

On the Stability of Token Systems

Süleyman Kerimov*

October 14, 2026

Abstract

Token systems offer a market solution for non-monetary economies, where monetary transactions are undesirable or restricted. In such systems, agents request and provide services (trade favors) over time with possibly different rates, and tokens are used to pay for service provision. The stability and effectiveness of token systems depend on whether token accumulation and debts of agents can be controlled. This control becomes more challenging when the service availability is minimal. When for every service request only a few agents can provide the service, the resulting low liquidity makes it harder to reward agents for their contributions.

We characterize necessary and sufficient conditions for the stability of token systems under the minimum token selection rule, where the agent with the minimum amount of tokens is chosen to provide service among the available providers. For homogeneous token systems in which each service costs the same amount of tokens, these conditions require that for any subset of agents, their total service request rate is strictly greater than their unavoidable token earning rate. We show that if any of these conditions is violated, no stationary provider selection rule can stabilize the system.

We extend our characterization to systems where services are heterogeneous, i.e., the cost of each service may be random and requester dependent, and we provide explicit tail bounds on stationary token endowments of agents. Our results shed light on how much asymmetry and heterogeneity a token system can tolerate to sustain cooperation. We discuss several market design levers for the social planner, while arguing how to relate stability to full participation in an ex-ante commitment game, where agents are risk sensitive.

1 Introduction

Numerous non-monetary economies such as kidney exchange (e.g., see [Ashlagi et al. \(2026\)](#)), peer-to-peer networks (e.g., see [Vishnumurthy et al. \(2003\)](#)), food bank donations (e.g., see

*Virani Undergraduate School of Business and Jones Graduate School of Business, Rice University, Houston TX, USA, 77005 kerimov@rice.edu.

Prendergast (2022)), and home exchange networks (e.g., see Goedde (2025)) use token systems to facilitate resource allocation and service exchanges. While there are differences in how these tokens are used in practice (food banks use tokens to bid for donated food, whereas hospitals receive tokens for their contributions to a patient-donor pair exchange pool, and they spend tokens to receive a patient-donor pair that enables an exchange within their own program), token systems serve a common purpose: coordinating the allocation of resources and increasing market efficiency when monetary transactions are not allowed due to repugnancy, or ethical and legal constraints.

In this paper, we study a trading favors setting, where agents request and provide services (trade favors) over time with possibly different rates, and tokens are used to pay for service provision. Agents provide services or contribute resources to earn tokens so that they can spend these tokens in the future when needed. Market thickness, in the sense that for any service request, there are many agents who can provide the service, helps operating token systems with more flexibility, as an agent who receives a service today need not be the one who eventually returns the favor to the same agent, and these two exchanges do not need to happen at the same time.

The effectiveness of a token system depends on whether token accumulation and debts of agents can be controlled. An agent who repeatedly provides service eventually accumulates many tokens, and if she cannot find enough opportunities to spend these tokens, then she may lose interest to contribute to the system, since tokens have no monetary value outside of the system. This creates a challenge for the social planner to incentivize such agents to contribute to the system and reward them back by fulfilling their service requests. Similarly, allowing agents who are in debt to spend tokens helps trading favors, but this makes controlling such imbalances crucial to avoid unraveling in the system.

We address these issues through a notion of stability that requires two conditions. First, stability requires the amount of tokens owned or owed by agents to be uniformly bounded with high probability. Second, it requires the market to clear in finite expected time, i.e., agents return to their initial endowment continuously over time. While these conditions do not address strategic environments directly, we show that these conditions yield a full participation in an appropriately defined strategic game.

Intuitively, when agents request and provide services with similar rates, the token system should behave well, since no agent should be more likely to accumulate tokens or be in a great debt compared to other agents. With asymmetric agents, opportunities to earn tokens no longer

need to align with the rate of spending tokens, and agents who frequently request service but do rarely provide service can destabilize the token system. Moreover, if the service availability is limited, the planner may be forced to select an agent who already owns a large amount of tokens as a service provider, which makes it difficult to move tokens from *richer* agents to agents who are in debt.

Kidney exchange illustrates these aforementioned concerns. Kidney exchange platforms help arranging swaps between incompatible patient-donor pairs, where compatibility between a patient and a donor determined by blood- and tissue-type compatibility (Roth et al. 2007). In the United States, hospitals can choose whether to participate in these platforms to enroll their incompatible patient-donor pairs. Generally speaking, these pairs can be categorized into two groups (easy-to-match and hard-to-match) based on their blood types and sensitivity levels. The majority of the pairs in the exchange pool are hard-to-match, and hospitals differ in size and in the composition of easy-to-match and hard-to-match pairs.

In order to incentivize hospitals to participate and contribute to the platform, token-based economies have been introduced, for example, by the National Kidney Registry, where the Center Liquidity Contribution program scores hospitals based on their contribution to the exchange pool, and hospitals with positive scores are given matching priority.¹ Intuitively, a hospital who contributes an easy-to-match pair to the pool earns tokens, and a hospital who finds a match to its hard-to-match pair through the exchange pool is deducted tokens. Interestingly, token values differ across patient-donor pair types to reflect their marginal contribution to the pool (e.g., see Agarwal et al. (2019)).² Here, the practical concern is that whether the difference in the composition of patient-donor pair types and size of the hospitals, as well as rare matching opportunities can allow sustaining cooperation between hospitals and increase efficiency? As noted in Agarwal et al. (2019), the kidney exchange market is highly fragmented, and most exchanges are arranged within hospitals instead of through platforms. Moreover, hospitals often match easy-to-match pairs with each other (instead of sharing it with the platform), which is not aligned with market efficiency.

These observations motivate us to raise the following question for token systems:

How much asymmetry between service requests and service provisions, and heterogeneity in service costs can a token system tolerate, and can a simple provider selec-

¹See Ashlagi and Roth (2021) and <https://www.kidneyregistry.com/wp-content/uploads/2022/02/nkr-report-2021-Q4-v10.pdf>.

²For a more extensive discussion of hospital participation in kidney exchange, see Agarwal et al. (2019), Ashlagi and Roth (2021), Ashlagi et al. (2026).

tion rule stabilize the token system?

Overview of results. We study a stylized model that isolates how asymmetry in contributions and needs of agents, heterogeneity in token charges, and limited service availability jointly determine whether the token system can remain stable. At each time period, one agent requests service, a limited number of agents are available to provide the requested service, and the service requester pays the service cost to the service provider in tokens. For the service provider selection rule, we study the minimum token selection rule (e.g., see [Johnson et al. \(2014\)](#), [Ashlagi et al. \(2026\)](#)), where an available agent with the minimum amount of tokens is chosen to provide service. We say that the token system is stable if the process that captures the token endowments of agents is positive recurrent: agents return to their initial endowment in finite time in expectation, and the number of tokens owned or owed by the agents are uniformly bounded with high probability in the long-run.

For homogeneous token systems, where each service request costs the same amount of tokens, we first characterize the maximal stability region in terms of market primitives (Theorem 1). In words, to ensure stability, the service request rate of every group of agents must be strictly larger than their unavoidable token earning rate: a group of agents earn tokens unavoidably when all available service providers belong to the group. We show that these conditions are sufficient under the minimum token selection rule and necessary under every feasible stationary provider selection rule. We then complement our stability result by providing a tail bound on the token distribution: we show that in the long-run, the probability of agents having more than M or less than $-M$ tokens is at most $O(1/M)$ (Theorem 2), generalizing a result established for symmetric agents in [Ashlagi et al. \(2026\)](#). We then study heterogeneous token systems, where service costs are possibly random and requester dependent, and we prove analogous stability and tail bound results for this case (Theorems 3 and 4).

Finally, we discuss several market design levers for the social planner, and we study a strategic environment that connects our stability notion to participation of agents in an ex-ante commitment game with mean-variance preferences (see §5.2).

Related literature. A stream of papers study token-based economies with strategic agents to incentivize trading favors (e.g., see [Möbius \(2001\)](#), [Kash et al. \(2012, 2015\)](#)). [Johnson et al. \(2014\)](#) study the minimum token selection rule that we also analyze in our paper, and they show that always-trade strategy can be supported in equilibrium when refusing to provide service can be punished and the market is sufficiently thick, in the sense that a constant fraction of agents

are available to provide each requested service. Following their work, [Bo et al. \(2018\)](#) study large markets and show an existence of an ϵ -Nash equilibrium while making punishments infeasible.

The closest papers to our work are [Ashlagi and Kerimov \(2019\)](#) and [Ashlagi et al. \(2026\)](#), which study token systems under minimal service availability. These papers establish that when agents are symmetric, i.e., agents request and become available to provide service with the same rates, having only two random available service providers for every service request suffices to stabilize the token system under the minimum token selection rule. [Ashlagi et al. \(2026\)](#) also examine an asymmetric setting through numerical studies based on kidney exchange data. Their study validates rare matching opportunities in kidney exchange, and it illustrates high concentration of agents' (hospitals) token endowments throughout the horizon. Our results complement this numerical evidence by theoretically characterizing the maximal stability region for token systems that exhibit asymmetry in service request and provision rates, as well as heterogeneous service costs.

A recent work [Hu et al. \(2025\)](#) study dynamic token systems, where they explicitly account for dynamic entry and exit of agents. Their work proposes mechanisms that allow agents to buy, sell, or recycle tokens, and they identify conditions under which always-trade strategy can be supported as a Markov perfect equilibrium. Motivated by a large home exchange platform, [Goedde \(2025\)](#) study a pricing problem in a token-based economy, where supply is endogenous. Their study also notes that agents can lose incentive to contribute to the system when the opportunities to spend tokens on the platform are limited. Another recent study [Chen et al. \(2026\)](#) analyze the problem of incentivizing resource pooling among strategic servers in the context of queueing systems via simple token-based mechanisms.

Notation. We use $\mathbb{Z}_{\geq 0}$ and $\mathbb{Z}_{\geq 1}$ to denote the set of nonnegative integers and the set of positive integers, respectively. We use $\mathbb{E}_{\pi}[\cdot]$ and $\mathbb{P}_{\pi}(\cdot)$ to denote the expectation and probability with respect to a distribution π , respectively.

2 Model

We study a non-monetary economy in which agents trade services (favors) using tokens (also known as artificial currency, scrips, or chips). There is a finite set of agents $\mathcal{A} = \{1, 2, \dots, n\}$, $n \geq 2$. Time is discrete and indexed by $t \in \mathbb{Z}_{\geq 0}$. We denote the number of tokens agent $i \in \mathcal{A}$ has at time t by $s_i^t \in \mathbb{Z}$. *Token distribution vector* at time t is denoted by $\mathbf{s}^t = (s_i^t)_{i \in \mathcal{A}}$. For convenience, we assume that $\mathbf{s}^0 = \mathbf{0}$.

