{1-\gamma} - c\hat{X}}{r}\right) \leq \log\left(\frac{X^{token^{1-\gamma}} - cX^{token}}{r + \lambda}\right). \quad (\text{E.7})$$

The left-hand side is the (log) payoff earned by investors if they maintain control of the platform by setting $X_{t_k} = \hat{X}$ for all k . The right-hand side is a lower bound on the expected payoff earned by investors if they set $X_0 = X^{token}$ and permit users to control the platform thereafter.

In the limit $\gamma \rightarrow 0$, inequality (E.7) can be rewritten as

$$\begin{aligned} 0 &\leq \lim_{\gamma \rightarrow 0} \log\left(\frac{\gamma}{1-\gamma} \frac{c}{r + \lambda} \times \left(\frac{c}{1-\gamma}\right)^{-\frac{1}{\gamma}}\right) - \log\left(2 \times (2r + c)^{-\frac{1}{\gamma}}\right) \\ &= \lim_{\gamma \rightarrow 0} \log \gamma + \log \frac{c}{2(r + \lambda)} + \frac{1}{\gamma} \log\left(1 + \frac{2r}{c}\right). \end{aligned}$$

Note that as $\gamma \rightarrow 0$, the right-hand side goes to infinity. This is because the function $\frac{1}{\gamma}$ diverges at a faster rate than $\log \gamma$. Therefore, for small enough γ , investors prefer to give

users control of the platform in the first period, and users control the platform in the long run.

□