

Efficient online exchange via fiat money

Mihaela van der Schaar, Jie Xu & William Zame

Economic Theory

ISSN 0938-2259

Econ Theory

DOI 10.1007/s00199-013-0744-4



Your article is protected by copyright and all rights are held exclusively by Springer-Verlag Berlin Heidelberg. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Efficient online exchange via fiat money

Mihaela van der Schaar · Jie Xu · William Zame

Received: 5 August 2011 / Accepted: 10 January 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract In many online systems, individuals provide services for each other; the recipient of the service obtains a benefit but the provider of the service incurs a cost. If benefit exceeds cost, provision of the service increases social welfare and should therefore be encouraged—but the individuals providing the service gain no (immediate) benefit from providing the service and hence have an incentive to withhold service. Hence, there is scope for designing a protocol that improves welfare by encouraging exchange. To operate successfully within the confines of the online environment, such a protocol should be distributed, robust, and consistent with individual incentives. This paper proposes and analyzes protocols that rely solely on the exchange of *fiat money* or *tokens*. The analysis has much in common with work on search models of money but the requirements of the environment also lead to many differences from previous analyses—and some surprises; in particular, existence of equilibrium becomes a thorny problem and the optimal quantity of money is different.

Keywords Online exchange · Token exchange

JEL Classification D51 · E40

1 Introduction

This paper is motivated by the problem of online exchange of files (or data or services). In typical systems that serve this purpose—Napster, now defunct, is the most familiar example but there are many in current operation, including Gnutella

M. van der Schaar · J. Xu
Department of Electrical Engineering, University of California, Los Angeles, CA 90095, USA

W. Zame (✉)
Department of Economics, University of California, Los Angeles, CA 90095, USA
e-mail: zame@econ.ucla.edu

and Kazaa (file sharing), Seti@home (computational assistance), Slashdot and Yahoo Answers (answers to queries)—a single interaction involves an agent who wants a file (or data or service) and an agent who can provide it. The former benefits from obtaining the file but the latter bears the (often non-trivial) cost of providing it and so has an incentive to free-ride.¹ Assuming that benefit exceeds the cost, provision of the service increases social welfare and should therefore be encouraged—but how?

This problem is a particular instance of trade in the absence of a double coincidence of wants, which has motivated a large literature on search models of money. Indeed, we shall formalize our problem in the same terms, and the “solution” we develop is for a (benevolent) designer to institute a system that relies on *fiat money* or *tokens* (we use the terms interchangeably), to introduce a quantity of tokens into the system, and to recommend strategies to the participants for requesting and providing service. Because our agents are self-interested, the designer must recommend strategies that constitute an equilibrium—but our environment also imposes other constraints on the designer: the system must be anonymous and distributed, must take account of the fact that agents meet only electronically (and not face-to-face), that files and tokens are indivisible, that the designer cannot know the precise parameters of the population and, perhaps most importantly, that the designer cannot constrain the number of tokens that agents hold.^{2,3}

This paper asks how much a designer can accomplish, given these constraints, by judicious choice of the *protocol*—the quantity of tokens and the recommended strategies. To answer this question, we characterize equilibrium protocols and among these, the ones that are robust to small perturbations of the population parameters (the designer’s slight misperceptions of these parameters); we prove that robust equilibrium protocols exist; we provide bounds for the efficiency of robust equilibrium protocols; we show that the “optimal quantity of money” in our setting is different than in other settings considered in the literature; we provide an *effective* procedure for choosing a robust equilibrium protocol whose efficiency is at least good (if not optimal); and we provide numerical simulations to illustrate some of the theorems and also to demonstrate that *design matters*: a great deal of efficiency may be lost if the designer chooses the “wrong” protocol.

As in the familiar search models of money, our environment is populated by a continuum of agents each of whom is initially endowed with a unique file that can

¹ Empirical studies show that this free-riding problem can be quite severe: in the Gnutella system for instance, almost 70 % of users share no files at all (Adar and Huberman 2000).

² It might be useful to note that none of the systems mentioned above involve a central authority or central monitoring agency. Napster, for instance, merely maintained many partial lists (distributed across many servers) of music files available and contact information for subscribers who had these files; users seeking files could simply search these lists and then contact the file-holder directly. In practice, the absence of a central agency is crucial, since it could not handle the volume of traffic that would be generated and would be exceedingly vulnerable to attack. Hence, a distributed system is a *sine qua non*.