Market primitives. At each time period t , an agent is drawn independently from a prob-

ability distribution $P = (p_i)_{i \in \mathcal{A}} > \mathbf{0}$ to be a *service requester*, denoted by k_t . The *service availability density* is denoted by $d \in \mathbb{Z}_{\geq 1}$, and the set of agents who are *available* to provide the requested service is determined by drawing d candidates independently and with replacement from a probability distribution $Q = (q_i)_{i \in \mathcal{A}} > \mathbf{0}$.³ Hence, at most d agents can provide the requested service.⁴ We assume that the service requester and available service provider draws are independent within and across time periods.

Token exchange dynamics. Let $\mathcal{B}^t \subset \mathcal{A}$ be the set of agents who are available to provide the requested service at time t . Note that $|\mathcal{B}^t| \leq d$. After realizing the service requester k_t and the service provider $l_t \in \mathcal{B}^t$ (we explain the service provider selection rule momentarily below), the service requester pays $c^t \in \mathbb{Z}_{\geq 1}$ many tokens to the service provider, and the rest of the agents remain unaffected by this exchange, i.e.,

$$\mathbf{s}^{t+1} = \mathbf{s}^t - c^t \mathbf{e}_{k_t} + c^t \mathbf{e}_{l_t} \quad \text{for all } t \in \mathbb{Z}_{\geq 0},$$

where $\mathbf{e}_i$ is the i 'th standard basis. Let $\mathbf{c} = (c^t)_{t \geq 0}$. *Token system* is the tuple $(n, P, Q, d, \mathbf{c})$. The state $\mathbf{s}^t$ is observed before the nature picks the service requester k_t , the service cost c^t , and the service provider l_t , and after the requester pays the provider, the system transitions to $\mathbf{s}^{t+1}$. The pairs (k_t, c^t) are assumed to be i.i.d. across time periods and independent of $\mathcal{B}^t$. As we discuss in §4, we allow c^t to depend on k_t .

Note that the total number of tokens endowed by the agents remains fixed throughout the horizon, i.e., $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all $t \in \mathbb{Z}_{\geq 0}$. Therefore, the state space on which $\mathbf{s}^t$ evolves is

$$\mathcal{S} := \left\{ \mathbf{s} \in \mathbb{Z}^n : \sum_{i \in \mathcal{A}} s_i = 0 \right\}.$$

For example, when $c^t = 1$ for all $t \in \mathbb{Z}_{\geq 0}$, any state in $\mathcal{S}$ is reachable from the initial state $\mathbf{0}$. Therefore, for simplicity, we assume that $\mathbb{P}(c^t = 1) > 0$ for all $t \in \mathbb{Z}_{\geq 0}$.⁵ Importantly, an agent who has negative amount of tokens can still request a service as we treat tokens as a *reputation* device: for example, when services cost the same amount of tokens, agents with $s_i^t > 0$ have

³In this paper, we are assuming that the market is thin and it has low service availability, i.e., d is a fixed constant and independent of n . One can capture thicker markets by setting $d = kn$ where $0 < k \leq 1$ (e.g., see Johnson et al. (2014), Kash et al. (2015), Bo et al. (2018)).

⁴Sampling with replacement is assumed to simplify the exposition. We also note that an agent can provide service to herself, which captures self-service cases such as in kidney exchange markets, patient-donor pairs from the same hospital can be matched internally when the patient-donor pair pools are merged.

⁵Equivalently, there exists $i \in \mathcal{A}$ such that $\mathbb{P}\{c^t = 1 \mid k_t = i\} > 0$. Without this assumption, let $g = \gcd\{c \geq 1 : \mathbb{P}(c^t = c) > 0\}$, where $\gcd$ denotes the greatest common divisor. Then the states that are reachable become $\mathcal{S}_g := \{\mathbf{s} \in (\mathbb{Z})^n : \sum_{i \in \mathcal{A}} s_i = 0\}$.

a larger net contribution to the system than agents with $s_i^t < 0$, and we study whether token systems can be effective to sustain cooperation through a notion of stability (see Definition 1).

Service provider selection rules. We consider service provider selection rules that are time homogeneous and only select an available agent as a service provider at each time period. More formally, let $\mathcal{A}^d$ denote the collection of nonempty subsets of $\mathcal{A}$ with size at most d , and let $\Delta(\mathcal{A})$ denote the probability simplex over $\mathcal{A}$. A provider selection rule ψ is a time-independent mapping

$$\psi : \mathcal{S} \times \mathcal{A} \times \mathbb{Z}_{\geq 1} \times \mathcal{A}^d \rightarrow \Delta(\mathcal{A}),$$

where $\psi_l(s, k, c, \mathcal{B}) := \mathbb{P}(l_t = l \mid \mathbf{s}^t = s, k_t = k, c^t = c, \mathcal{B}^t = \mathcal{B})$ is the conditional probability that $l \in \mathcal{A}$ is the service provider at time t , given that at time t the token distribution vector is $\mathbf{s} \in \mathcal{S}$, the service requester is $k \in \mathcal{A}$, the cost of the service is $c \in \mathbb{Z}_{\geq 1}$, and the set of available agents is $\mathcal{B} \in \mathcal{A}^d$. Let Ψ be the set of all feasible stationary service provider selection rules, where $\psi \in \Psi$ if and only if:

1. $\psi_l(s, k, c, \mathcal{B}) \geq 0$ for all $k, l \in \mathcal{A}$, $\mathbf{s} \in \mathcal{S}$, $c \in \mathbb{Z}_{\geq 1}$, and $\mathcal{B} \in \mathcal{A}^d$.
2. $\sum_{l \in \mathcal{A}} \psi_l(s, k, c, \mathcal{B}) = 1$ for all $k \in \mathcal{A}$, $\mathbf{s} \in \mathcal{S}$, $c \in \mathbb{Z}_{\geq 1}$, and $\mathcal{B} \in \mathcal{A}^d$.
3. $\psi_l(s, k, c, \mathcal{B}) = 0$ for all $k \in \mathcal{A}$, $\mathbf{s} \in \mathcal{S}$, $c \in \mathbb{Z}_{\geq 1}$, $\mathcal{B} \in \mathcal{A}^d$, and $l \notin \mathcal{B}$.

The third feasibility condition ensures that the service provider is always an available agent. In this paper, unless stated otherwise, we focus on *minimum token selection rule* as the service provider rule (e.g., see Johnson et al. (2014), Ashlagi et al. (2026)). Under the minimum token selection rule, an available agent who has the least amount of tokens is selected as the service provider, i.e., $l_t \in \arg \min_{i \in \mathcal{B}^t} s_i^t$, where any tie is broken uniformly at random, i.e., the rule chooses an agent uniformly in $\{i \in \mathcal{B}^t : s_i^t = \min_{j \in \mathcal{B}^t} s_j^t\}$ at any time t . We say that agents are *symmetric* if $p_i = q_i = 1/n$ for all $i \in \mathcal{A}$, and agents are *asymmetric* otherwise. We say that the token system is *homogeneous* if $c^t = 1$ for all $t \in \mathbb{Z}_{\geq 0}$, capturing markets where every service has the same value, and this value is normalized to be 1. We say that the token system is *heterogeneous* otherwise, which we analyze in §4.

2.1 Stability

Under any feasible selection rule $\psi \in \Psi$, since $p_i, q_i > 0$ for all $i \in \mathcal{A}$, $\mathbb{P}(c^t = 1) > 0$ for all t , and there is a strictly positive probability that an agent can serve herself, the Markov chain $(\mathbf{s}^t : t \geq 0)$ is an irreducible and aperiodic Markov chain on $\mathcal{S} = \{\mathbf{s} \in \mathbb{Z}^n : \sum_{i \in \mathcal{A}} s_i = 0\}$. We

say that the token system $(n, P, Q, d, \mathbf{c})$ is *stable* if the Markov chain $\mathbf{s}^t$ is positive recurrent. One can view the existence of a unique stationary distribution, which is implied by positive recurrence, as a necessity for implementing an effective token system. This is because stability implies $\mathbb{E}[T_0] < \infty$, where $T_0 = \inf\{t \geq 1 : \mathbf{s}^t = \mathbf{0}\}$. Hence, if the token system is stable, then agents return to their initial endowment in finite time in expectation, where we informally interpret this as *fairness* is achieved. Stability also implies the following *boundedness* condition: for every $\epsilon > 0$, there exists $M, t_0 < \infty$ such that for all $t \geq t_0$, we have $\mathbb{P}(\max_{i \in \mathcal{A}} |s_i^t| > M) < \epsilon$, so that agents' token balances do not diverge to positive or negative infinity. In order to study how concentrated the token system is, we will analyze stability through the following definition.

Definition 1 (Stability). The token system $(n, P, Q, d, \mathbf{c})$ is stable if the Markov chain $\mathbf{s}^t$ is positive recurrent. Equivalently, there exists a function $f : \mathbb{Z}_{\geq 0} \rightarrow [0, 1]$ and $t_0 < \infty$ such that $\lim_{M \rightarrow \infty} f(M) = 0$, and $\mathbb{P}(|s_i^t| > M) < f(M)$ for all $i \in \mathcal{A}$, for all $M \in \mathbb{Z}_{\geq 0}$, and for all $t \geq t_0$.

Indeed, it is ideal to show the existence of an $f(\cdot)$ that converges to 0 fast, which in turn will imply that the token distribution is highly concentrated around the initial point $\mathbf{s}^0 = \mathbf{0}$. It is worth mentioning that while this definition of stability does not capture strategic environments directly, where agents decide whether or not to provide service, stability is a sufficient condition to sustain cooperation in a risk aversion model in which agents take actions via mean-variance preferences (e.g., see §5.2 and Ashlagi et al. (2026)[Section 5]).

2.2 Groups

For any nonempty group of agents $\mathcal{G} \subsetneq \mathcal{A}$, let $p(\mathcal{G}) = \sum_{i \in \mathcal{G}} p_i$ be the *total service request rate* of $\mathcal{G}$, $q(\mathcal{G}) = \sum_{i \in \mathcal{G}} q_i$ be the *total service availability rate* of $\mathcal{G}$ per one available candidate draw, and $S^t(\mathcal{G}) = \sum_{i \in \mathcal{G}} s_i^t$ be the total token balance of $\mathcal{G}$ at time t . Let $\Delta S^t(\mathcal{G}) = S^{t+1}(\mathcal{G}) - S^t(\mathcal{G})$. Note that, we have

$$\Delta S^t(\mathcal{G}) = c^t(\mathbb{1}\{l_t \in \mathcal{G}\} - \mathbb{1}\{k_t \in \mathcal{G}\}). \quad (1)$$

For homogeneous token systems, when $d = 1$, it is easy to see that the relevant groups that determine the evolution of $\mathbf{s}^t$ consists of singletons, where for each $i \in \mathcal{A}$, s_i^t behaves like a lazy random walk with drift $q_i - p_i$ and violates positive recurrence depending on whether $q_i = p_i$ (which implies null recurrence) or $q_i \neq p_i$ (which implies no stationary distribution exists).⁶ Therefore no homogeneous token system is stable when $d = 1$. We will complement this result by showing that the same result holds for heterogeneous token systems.

⁶See also Ashlagi et al. (2026)[Section 2.1].

When $d \geq 2$, in the evolution of $\mathbf{s}^t$, the relevant groups no longer consist solely of singletons, and the behavior of the token system under the minimum token selection rule depends on the ordering and the gaps between the amount of tokens groups of agents endow. To that end, we define the following useful notions.

For a fixed state $\mathbf{s} \in \mathcal{S}$, consider a permutation $\sigma : \{1, 2, \dots, n\} \rightarrow \mathcal{A}$ such that

$$s_{\sigma(1)} \leq s_{\sigma(2)} \leq \dots \leq s_{\sigma(n)}. \quad (2)$$

Note that such a permutation is not necessarily unique and any ties can be broken arbitrarily. Define the *token gaps* $\delta_m := s_{\sigma(m+1)} - s_{\sigma(m)}$ for all $1 \leq m \leq n-1$, and let $\mathcal{U}_m := \{\sigma(m+1), \sigma(m+2), \dots, \sigma(n)\}$ be the m 'th *upper group*. Note that for all $i \in \mathcal{A}$, we have

$$s_i = s_{\sigma(1)} + \sum_{m=1}^{n-1} \delta_m \mathbb{1}\{i \in \mathcal{U}_m\}, \quad (3)$$

which implies

$$\max_{i \in \mathcal{A}} s_i - \min_{i \in \mathcal{A}} s_i = \sum_{m=1}^{n-1} \delta_m. \quad (4)$$

We define the *range* of token amounts by $R(\mathbf{s}) := \max_{i \in \mathcal{A}} s_i - \min_{i \in \mathcal{A}} s_i$. These notions are illustrated in Figure 1.




Figure 1: An example of a token distribution $\mathbf{s} \in \mathcal{S}$ of 9 agents, where $\sum_{i=1}^9 s_i = 0$. The fifth upper group $\mathcal{U}_5$ consists of four agents $\{\sigma(6), \sigma(7), \sigma(8), \sigma(9)\}$. Since $\min_{i \in \mathcal{A}} s_i = -4$ and $\max_{i \in \mathcal{A}} s_i = 3$, we have $R(\mathbf{s}) = 7$. The sequence of token gaps is $(\delta_m)_{m=1}^8 = (2, 0, 1, 1, 1, 1, 1, 0)$.

Finally, for homogeneous token systems, we define the *service slack* as $\epsilon(\mathcal{G}) := p(\mathcal{G}) - q(\mathcal{G})^d$, where $q(\mathcal{G})^d = \mathbb{P}(\mathcal{B}^t \subset \mathcal{G})$ is the probability that the available agents who can provide service at

any time belong to group $\mathcal{G}$. Let $\epsilon := \min_{\emptyset \neq \mathcal{G} \subseteq \mathcal{A}} \epsilon(\mathcal{G})$ be the *minimum service slack*. We define their heterogeneous token system counterparts in §4. Our results show that ϵ acts as a stability parameter: when $\epsilon > 0$, the token system can be stabilized under the minimum token selection rule, otherwise stability cannot be achieved under any feasible provider rule.

3 Homogeneous token systems

We start by presenting our first result that characterizes the maximal stability region for homogeneous token systems, where recall that agents may have different service request and provision rates, but every service costs one token, i.e., $c^t = 1$ for all t . Intuitively, when there is a group of agents who own a large amount of tokens, while their service request rate is independent of the provider selection rule (which helps the group to spend tokens), their token earning rate may depend on the rule, and in particular, if all available service providers own a large amount of tokens, then the social planner cannot move tokens from such *rich* agents to the rest of the agents. Thus, stability should require this unavoidable token earning rate to be smaller than the rate of spending tokens, and the following result formalizes this intuition.

Theorem 1. *The homogeneous token system $(n, P, Q, d, \mathbf{1})$ is stable under the minimum token selection rule if and only if for all nonempty group of agents $\mathcal{G} \subsetneq \mathcal{A}$, we have*

$$p(\mathcal{G}) > q(\mathcal{G})^d, \tag{5}$$

i.e., $\epsilon > 0$. Moreover, if at least one inequality in (5) fails to hold, i.e., $\epsilon \leq 0$, then under no feasible stationary provider selection rule the token system is stable.

When agents are symmetric, i.e., $p_i = q_i = 1/n$ for all $i \in \mathcal{A}$, [Ashlagi et al. \(2026\)](#) already established the stability of the token system under minimum token selection rule when $n \geq 2$ and $d \geq 2$. Their result immediately follows from Theorem 1, since for any group $\mathcal{G}$ with size $1 \leq g < n$, (5) holds:

$$p(\mathcal{G}) = \frac{g}{n} > \left(\frac{g}{n}\right)^d = q(\mathcal{G})^d.$$

Indeed, it is natural to expect that when the service request and provision rates are close to each other, agents should benefit from trading favors and are able to balance their contributions over time, and Theorem 1 provides the maximal region that characterizes how much asymmetry the token system can tolerate to ensure stability. The fact that no token system is stable when $d = 1$ immediately follows from Theorem 1 as well, since $p_i > q_i$ for all $i \in \mathcal{A}$ contradicts

$$\sum_{i \in \mathcal{A}} p_i = \sum_{i \in \mathcal{A}} q_i = 1.$$

Another interesting setting in which (5) holds is when agents are *aligned*, i.e., $p_i = q_i$ for all $i \in \mathcal{A}$ and $d \geq 2$. In the context of kidney exchange, hospitals can have huge differences in terms of their sizes, i.e., large values of $|p_i - p_j|$'s or $|q_i - q_j|$'s, whereas what matters more for the stability of token systems is the ratio between the service request and provision rates, which can have a stylized interpretation of the composition of patient-donor pairs (the ratio between easy-to-match and hard-to-match pairs).

To provide some intuition for why condition (5) is necessary and sufficient, note that the service slack $\epsilon(\mathcal{G}) = p(\mathcal{G}) - q(\mathcal{G})^d$ is the difference between the total service request rate and the unavoidable token earning rate of group $\mathcal{G}$. Whenever all available service providers belong to group $\mathcal{G}$, which happens with probability $q(\mathcal{G})^d$, the service provider belongs to $\mathcal{G}$ under every feasible stationary provider selection rule. Therefore, when $q(\mathcal{G})^d > p(\mathcal{G})$, no feasible selection rule can control this *positive drift*. The *zero drift* case, i.e., $q(\mathcal{G})^d = p(\mathcal{G})$, is more delicate, as stability does not hold due to symmetric random walk type arguments provided by Lemma 2.

Our second result complements the stability arguments by providing a tail bound on the token distribution, in terms of the minimum service slack ϵ and $\tau := 1 - \sum_{i \in \mathcal{A}} p_i^2$. As the proof of Theorem 2 shows, τ can be interpreted as the fraction of times in the steady-state that a trade between the service requester and service provider affects the token balances, i.e., there is no self-service.

Theorem 2. *Assume that the homogeneous token system $(n, P, Q, d, \mathbf{1})$ is stable and let π be the steady-state distribution of $(\mathbf{s}^t : t \geq 0)$ under the minimum token selection rule. Let $\tau = 1 - \sum_{i \in \mathcal{A}} p_i^2$. Then for all $i \in \mathcal{A}$ and $M \in \mathbb{Z}_{\geq 0}$, we have*

$$\mathbb{P}_\pi(|s_i^t| > M) \leq \min \left\{ 1, \frac{\tau}{\epsilon(M+1)} \right\}. \quad (6)$$

Moreover, we have the following bound on the expected range of the token amounts:

$$\mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\tau}{\epsilon}. \quad (7)$$

In Theorem 2, the ratio τ/ϵ that dictates both upper bounds in (6) and (7) compares the fraction of times there is no self-service to the “weakest” group of agents’ service slack. When τ is fixed, smaller the ϵ , weaker the expected range bound in the steady-state.

3.1 Analysis

In order to prove our stability results, we use the following two standard Lyapunov drift results, which are reformulated for Markov chains with countable state spaces.

Lemma 1 (Theorem 11.3.4 in [Meyn and Tweedie \(2012\)](#)). *Let X^t be an irreducible, time homogeneous Markov chain on a countable state space $\mathcal{X}$. Assume that there exists a real-valued function $V : \mathcal{X} \rightarrow [0, \infty)$, a nonempty finite set $C \subset \mathcal{X}$, and some constant $a > 0$ such that*

$$\mathbb{E}[V(X^{t+1}) - V(X^t) \mid X^t = x] \leq -a \quad \text{for all } x \notin C, \quad (8)$$

$$\mathbb{E}[V(X^{t+1}) \mid X^t = x] < \infty \quad \text{for all } x \in C. \quad (9)$$

Then X^t is positive recurrent.

Lemma 2 (Theorem 11.5.2 in [Meyn and Tweedie \(2012\)](#)). *Let X^t be an irreducible, time homogeneous Markov chain on a countable state space $\mathcal{X}$. Assume that there exists a real-valued function $V : \mathcal{X} \rightarrow [0, \infty)$, a nonempty (not necessarily finite) set $C \subset \mathcal{X}$, such that*

$$\mathbb{E}[V(X^{t+1}) - V(X^t) \mid X^t = x] \geq 0 \quad \text{for all } x \notin C, \quad (10)$$

$$\sup_{x \in \mathcal{X}} \mathbb{E}[|V(X^{t+1}) - V(X^t)| \mid X^t = x] < \infty, \quad (11)$$

$$V(y) > \sup_{x \in C} V(x) \quad \text{for some } y \notin C. \quad (12)$$

Then X^t is not positive recurrent.

Once a suitable Lyapunov function $V(\cdot)$ is established, we can deduce whether the token system is stable or not per Lemmas 1 and 2. Since we are also interested in proving tail bounds under the minimum token selection rule, it is also crucial to characterize the negative drift (8) in terms of the market primitives. The next proposition is a useful step towards providing this.

Proposition 1. *Under the minimum token selection rule, we have*

$$\mathbb{E}[s_{k^t}^t - s_{l^t}^t \mid \mathbf{s}^t = \mathbf{s}] = \sum_{m=1}^{n-1} \epsilon(U_m) \delta_m \quad \text{for all } \mathbf{s} \in \mathcal{S}.$$

In particular, if $\epsilon > 0$, then $\mathbb{E}[s_{k^t}^t - s_{l^t}^t \mid \mathbf{s}^t = \mathbf{s}] \geq \epsilon R(\mathbf{s})$.

Proof of Proposition 1. Fix $\mathbf{s} \in \mathcal{S}$, a permutation σ satisfying (2), and the corresponding upper

groups $\mathcal{U}_m$'s. For the service requester, per (3) we have

$$\mathbb{E}[s_{k_t}^t \mid \mathbf{s}^t = \mathbf{s}] = s_{\sigma(1)} + \sum_{m=1}^{n-1} \delta_m \mathbb{P}(k_t \in \mathcal{U}_m \mid \mathbf{s}^t = \mathbf{s}). \quad (13)$$

Since the realization of the service requester k_t is independent of $\mathbf{s}^t$, we have

$$\mathbb{P}(k_t \in \mathcal{U}_m \mid \mathbf{s}^t = \mathbf{s}) = \mathbb{P}(k_t \in \mathcal{U}_m) = \sum_{i \in \mathcal{U}_m} p_i = p(\mathcal{U}_m). \quad (14)$$

For the service provider, per (3) we have

$$\mathbb{E}[s_{l_t}^t \mid \mathbf{s}^t = \mathbf{s}] = s_{\sigma(1)} + \sum_{m=1}^{n-1} \delta_m \mathbb{P}(l_t \in \mathcal{U}_m \mid \mathbf{s}^t = \mathbf{s}).$$

When $\delta_m > 0$, any agent who belongs to $\mathcal{U}_m$ has strictly more tokens than every agent who belongs to $\mathcal{U}_m^c$. Therefore, under the minimum token selection rule, if there exists an available agent who does not belong to $\mathcal{U}_m$, then we must have $l_t \notin \mathcal{U}_m$. Thus, we have $\mathbb{P}(l_t \in \mathcal{U}_m \mid \mathbf{s}^t = \mathbf{s}) = q(\mathcal{U}_m)^d$ so that

$$\mathbb{E}[s_{l_t}^t \mid \mathbf{s}^t = \mathbf{s}] = s_{\sigma(1)} + \sum_{m=1}^{n-1} \delta_m q(\mathcal{U}_m)^d. \quad (15)$$

Combining (13), (14) and (15), we get

$$\mathbb{E}[s_{k_t}^t - s_{l_t}^t \mid \mathbf{s}^t = \mathbf{s}] = \sum_{m=1}^{n-1} (p(\mathcal{U}_m) - q(\mathcal{U}_m)^d) \delta_m = \sum_{m=1}^{n-1} \epsilon(\mathcal{U}_m) \delta_m. \quad (16)$$

When $\epsilon > 0$, since $R(\mathbf{s}) = \sum_{m=1}^{n-1} \delta_m$ per (4), we have $\mathbb{E}[s_{k_t}^t - s_{l_t}^t \mid \mathbf{s}^t = \mathbf{s}] \geq \epsilon R(\mathbf{s})$. $\square$

Now we use Proposition 1, Lemma 1, and Lemma 2 to prove Theorems 1 and 2.

Proof of Theorem 1. We first prove the sufficiency. Assume that $\epsilon > 0$. Consider the quadratic Lyapunov function $V(\mathbf{s}^t) := \sum_{i \in \mathcal{A}} (s_i^t)^2$. Define the event $E_t := \{k_t \neq l_t\}$. The expected one-step difference of $V(\mathbf{s}^t)$ is

$$\begin{aligned} \mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}, E_t] &= \mathbb{E}[(s_{k_t}^t - 1)^2 - (s_{k_t}^t)^2 + (s_{l_t}^t + 1)^2 - (s_{l_t}^t)^2 \mid \mathbf{s}^t = \mathbf{s}, E_t] \\ &= 2\mathbb{E}[s_{l_t}^t - s_{k_t}^t \mid \mathbf{s}^t = \mathbf{s}, E_t] + 2. \end{aligned}$$

Under the event $E_t^c = \{k_t = l_t\}$, $V(\mathbf{s}^t)$ remains the same. Therefore, per Proposition 1, we have

$$\begin{aligned}\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] &= 2\mathbb{E}[s_{l_t}^t - s_{k_t}^t \mid \mathbf{s}^t = \mathbf{s}] + 2\mathbb{P}(E_t \mid \mathbf{s}^t = \mathbf{s}) \\ &= -2 \sum_{m=1}^{n-1} \epsilon(\mathcal{U}_m) \delta_m + 2\mathbb{P}(E_t \mid \mathbf{s}^t = \mathbf{s}) \\ &\leq -2\epsilon R(\mathbf{s}) + 2.\end{aligned}\tag{17}$$

Now we use Lemma 1 to conclude stability. Consider the set $C = \{\mathbf{s} \in \mathcal{S} : R(\mathbf{s}) \leq 2/\epsilon\}$. Clearly C is nonempty since $\mathbf{0} \in C$, and C is a finite set since $\sum_{i \in \mathcal{A}} s_i = 0$. When $\mathbf{s} \notin C$, we have

$$\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] \leq -2\epsilon R(\mathbf{s}) + 2 < -2 < 0,$$

so that (8) is satisfied with $a = 2$. When $\mathbf{s} \in C$, since any agent can earn or lose at most one token every period, we have $\mathbb{E}[V(\mathbf{s}^{t+1}) \mid \mathbf{s}^t = \mathbf{s}] \leq \sum_{i \in \mathcal{A}} (|s_i| + 1)^2 < \infty$, so that (9) is satisfied. Per Lemma 1, we have stability under the minimum token selection rule.

Now we prove the necessity. Assume that $\epsilon \leq 0$. Thus, there exists a nonempty group $\mathcal{G} \subsetneq \mathcal{A}$ such that $p(\mathcal{G}) \leq q(\mathcal{G})^d$. Consider the system under an arbitrary feasible stationary provider selection rule $\psi \in \Psi$. Recall (1) that states $S^{t+1}(\mathcal{G}) - S^t(\mathcal{G}) = \mathbb{1}\{l_t \in \mathcal{G}\} - \mathbb{1}\{k_t \in \mathcal{G}\}$. Taking expectations, we have

$$\begin{aligned}\mathbb{E}[S^{t+1}(\mathcal{G}) - S^t(\mathcal{G}) \mid \mathbf{s}^t = \mathbf{s}] &= \mathbb{P}(l_t \in \mathcal{G} \mid \mathbf{s}^t = \mathbf{s}) - \mathbb{P}(k_t \in \mathcal{G} \mid \mathbf{s}^t = \mathbf{s}) \\ &\geq \mathbb{P}(\mathcal{B}^t \subset \mathcal{G} \mid \mathbf{s}^t = \mathbf{s}) - p(\mathcal{G})\end{aligned}\tag{18}$$

$$\begin{aligned}&\geq q(\mathcal{G})^d - p(\mathcal{G}) \\ &\geq 0,\end{aligned}\tag{19}$$

where (18) follows since on the event $\mathcal{B}^t \subset \mathcal{G}$, we must have $l_t \in \mathcal{G}$, and k_t is independent of $\mathbf{s}^t$. Consider the Lyapunov function $V(\mathbf{s}^t) = (S^t(\mathcal{G}))^+$ and the set $C = \{\mathbf{s} \in \mathcal{S} : \sum_{i \in \mathcal{G}} s_i \leq 0\}$.

Now we use Lemma 2 to conclude that stability fails. First, note that (10) holds since when $\mathbf{s}^t \notin C$, we have $S^t(\mathcal{G}) \geq 1$, which implies that $S^{t+1}(\mathcal{G}) \geq 0$. Therefore, we have $V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) = S^{t+1}(\mathcal{G}) - S^t(\mathcal{G})$ for all $\mathbf{s}^t \notin C$ so that per (19), we have

$$\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] = \mathbb{E}[S^{t+1}(\mathcal{G}) - S^t(\mathcal{G}) \mid \mathbf{s}^t = \mathbf{s}] \geq 0 \quad \text{for all } \mathbf{s} \notin C.$$

Condition (11) holds since the mapping $x \rightarrow (x)^+$ is 1-Lipschitz, i.e., $|x^+ - y^+| \leq |x - y|$ for all

$x, y \in \mathbb{R}$, so that for all $s \in \mathcal{S}$,

$$\begin{aligned}\mathbb{E}[|V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t)| \mid \mathbf{s}^t = \mathbf{s}] &= \mathbb{E}[|(S^{t+1}(\mathcal{G}))^+ - (S^t(\mathcal{G}))^+| \mid \mathbf{s}^t = \mathbf{s}] \\ &\leq \mathbb{E}[|S^{t+1}(\mathcal{G}) - S^t(\mathcal{G})| \mid \mathbf{s}^t = \mathbf{s}] \\ &\leq \mathbb{E}[|\mathbb{1}\{l_t \in \mathcal{G}\} - \mathbb{1}\{k_t \in \mathcal{G}\}| \mid \mathbf{s}^t = \mathbf{s}] \\ &\leq 1.\end{aligned}$$

Finally for (12), first note that $\sup_{\mathbf{s} \in \mathcal{C}} V(\mathbf{s}) = 0$, and consider $\mathbf{y} = \mathbf{e}_i - \mathbf{e}_j \in \mathcal{S}$ (agent i has 1 token and agent j has -1 token and remaining agents have 0 tokens), where $i \in \mathcal{G}$ and $j \notin \mathcal{G}$. Then $V(\mathbf{y}) = 1 > 0$, where $\mathbf{y} \notin \mathcal{C}$. Thus, per Lemma 2, we conclude that no feasible stationary provider selection rule stabilizes the token system when $\epsilon \leq 0$. $\square$

Proof of Theorem 2. Since the token system is stable, we have $\epsilon > 0$ per Theorem 1. Consider the request rate shifted quadratic Lyapunov function $V(\mathbf{s}^t) = \sum_{i \in \mathcal{A}} (s_i^t + p_i)^2$. This shift helps to cancel out the self-service term in the drift calculations, which will be evident momentarily below. The one-step difference of $V(\mathbf{s}^t)$ is

$$V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) = 2(s_{l_t}^t - s_{k_t}^t) + 2(p_{l_t} - p_{k_t}) + 2\mathbb{1}\{k_t \neq l_t\},$$

where note that the one-step difference is zero on the event $\{k_t = l_t\}$. Since the realization of the service requester k_t is independent of $\mathbf{s}^t$, we have

$$\mathbb{E}[p_{k_t} \mid \mathbf{s}^t] = \sum_{i \in \mathcal{A}} p_i^2 = 1 - \tau. \quad (20)$$

Moreover, k_t is independent of l_t so that

$$\mathbb{P}(k_t = l_t \mid \mathbf{s}^t) = \sum_{i \in \mathcal{A}} p_i \mathbb{P}(l_t = i \mid \mathbf{s}^t) = \mathbb{E}[p_{l_t} \mid \mathbf{s}^t]. \quad (21)$$

Per (20), (21), and Proposition 1, the expected one-step difference of $V(\mathbf{s}^t)$ is

$$\begin{aligned}\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t] &= -2\mathbb{E}[s_{k_t}^t - s_{l_t}^t \mid \mathbf{s}^t] - 2\mathbb{E}[p_{k_t} \mid \mathbf{s}^t] + 2\mathbb{E}[p_{l_t} \mid \mathbf{s}^t] + 2\mathbb{P}(k_t \neq l_t \mid \mathbf{s}^t) \\ &= -2\mathbb{E}[s_{k_t}^t - s_{l_t}^t \mid \mathbf{s}^t] - 2\mathbb{E}[p_{k_t} \mid \mathbf{s}^t] + 2\mathbb{E}[p_{l_t} \mid \mathbf{s}^t] + 2(1 - \mathbb{E}[p_{l_t} \mid \mathbf{s}^t]) \\ &= -2\mathbb{E}[s_{k_t}^t - s_{l_t}^t \mid \mathbf{s}^t] - 2(1 - \tau) + 2 \\ &\leq -2\epsilon R(\mathbf{s}^t) + 2\tau\end{aligned}$$

To conclude a stationary bound, we use the following standard truncation argument. Consider the truncation $\tilde{V}(\mathbf{s}^t) = \min\{V(\mathbf{s}^t), B\}$ for some constant $B > 0$. Note that for any sample path, we have

$$(V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t)) \mathbb{1}\{V(\mathbf{s}^t) < B\} \geq \tilde{V}(\mathbf{s}^{t+1}) - \tilde{V}(\mathbf{s}^t)$$

Therefore, we get

$$\begin{aligned} \mathbb{E}[\tilde{V}(\mathbf{s}^{t+1}) - \tilde{V}(\mathbf{s}^t) \mid \mathbf{s}^t] &\leq \mathbb{1}\{V(\mathbf{s}^t) < B\} \mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t] \\ &\leq \mathbb{1}\{V(\mathbf{s}^t) < B\} (2\tau - 2\epsilon R(\mathbf{s}^t)). \end{aligned} \quad (22)$$

Under stationarity, $\mathbf{s}^{t+1}$ and $\mathbf{s}^t$ have the same distribution, and since $0 \leq \tilde{V}(\mathbf{s}^t) \leq B$, we have $\mathbb{E}_\pi[\tilde{V}(\mathbf{s}^{t+1})] = \mathbb{E}_\pi[\tilde{V}(\mathbf{s}^t)] < \infty$. Thus, the truncated potential has zero expected increment: $\mathbb{E}_\pi[\tilde{V}(\mathbf{s}^{t+1}) - \tilde{V}(\mathbf{s}^t)] = 0$. Per (22), we have

$$0 \leq 2\tau \mathbb{P}_\pi(V(\mathbf{s}^t) < B) - 2\epsilon \mathbb{E}_\pi[R(\mathbf{s}^t) \mathbb{1}\{V(\mathbf{s}^t) < B\}],$$

which implies

$$\mathbb{E}_\pi[R(\mathbf{s}^t) \mathbb{1}\{V(\mathbf{s}^t) < B\}] \leq \frac{\tau}{\epsilon} \mathbb{P}_\pi(V(\mathbf{s}^t) < B) \leq \frac{\tau}{\epsilon}.$$

Using monotone convergence theorem, as $B \rightarrow \infty$, we have

$$\mathbb{E}_\pi R(\mathbf{s}^t) \leq \frac{\tau}{\epsilon}. \quad (23)$$

Since $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all t , we have $s_i^t = \frac{1}{n} \sum_{j \in \mathcal{A}} (s_i^t - s_j^t)$ for all $i \in \mathcal{A}$. Using triangle inequality, we get

$$|s_i^t| = \left| \frac{1}{n} \sum_{j \in \mathcal{A}} (s_i^t - s_j^t) \right| \leq \frac{1}{n} \sum_{j \in \mathcal{A}} |s_i^t - s_j^t| \leq \frac{n}{n} R(\mathbf{s}^t) = R(\mathbf{s}^t).$$

Therefore, per (23), we have

$$\mathbb{E}_\pi[|s_i^t|] \leq \mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\tau}{\epsilon}.$$

Finally, by Markov's inequality, for all $M \in \mathbb{Z}_{\geq 0}$, we have

$$\mathbb{P}_\pi(|s_i^t| > M) = \mathbb{P}_\pi(|s_i^t| \geq M + 1) \leq \frac{\mathbb{E}_\pi[|s_i^t|]}{M + 1} \leq \frac{\tau}{\epsilon(M + 1)}.$$

We finish the proof by verifying the interpretation of τ as the fraction of times in steady-state that there is no self-service. Write

$$\mathbb{E}_\pi[s_i^{t+1} - s_i^t] = 0 = \mathbb{P}_\pi(l_t = i) - p_i.$$

Thus, taking stationary expectations in (21) yields $\mathbb{P}_\pi(k_t = l_t) = \sum_{i \in \mathcal{A}} p_i^2 = 1 - \tau$. $\square$

Remark 1. *The proof of Theorem 2 reveals that one can improve upon the tail bound previously given by Ashlagi et al. (2026) for symmetric agents. In particular, when $p_i = q_i = 1/n$ and $d \geq 2$, then for all $i \in \mathcal{A}$ and $M \in \mathbb{Z}_{\geq 0}$, we have*

$$\mathbb{P}_\pi(|s_i^t| > M) \leq \min \left\{ 1, \frac{2(1 - 1/n)}{(M + 1)} \right\},$$

where the previous bound was $\min\{1, 5/(M + 1)\}$ (Ashlagi et al. 2026)[Theorem 1]. The proof is given in Appendix A.

4 Heterogeneous token systems

In general, services may cost different amounts of tokens. In the home exchange platform studied in Goedde (2025), where agents request available homes that they like and pay hosts a per-night price in tokens, median homes are priced around 150 tokens and the top-end ones at 300 tokens. The Center Liquidity Contribution program of National Kidney Registry deducts more tokens (points) from a hospital if their patient-donor pair has a unfavorable blood type donor and highly sensitized patient.

In this section, we investigate to what extent this added heterogeneity affects the necessary and sufficient stability conditions in (5). Intuitively, an agent who requests service infrequently can be a part of a stable homogeneous token system, but in the heterogeneous setting, she may destabilize the system if her service requests are expensive. We now analyze heterogeneous token systems, where the service cost sequence $\mathbf{c} = (c^t)_{t \geq 0}$ is identically and independently distributed and may depend on the service requester at each time period.

4.1 Preliminaries

For each agent $i \in \mathcal{A}$, we associate a probability distribution F_i on $\mathbb{Z}_{\geq 1}$ that governs the cost of the service requested by agent i . Given the service requester k_t at time t , the cost of the service is drawn independently: $c^t \mid \{k^t = i\} \sim F_i$. Hence the pairs (k_t, c^t) are i.i.d. throughout the time horizon, where for any $X \subset \mathbb{Z}_{\geq 1}$, we have $\mathbb{P}(k_t = i, c^t \in X) = p_i F_i(X)$. Note that $(c^t)_{t \geq 0}$

is i.i.d. with the distribution $F(X) = \mathbb{P}(c^t \in X) = \sum_{i \in \mathcal{A}} p_i F_i(X)$, and our model allows the service requester and her service cost to be dependent. We provide some examples in Example 1. Throughout this section, we assume that the service costs have a finite second moment, i.e., $\nu := \mathbb{E}[(c^t)^2] < \infty$.

Let $\mu_i := \mathbb{E}[c^t \mid k_t = i]$ and $\mu := \mathbb{E}[c^t] = \sum_{i \in \mathcal{A}} p_i \mu_i < \infty$. Thus, agent i 's expected service cost per request is μ_i , and the system's aggregate service cost is μ . For each agent $i \in \mathcal{A}$, define her expected *expenditure* share by

$$\bar{p}_i := \frac{\mathbb{E}[c^t \mathbb{1}\{k_t = i\}]}{\mu} = \frac{p_i \mathbb{E}[c^t \mid k_t = i]}{\mu} = \frac{p_i \mu_i}{\sum_{j \in \mathcal{A}} p_j \mu_j},$$

and for any nonempty group of agents $\mathcal{G} \subsetneq \mathcal{A}$, we let $\bar{p}(\mathcal{G}) = \sum_{i \in \mathcal{G}} \bar{p}_i$. Note that $(\bar{p}_i)_{i \in \mathcal{A}} > 0$ is a probability vector, where $\bar{p}(\mathcal{G})$ measures the expected expenditure share of group $\mathcal{G}$. It is useful to note that if k_t and c^t are independent for all t , then we have $\bar{p}_i = p_i$ for all $i \in \mathcal{A}$.

Similar to the homogeneous case in §3, we now define service slacks appropriately. Note that any group $\mathcal{G}$ spends in expectation $\sum_{i \in \mathcal{G}} p_i \mu_i = \mu \bar{p}(\mathcal{G})$ many tokens at each time period. Whenever all available providers belong to $\mathcal{G}$, which happens with probability $q(\mathcal{G})^d$, the service provider belongs to $\mathcal{G}$ under any stationary provider selection rule, so that in expectation, the unavoidable token earning rate of $\mathcal{G}$ is $\mu q(\mathcal{G})^d$. Therefore, for each nonempty group $\mathcal{G} \subsetneq \mathcal{A}$, we define their *service slack* as $\bar{\epsilon}(\mathcal{G}) := \bar{p}(\mathcal{G}) - q(\mathcal{G})^d$, and let $\bar{\epsilon} := \min_{\emptyset \neq \mathcal{G} \subsetneq \mathcal{A}} \bar{\epsilon}(\mathcal{G})$ be the *minimum service slack*.

Example 1. Consider two agents with $p_1 = p_2 = 1/2$. Assume that F_1 assigns probability 1 to cost of 1, and F_2 assigns equal probabilities to costs of 1 and 3. This gives $\mu_1 = 1$, $\mu_2 = 2$, and $\mu = 1 \cdot (1/2) + 2 \cdot (1/2) = 3/2$. Note that the service requester and the service cost are dependent, since only the second agent can have a service of cost 3. However, the sequence $(c^t)_{t \geq 0}$ is i.i.d. with the distribution that assigns 3/4 to cost of 1 and 1/4 to cost of 3. If instead we consider $F_1 = F_2$, then the service requester and her service cost become independent.

4.2 The maximal stability region and tail bounds

Our main results for the heterogeneous token systems interestingly mirror the ones established for the homogeneous case.

Theorem 3 (Maximal stability region). *Assume that service costs have a finite second moment, i.e., $\mathbb{E}[(c^t)^2] < \infty$. Then the heterogeneous token system $(n, P, Q, d, \mathbf{c})$ is stable under the minimum token selection rule if and only if for all nonempty group of agents $\mathcal{G} \subsetneq \mathcal{A}$, we have*

$$\bar{p}(\mathcal{G}) > q(\mathcal{G})^d \quad (24)$$

i.e., $\bar{\epsilon} > 0$. Moreover, if at least one inequality in (24) fails to hold, i.e., $\bar{\epsilon} \leq 0$, then under no feasible stationary provider selection rule the token system is stable.

The proof of Theorem 3 mostly follows the weighted version of the proof of Theorem 1, and thus it is delegated to Appendix B. For our tail bound results, recall $\nu := \mathbb{E}[(c^t)^2] < \infty$ and define $\nu_i := \mathbb{E}[(c^t)^2 \mathbb{1}\{k_t = i\}]$ for all $i \in \mathcal{A}$. Note that $\nu = \sum_{i \in \mathcal{A}} \nu_i$. Analogous to τ , the parameter that drives the upper bound result in the heterogeneous case is $\gamma := \nu - \sum_{i \in \mathcal{A}} \bar{p}_i \nu_i = \sum_{i \in \mathcal{A}} (1 - \bar{p}_i) \nu_i$, where γ is the expected squared token transfer size between the service requester and the service provider in the steady-state as shown in the proof of Theorem 4.

Theorem 4. Assume that the heterogeneous token system is stable and let π be the steady-state distribution of $(\mathbf{s}^t : t \geq 0)$ under the minimum token selection rule. Let $\gamma = \sum_{i \in \mathcal{A}} (1 - \bar{p}_i) \nu_i$. Then for all $i \in \mathcal{A}$ and $M \in \mathbb{Z}_{\geq 0}$, we have

$$\mathbb{P}_\pi(|s_i^t| > M) \leq \min \left\{ 1, \frac{\gamma}{\mu \bar{\epsilon} (M + 1)} \right\}.$$

Moreover, we have the following bound on the expected range of the token amounts:

$$\mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\gamma}{\mu \bar{\epsilon}}.$$

The proof of Theorem 4 can be found in Appendix B. The main change compared to the homogeneous case is that the Lyapunov function is shifted by ν_i/μ instead of p_i , which helps canceling out the state-dependent token transfer terms as in the proof of Theorem 2. We finish this section with the following corollary, that shows that when the service costs are independent of the service requester, the maximal stability region of a heterogeneous token system coincides with the homogeneous case. The proof is given in Appendix B.

Corollary 1. Assume that $F_i = F$ for all $i \in \mathcal{A}$ in the heterogeneous token system so that c^t and k_t are independent for all $t \in \mathbb{Z}_{\geq 0}$. Then the system is stable under the minimum token selection rule if and only if for all nonempty group of agents $\mathcal{G} \subsetneq \mathcal{A}$, we have

$$p(\mathcal{G}) > q(\mathcal{G})^d \quad (25)$$

i.e., $\bar{\epsilon} > 0$. If at least one inequality in (25) fails to hold, i.e., $\bar{\epsilon} \leq 0$, then under no feasible

stationary provider selection rule the token system is stable. Moreover, when (25) hold, for all $i \in \mathcal{A}$ and $M \in \mathbb{Z}_{\geq 0}$, we have

$$\mathbb{P}_\pi(|s_i^t| > M) \leq \min \left\{ 1, \frac{\nu\tau}{\mu\epsilon(M+1)} \right\}, \quad \text{and} \quad \mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\nu\tau}{\mu\epsilon}.$$

5 Token system market design

In this section, we discuss several market design levers for the social planner. In §5.1, we provide a simple stability test that helps the social planner to verify whether the token system is stable under the minimum token selection rule. As a corollary, we also characterize the minimum service availability density d needed to guarantee stability. In §5.2, we study a simple ex-ante commitment game, where agents take actions based on mean-variance preferences, and we show that stability is sufficient to ensure full participation of agents.

In what follows, let $\tilde{p}_i = \bar{p}_i$ if the token system is heterogeneous, and let $\tilde{p}_i = p_i$ if the token system is homogeneous. For any nonempty group $\mathcal{G} \subset \mathcal{A}$, we accordingly define $\tilde{p}(\mathcal{G}) := \sum_{i \in \mathcal{G}} \tilde{p}_i$, and recall that $q(\mathcal{G}) = \sum_{i \in \mathcal{G}} q_i$.

5.1 Stability test and minimal service availability density

Note that for the social planner to verify the stability conditions (5) (or, (24)), she needs to check $2^n - 2$ many inequalities. The next proposition provides an efficient stability test for the planner where $n - 1$ comparisons suffice to verify stability after appropriately sorting the agents based on market primitives.

Proposition 2. *Let $\tilde{\sigma} : \{1, 2, \dots, n\} \rightarrow \mathcal{A}$ such that*

$$\frac{\tilde{p}_{\tilde{\sigma}(1)}}{q_{\tilde{\sigma}(1)}} \leq \frac{\tilde{p}_{\tilde{\sigma}(2)}}{q_{\tilde{\sigma}(2)}} \leq \dots \leq \frac{\tilde{p}_{\tilde{\sigma}(n)}}{q_{\tilde{\sigma}(n)}},$$

where any ties are broken arbitrarily. For all $1 \leq k < n$, let $\mathcal{G}_k := \{\tilde{\sigma}(1), \dots, \tilde{\sigma}(k)\}$. Then the token system $(n, P, Q, d, \mathbf{c})$ with $d > 1$ is stable under the minimum token selection rule if and only if for all $1 \leq k < n$, we have

$$\tilde{p}(\mathcal{G}_k) > q(\mathcal{G}_k)^d. \tag{26}$$

What Proposition 2 implies is that the social planner can sort the agents by their $\tilde{p}_i/q_i$'s, which requires $O(n \log n)$ comparisons, and check $n - 1$ many inequalities to verify stability, instead of checking exponentially many inequalities.

Proof of Proposition 2. The necessity part follows immediately since each $\mathcal{G}_k$ is a nonempty

subset of $\mathcal{A}$ and $\mathcal{G}_k \neq \mathcal{A}$. We prove the sufficiency part via contraposition. Fix some $d > 1$ and assume that the stability conditions (5) (or, (24)) are violated by some group. Consider the service slack $\tilde{\epsilon}(\mathcal{G}) := \tilde{p}(\mathcal{G}) - q(\mathcal{G})^d$ of each subset of $\mathcal{A}$, including $\mathcal{A}$ and the empty set $\emptyset$. With a slight abuse of notation, let $\mathcal{G}$ be the subset minimizing $\tilde{\epsilon}(\mathcal{G})$ over all subsets of $\mathcal{A}$. If the minimum is negative, then $\mathcal{G}$ cannot be $\mathcal{A}$ or $\emptyset$, since the service slacks of $\mathcal{A}$ and $\emptyset$ are both 0. If the minimum is zero, then $\emptyset \neq \mathcal{G} \subsetneq \mathcal{A}$ violates the stability conditions (5) (or, (24)).

Let $x := q(\mathcal{G}) \in (0, 1)$. For any $i \in \mathcal{G}$, we must have $\tilde{\epsilon}(\mathcal{G} \setminus \{i\}) \geq \tilde{\epsilon}(\mathcal{G})$ since $\mathcal{G}$ is a minimizer, which implies

$$\tilde{p}_i \leq x^d - (x - q_i)^d < dx^{d-1}q_i, \quad (27)$$

where the strict inequality follows from the strict convexity of $f(y) = y^d$. Similarly, for any $j \notin \mathcal{G}$, we must have $\tilde{\epsilon}(\mathcal{G} \cup \{j\}) \geq \tilde{\epsilon}(\mathcal{G})$, which implies

$$\tilde{p}_j \geq (x + q_j)^d - x^d > dx^{d-1}q_j, \quad (28)$$

Combining (27) and (28) gives

$$\frac{\tilde{p}_i}{q_i} < \frac{\tilde{p}_j}{q_j} \quad \text{for all } i \in \mathcal{G}, j \notin \mathcal{G}.$$

Therefore, $\mathcal{G} = \mathcal{G}_k$ for some $1 \leq k < n$ and $\mathcal{G}$ violates (26). $\square$

As a corollary of Proposition 2, the social planner can determine the minimum service availability density required to ensure stability. Per (26), since $\tilde{p}(\mathcal{G}_k), q(\mathcal{G}_k) \in (0, 1)$, taking logarithms yields $d > \log(\tilde{p}(\mathcal{G}_k)) / \log(q(\mathcal{G}_k))$, which gives the following result.

Corollary 2. *The token system is stable under the minimum token selection rule if and only if*

$$d \geq 1 + \left\lfloor \max_{1 \leq k < n} \frac{\log(\tilde{p}(\mathcal{G}_k))}{\log(q(\mathcal{G}_k))} \right\rfloor.$$

5.2 A commitment game

Following the mean-variance preferences framework studied in Ashlagi et al. (2026)[Section 5], we now generalize their commitment game in the case of homogeneous token systems.⁷ Fix a time horizon $T \in \mathbb{Z}_{\geq 1}$. We assume that before the system starts, each agent can either commit to participate in the token system for T periods, or they can opt out of the token system. For

⁷This is for ease of exposition. Our arguments continue to hold for heterogeneous token systems with minor modifications.

simplicity, we assume that opting out gives a payoff of zero, and in general, our results allow agents' outside options to be constants that are independent of the time horizon T . Suppose that n agents commit to participate. If agent $i \in \mathcal{A}$ receives service, she gains a payoff of $r_i > 0$, and if she provides service, she incurs a cost $c_i \in (0, r_i)$. Let

$$A_i(T) = \sum_{t=0}^{T-1} \mathbb{1}\{k_t = i\} \quad \text{and} \quad B_i(T) = \sum_{t=0}^{T-1} \mathbb{1}\{l_t = i\},$$

denote the total number of times agent i is a service requester and a service provider, respectively. Hence, the total payoff to agent i is $U_i(T) := r_i A_i(T) - c_i B_i(T)$.

We now follow a simple model of mean-variance preferences (e.g., see [Markowitz \(1952\)](#), [Tobin \(1958\)](#), [Pomatto et al. \(2020\)](#)). Assume that agent i chooses to participate in the token system if $V_i(T) := \mathbb{E}[U_i(T)] - \kappa_i \text{Var}(U_i(T)) > 0$, where $\kappa_i > 0$ represents their risk sensitivity. We say that full participation is strictly individually rational if $V_i(T) > 0$ for all $i \in \mathcal{A}$. Thus, full participation is a strict Nash equilibrium if every agent prefers commitment over opting out.

Theorem 5. *Assume that the homogeneous token system satisfies $\epsilon > 0$. Then under the minimum token selection rule, for all $i \in \mathcal{A}$, we have*

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[U_i(T)]}{T} = p_i(r_i - c_i), \tag{29}$$

$$\lim_{T \rightarrow \infty} \frac{\text{Var}(U_i(T))}{T} = p_i(1 - p_i)(r_i - c_i)^2, \tag{30}$$

$$\lim_{T \rightarrow \infty} \frac{V_i(T)}{T} = p_i(r_i - c_i)[1 - \kappa_i(r_i - c_i)(1 - p_i)]. \tag{31}$$

Therefore, if

$$\kappa_i < \frac{1}{(r_i - c_i)(1 - p_i)} \quad \text{for all } i \in \mathcal{A}, \tag{32}$$

then for all sufficiently large T , full participation is a strict Nash equilibrium.

Theorem 5 relates our stability definition to ex-ante participation in a commitment game. As noted in ([Ashlagi et al. 2026](#), Section 5), stability helps control agents' payoff risk. In our model, stability keeps the token balances uniformly bounded in expectation in second moment. When each agent's risk sensitivity parameter satisfies (32), her long-run expected net benefit exceeds the variance term in her payoff. Therefore, when the time horizon T is large enough, full participation is strictly individually rational and yields a strict Nash equilibrium of the commitment game.

Proof of Theorem 5. We will use the fact that

$$\sup_{t \geq 0} \mathbb{E} \left[\sum_{i \in \mathcal{A}} (s_i^t)^2 \right] \leq \Gamma, \quad (33)$$

for some constant $\Gamma < \infty$, which we prove in Appendix C. Since each service costs one token, we have $B_i(T) - A_i(T) = s_i^T$. Thus, we can write agent i 's total payoff as

$$U_i(T) = r_i A_i(T) - c_i B_i(T) = (r_i - c_i) A_i(T) - c_i s_i^T.$$

Since $A_i(T) \sim \text{Binomial}(T, p_i)$, $\mathbb{E}[A_i(T)] = p_i T$, and per (33), we have

$$|\mathbb{E}[U_i(T)] - (r_i - c_i) p_i T| = |c_i \mathbb{E}[s_i^T]| \leq c_i \sqrt{\Gamma},$$

where dividing both sides with T and taking limit proves (29). Since $\text{Var}(A_i(T)) = p_i(1 - p_i)T$, we also have

$$\begin{aligned} \text{Var}(U_i(T)) &= (r_i - c_i)^2 T p_i(1 - p_i) + c_i^2 \text{Var}(s_i^T) \\ &\quad - 2c_i(r_i - c_i) \text{Cov}(A_i(T), s_i^T). \end{aligned}$$

Using Cauchy-Schwarz, we get $|\text{Cov}(A_i(T), s_i^T)| \leq \sqrt{\text{Var}(A_i(T)) \text{Var}(s_i^T)}$. Per (33), we have $\sqrt{\text{Var}(s_i^T)} \leq \sqrt{\Gamma}$ so that

$$|\text{Var}(U_i(T)) - (r_i - c_i)^2 T p_i(1 - p_i)| \leq c_i^2 \Gamma + 2c_i(r_i - c_i) \sqrt{T p_i(1 - p_i) \Gamma},$$

where dividing both sides by T and taking the limit proves (30). Since $V_i(T) = \mathbb{E}[U_i(T)] - \kappa_i \text{Var}(U_i(T))$, we have

$$\lim_{T \rightarrow \infty} \frac{V_i(T)}{T} = p_i(r_i - c_i) [1 - \kappa_i(r_i - c_i)(1 - p_i)],$$

which proves (31). The limit (31) is strictly positive when (32) is satisfied. Therefore, for sufficiently large T , every agent strictly prefers commitment to opting out, given that all other agents are committing. □

6 Concluding remarks

We study the stability of token systems, where agents are asymmetric in their service request and service provision rates, there is limited service availability, and service costs are possibly heterogeneous. We characterize the maximal stability region and show that the minimum token selection rule achieves it, where among the agents who are available to provide the requested service, the agent with the minimum amount of tokens is chosen to provide service.

Our analysis reveals a stability parameter ϵ that measures the minimum difference between the total service request rate and unavoidable token earning rate over all possible groups of agents. We show that when ϵ is strictly greater than zero, the token system can be stabilized, otherwise no stationary provider selection rule can stabilize the system. Our results quantify how much asymmetry and heterogeneity a token system can tolerate. We also provide explicit stationary tail bounds that measure how concentrated the token distribution is in the long-run.

Many directions remain open. In many applications, service availability is determined by a compatibility network. It would be interesting to study the question of how does the network structure affect the maximal stability region, to guide practitioners to effectively operate token systems in applications such as kidney exchange. In this paper, we considered service costs that are possibly random and dependent on service requester. It would be also interesting to study more general settings, where service costs can jointly depend on the service requester and provider.

References

- Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman. Market failure in kidney exchange. *American Economic Review*, 109(11):4026–4070, 2019.
- Itai Ashlagi and Süleyman Kerimov. Scrip systems with minimal availability. In *WINE*, page 335, 2019.
- Itai Ashlagi and Alvin E Roth. Kidney exchange: An operations perspective. *Management Science*, 67(9):5455–5478, 2021.
- Itai Ashlagi, Süleyman Kerimov, Omer Tamuz, and Geng Zhao. The power of two in token systems. *Management Science*, 2026.
- Yang Bo, Milind Dawande, and Ganesh Janakiraman. Analysis of scrip systems: On an open question in johnson et al.(2014). *Operations Research*, 2018.
- Chen Chen, Yilun Chen, and Pengyu Qian. Incentivizing resource pooling. *Management Science*, 72(8):6454–6471, 2026.
- Julius Goedde. Pricing in markets without money: Theory and evidence from home exchanges. 2025.

- Xing Hu, Zhixi Wan, and Qichuan Yin. Token design for favor trading with dynamic membership. *Available at SSRN 6175079*, 2025.
- Kris Johnson, David Simchi-Levi, and Peng Sun. Analyzing scrip systems. *Operations Research*, 62(3): 524–534, 2014.
- Ian A Kash, Eric J Friedman, and Joseph Y Halpern. Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion. *Distributed Computing*, 25(5):335–357, 2012.
- Ian A Kash, Eric J Friedman, and Joseph Y Halpern. An equilibrium analysis of scrip systems. *ACM Transactions on Economics and Computation*, 3(3):13, 2015.
- Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.
- Sean P Meyn and Richard L Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.
- Markus Möbius. Trading favors. *Mimeo, Microsoft Research*, 2001.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. Stochastic dominance under independent noise. *Journal of Political Economy*, 128(5):1877–1900, 2020.
- Canice Prendergast. The allocation of food to food banks. *Journal of Political Economy*, 130(8):1993–2017, 2022.
- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007.
- James Tobin. Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2): 65–86, 1958.
- Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gun Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems*, volume 35, 2003.

A Proofs from Section 3

Here we present the proof for Remark 1 that modifies the proof of Theorem 2 accordingly.

Proof of Remark 1. Assume that agents are symmetric, i.e., $p_i = q_i = 1/n$ for every $i \in A$ and $d \geq 2$. We have $\tau = 1 - \sum_{i \in A} p_i^2 = 1 - n(1/n)^2 = 1 - 1/n$. Following the truncation and monotone convergence arguments in the proof of Theorem 2, per (23), we have

$$\mathbb{E}_\pi \left[\sum_{m=1}^{n-1} \epsilon(U_m) \delta_m \right] \leq \tau = 1 - \frac{1}{n}. \quad (34)$$

Since $d \geq 2$ and $|\mathcal{U}_m| = n - m$, we have

$$\epsilon(U_m) = \frac{n-m}{n} - \left(\frac{n-m}{n}\right)^d \geq \frac{m(n-m)}{n^2}. \quad (35)$$

Since $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all t , we have $s_i^t = \frac{1}{n} \sum_{j \in \mathcal{A}} (s_i^t - s_j^t)$ for all $i \in \mathcal{A}$. Therefore, using triangle inequality gives

$$\begin{aligned} \frac{1}{n} \sum_{i \in \mathcal{A}} |s_i^t| &\leq \frac{1}{n^2} \sum_{i,j \in \mathcal{A}} |s_i^t - s_j^t| = \frac{2}{n^2} \sum_{m=1}^{n-1} m(n-m) \delta_m \\ &\leq 2 \sum_{m=1}^{n-1} \epsilon(U_m) \delta_m, \end{aligned}$$

where the equality holds since δ_m appears in exactly $m(n-m)$ unordered pairwise differences, and the last inequality holds per (35). Taking stationary expectations, using (34) and the fact that agents are symmetric yields for all $i \in \mathcal{A}$,

$$\mathbb{E}_\pi[|s_i^t|] = \frac{1}{n} \mathbb{E}_\pi \left[\sum_{j \in \mathcal{A}} |s_j^t| \right] \leq 2 \left(1 - \frac{1}{n} \right).$$

Finally, by Markov's inequality, for all $M \in \mathbb{Z}_{\geq 0}$, we have

$$\mathbb{P}_\pi(|s_i^t| > M) = \mathbb{P}_\pi(|s_i^t| \geq M+1) \leq \frac{\mathbb{E}_\pi[|s_i^t|]}{M+1} \leq \frac{2(1-1/n)}{(M+1)}.$$

□

B Proofs from Section 4

Proof of Theorem 3. First, assume that $\bar{\epsilon} > 0$. Independence gives

$$\mathbb{E}[c^t \mathbb{1}\{k_t \in \mathcal{G}\} \mid \mathbf{s}^t = \mathbf{s}] = \mu \bar{p}(\mathcal{G}), \quad \mathbb{E}[c^t \mathbb{1}\{\mathcal{B}^t \subset \mathcal{G}\} \mid \mathbf{s}^t = \mathbf{s}] = \mu q(G)^d.$$

Per (16) in the proof of Proposition 1, we have

$$\mathbb{E}[c^t (s_{k_t}^t - s_{l_t}^t) \mid \mathbf{s}^t = \mathbf{s}] = \mu \sum_{m=1}^{n-1} \bar{\epsilon}(\mathcal{U}_m) \delta_m \geq \mu \bar{\epsilon} R(\mathbf{s}).$$

Consider again the quadratic Lyapunov function $V(\mathbf{s}^t) = \sum_i (s_i^t)^2$. The expected one-step difference of $V(\mathbf{s}^t)$ is

$$\begin{aligned}\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] &= 2\mathbb{E}[c^t(s_{l^t}^t - s_{k^t}^t) \mid \mathbf{s}^t = \mathbf{s}] + 2\mathbb{E}[(c^t)^2 \mathbb{1}\{k^t \neq l^t\} \mid \mathbf{s}^t = \mathbf{s}] \\ &\leq -2\mu\bar{\epsilon}R(\mathbf{s}) + 2\nu,\end{aligned}\tag{36}$$

where the inequality follows as in (17). Consider the set

$$C = \left\{ \mathbf{s} \in S : R(\mathbf{s}) \leq \frac{\nu + 1}{\mu\bar{\epsilon}} \right\},$$

which is nonempty since $\mathbf{0} \in C$, and C is finite since $\sum_{i \in \mathcal{A}} s_i = 0$. When $s \notin C$, the drift (36) is strictly less than -2 so that (8) of Lemma 1 holds. Moreover, we have $\mathbb{E}[V(\mathbf{s}^{t+1}) \mid \mathbf{s}^t = \mathbf{s}] \leq 2V(\mathbf{s}) + 4\nu < \infty$ so that (9) of Lemma 1 also holds. Thus, we have stability under the minimum token selection rule.

To prove necessity, we assume that $\bar{\epsilon}(\mathcal{G}) \leq 0$ for some nonempty proper group $\mathcal{G} \subsetneq \mathcal{A}$, and fix any feasible stationary provider rule. Following the necessity part of the proof of Theorem 1, we have

$$\mathbb{E}[S^{t+1}(\mathcal{G}) - S^t(\mathcal{G}) \mid \mathbf{s}^t = \mathbf{s}] \geq \mu(q(\mathcal{G}))^d - \mu\bar{p}(\mathcal{G}) = -\mu\bar{\epsilon}(\mathcal{G}) \geq 0.$$

Let $V(\mathbf{s}^t) = (S^t(\mathcal{G}))^+$ and the set $C = \{\mathbf{s} \in \mathcal{S} : \sum_{i \in \mathcal{G}} s_i \leq 0\}$. We have

$$V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \geq \Delta S^t(\mathcal{G}) \quad \text{for all } \mathbf{s}^t \notin C,$$

which validates (10). The condition (11) holds again since the mapping $x \rightarrow (x)^+$ is 1-Lipschitz so that for all $\mathbf{s} \in \mathcal{S}$, we have

$$\mathbb{E}[|V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t)| \mid \mathbf{s}^t = \mathbf{s}] \leq \mathbb{E}[|S^{t+1}(\mathcal{G}) - S^t(\mathcal{G})| \mid \mathbf{s}^t = \mathbf{s}] \leq \mu.$$

Finally, $\sup_{\mathbf{s} \in C} V(\mathbf{s}) = 0$, and $V(\mathbf{e}_i - \mathbf{e}_j) = 1$ for any $i \in \mathcal{G}$ and $j \notin \mathcal{G}$, which verifies (12). Hence, per Lemma 2, no feasible stationary provider selection rule stabilizes the token system when $\bar{\epsilon} \leq 0$. $\square$

Proof of Theorem 4. Per Theorem 3, the heterogeneous token system is stable so that $\bar{\epsilon} > 0$.

Consider the shifted quadratic Lyapunov function

$$V(\mathbf{s}^t) = \sum_{i \in A} \left(s_i^t + \frac{\nu_i}{\mu} \right)^2,$$

where its one-step difference is

$$V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) = 2c^t(s_{l_t}^t - s_{k_t}^t) + \frac{2c^t}{\mu}(\nu_{l_t} - \nu_{k_t}) + 2(c^t)^2 \mathbb{1}\{k_t \neq l_t\}.$$

Since (k_t, c^t) is independent of l_t , we have

$$\frac{1}{\mu} \mathbb{E}[c^t(\nu_{l_t} - \nu_{k_t}) \mid \mathbf{s}^t] = \mathbb{E}[\nu_{l_t} \mid \mathbf{s}^t] - \sum_i \bar{p}_i \nu_i,$$

and

$$\mathbb{E}[(c^t)^2 \mathbb{1}\{k_t \neq l_t\} \mid \mathbf{s}^t] = \nu - \mathbb{E}[\nu_{l_t} \mid \mathbf{s}^t].$$

Since $\gamma = \nu - \sum_i \bar{p}_i \nu_i$, per Proposition 1, we have

$$\begin{aligned} \mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] &= 2\gamma - 2\mu \sum_{m=1}^{n-1} \bar{\epsilon}(U_m) \delta_m \\ &\leq 2\gamma - 2\mu \bar{\epsilon} R(\mathbf{s}). \end{aligned}$$

We now follow the truncation argument from the proof of Theorem 2 by setting $\tilde{V}(\mathbf{s}^t) = \min\{V(\mathbf{s}^t), B\}$ for some constant $B > 0$. Stationarity implies

$$\mu \bar{\epsilon} \mathbb{E}_\pi [R(\mathbf{s}^t) \mathbb{1}\{V(\mathbf{s}^t) < B\}] \leq \gamma \mathbb{P}_\pi(V(\mathbf{s}^t) < B) \leq \gamma.$$

Using monotone convergence theorem, as $B \rightarrow \infty$, we have

$$\mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\gamma}{\mu \bar{\epsilon}}.$$

Using $|s_i^t| \leq R(\mathbf{s}^t)$ and Markov's inequality yield

$$\mathbb{P}_\pi(|s_i^t| > M) \leq \frac{\gamma}{\mu \bar{\epsilon}(M + 1)}.$$

The interpretation of γ proceeds as follows. Since $\mathbb{E}_\pi[s_i^{t+1} - s_i^t] = 0 = \mu \mathbb{P}_\pi(l_t = i) - \mu \bar{p}_i$, we

have $\mathbb{P}_\pi(l_t = i) = \bar{p}_i$ so that

$$\mathbb{E}_\pi[(c^t)^2 \mathbb{1}\{k^t \neq l_t\}] = \nu - \sum_i \nu_i \mathbb{P}_\pi(l_t = i) = \gamma.$$

□

Proof of Corollary 1. Now that c^t and k_t are independent, we have $\mu_i = \mu, \bar{p}_i = p_i, \nu_i = p_i \nu$. Therefore for any nonempty group $\mathcal{G} \subsetneq \mathcal{A}$, we have $\bar{\epsilon}(\mathcal{G}) = \epsilon(\mathcal{G})$, and $\bar{\epsilon} = \epsilon$. Thus,

$$\gamma = \nu - \sum_i p_i(p_i \nu) = \nu \left(1 - \sum_i p_i^2\right) = \nu \tau.$$

Substituting these identities into the bounds given in Theorem 4 gives

$$\mathbb{E}_\pi[R(\mathbf{s}^t)] \leq \frac{\nu \tau}{\mu \epsilon}, \quad \mathbb{P}_\pi(|s_i^t| > M) \leq \min \left\{ 1, \frac{\nu \tau}{\mu \epsilon (M + 1)} \right\},$$

as claimed. □

C Proofs from Section 5

In this section, we prove that

$$\sup_{t \geq 0} \mathbb{E} \left[\sum_{i \in \mathcal{A}} (s_i^t)^2 \right] \leq \Gamma.$$

for some constant $\Gamma < \infty$, to complete the proof of Theorem 5. Let $V(\mathbf{s}) = \sum_{i=1}^n s_i^2$ and $W(\mathbf{s}) = V(\mathbf{s})^2$. Per (17) from the proof of Theorem 1, we have

$$\mathbb{E}[V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] \leq -2\epsilon R(\mathbf{s}) + 2.$$

Since the service requester pays one token to the service provider, in the worst case this exchange happens between the *richest* and the *poorest* agents so that

$$|V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t)| \leq 2R(\mathbf{s}^t) + 2.$$

Moreover, $(R(\mathbf{s}^t))^2 = (\max_{i \in \mathcal{A}} s_i^t - \min_{i \in \mathcal{A}} s_i^t)^2 \leq 2(\max_{i \in \mathcal{A}} s_i^t)^2 + 2(\min_{i \in \mathcal{A}} s_i^t)^2 \leq 2V(\mathbf{s})$. Therefore, we have

$$(V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t))^2 \leq (2R(\mathbf{s}^t) + 2)^2 \leq 4R(\mathbf{s}^t)^2 + 8R(\mathbf{s}^t) + 4$$

$$\begin{aligned}
&\leq 12R(\mathbf{s}^t)^2 + 4 \\
&\leq 24V(\mathbf{s}^t) + 4.
\end{aligned}$$

Rewriting $V(\mathbf{s}^{t+1}) = V(\mathbf{s}^t) + (V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t))$, we get

$$\begin{aligned}
W(\mathbf{s}^{t+1}) - W(\mathbf{s}^t) &= (V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t))(V(\mathbf{s}^{t+1}) + V(\mathbf{s}^t)) \\
&= 2V(\mathbf{s}^t)(V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t)) + (V(\mathbf{s}^{t+1}) - V(\mathbf{s}^t))^2.
\end{aligned}$$

Thus, we obtain

$$\begin{aligned}
\mathbb{E}[W(\mathbf{s}^{t+1}) - W(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] &\leq 2V(\mathbf{s})(-2\epsilon R(\mathbf{s}) + 2) + 24V(\mathbf{s}) + 4 \\
&= 28V(\mathbf{s}) + 4 - 4\epsilon R(\mathbf{s})V(\mathbf{s}) \\
&= V(\mathbf{s})(29 - 4\epsilon R(\mathbf{s})) + 4 - V(\mathbf{s}). \tag{37}
\end{aligned}$$

When $R(\mathbf{s}) \geq 29/(2\epsilon)$, (37) is at most $4 - V(\mathbf{s})$. When $R(\mathbf{s}) < 29/(2\epsilon)$, $V(\mathbf{s})$ is bounded by some constant since $V(\mathbf{s}) \leq nR(\mathbf{s})^2$. Therefore, there exists a constant $\Gamma < \infty$ such that for all $\mathbf{s} \in \mathcal{S}$, we have

$$\mathbb{E}[W(\mathbf{s}^{t+1}) - W(\mathbf{s}^t) \mid \mathbf{s}^t = \mathbf{s}] \leq \Gamma - V(\mathbf{s}).$$

Let π be the stationary distribution of $(\mathbf{s}^t : t \geq 0)$ provided per Theorem 1. Applying the truncation and monotone convergence argument from the proof of Theorem 2 to $\widetilde{W}(\mathbf{s}^t) = \min\{W(\mathbf{s}^t), B\}$ for some constant $B > 0$ gives $\mathbb{E}_\pi[V(\mathbf{s}^t)] \leq \Gamma$, which implies for some redefined finite constant Γ , we have

$$\sup_{t \geq 0} \mathbb{E} \left[\sum_{i \in \mathcal{A}} (s_i^t)^2 \right] \leq \Gamma < \infty.$$