³ The reader might wonder how agents who do not meet in person can exchange tokens at all, since they can only exchange electronic files, and electronic files would seem to be easily duplicated. In fact, however, there are practicable, secure, and private procedures for online token exchange, utilizing hardware or software or both; see Buttyan and Hubaux (2001), Chandrakumar et al. (2003) and Cuffoletti (2010) for instance. Similar procedures can also serve as escrow accounts to assure that service that is promised is actually provided and that payment that is promised is actually made.

be duplicated and provided to others.⁴ In each period, a fraction of the population is matched; one member of each match—the *client*—must decide whether to request service (provision of a file or forwarding of a packet) and the other—the *server*—must decide whether to provide the service (if requested). The client who receives the service derives a benefit, the server who provides the service incurs a cost. To simplify the analysis, we assume here that, except for the uniqueness of the files they possess, all agents are identical and that all files are equally valuable to receive and equally costly to provide. (We discuss extensions in the Conclusion.) We assume benefit exceeds cost, so that social welfare is increased when the service is provided, but that cost is strictly positive, so that the server has a disincentive to provide it. The designer supplies a supply of tokens and recommends strategies (circumstances under which service should be requested or provided); together these constitute a *protocol*. We assume that the price of service is fixed at one token; this restriction seems natural in our environment and is made in much of the search literature; see also below and the Conclusion. We differ from much of the literature in three ways suggested by the motivation discussed. First, we do not impose an exogenous upper bound on money holdings: agents can store as much money as they wish. Second, we require that the protocol should induce an equilibrium (i.e., that the recommended strategies are best replies in the (unique) steady-state distribution) that is *robust* to small perturbations of the population parameters. Third, we allow the designer to control *both* the money supply *and* the price. As we shall show, each of these has significant implications.

Leaving aside degenerate protocols in which there is no trade, *all* robust equilibrium protocols are Markov (not history dependent), *symmetric* (the population plays a single pure strategy) and have a particularly simple form: clients request service whenever their token holding is above zero; servers provide service when their token holding is at or below a *threshold* K and do not provide service when the token holding is above K . We prove that robust equilibria exist but the absence of an exogenous upper holdings makes the proof surprisingly hard. (See Sect. 6). Having shown that robust equilibria exist we turn to our original question: which equilibrium protocols are the most efficient? We have shown that we can restrict attention to threshold strategies; among protocols that employ the threshold K the one that would be most efficient *if agents were compliant* has token supply $K/2$. However, these protocols need not be equilibria and the most efficient protocols may have token supplies different from $K/2$; $K/2$ *need not be the optimal quantity of money*. We go on to provide estimates for efficiency of various protocols and an effective procedure for the designer to choose a “good”—if not optimal—protocol. Simulations illustrate these results and some related points.

Following a discussion of the literature below, the remainder of the paper is organized in the following way. Section 3 introduces the model and defines strategies, the steady-state value function, best responses, equilibrium, and robust equilibrium. Section 4 discusses the nature of equilibrium and robust equilibrium and establishes existence of robust equilibrium. Simulations illustrate the nature of equilibrium. Section 5 discusses efficiency of protocols, shows that almost full efficiency is obtain-

⁴ In the real systems we have in mind, the population is in the tens of thousands or hundreds of thousands so a continuum model seems a reasonable approximation.

able by an equilibrium protocol if agents are sufficiently patient or the benefit/cost ratio is sufficiently large, establishes lower bounds for efficiency of equilibrium protocols and uses these bounds to provide an effective procedure for constructing efficient equilibrium protocols, and shows that the optimal quantity of money in our setting, with no exogenous bound on money holdings is different from the optimal quantity of money in settings with an exogenous bound on money holdings. Numerical simulations give some idea of the efficiency loss resulting from choosing the “wrong” protocol. Section 6 concludes and offers some directions for further research. Proofs are collected in the “Appendix”.

2 Literature

Following the seminal work of [Kiyotaki and Wright \(1989\)](#), there is a large literature on search models of money which has contributed enormously to our understanding of money in various environments. A portion of this literature—e.g. [Camera and Corbae \(1999\)](#), [Berentsen \(2002\)](#)—allows agents to accumulate more than one unit of money, while maintaining the assumption of [Kiyotaki and Wright \(1989\)](#) that there is an exogenously given upper bound on money holdings; this is precluded in our environment. A different portion of this literature—e.g. [Cavalcanti and Wallace \(1999a,b\)](#), [Berentsen et al. \(2007\)](#) and [Hu et al. \(2002\)](#)—assumes that agents have and can condition on (complete or partial) knowledge of the money holdings of (some of) their counter-parties in each match, which is again precluded in our environment. A particularly striking paper in this literature is [Kocherlakota \(2002\)](#), which shows that any individually rational outcome can be supported in equilibrium provided money is infinitely divisible and the common discount factor is above some minimum threshold. However, [Kocherlakota \(2002\)](#) also assumes that agents have a good deal of information about the money holdings of their counterparties. To quote the abstract: “The one-money theorem says that the allocation is achievable using only one money if that money is divisible and money holdings are observable. The two-money theorem says that the allocation is achievable using two divisible monies, even if money holdings are concealable.” To elaborate: the one-money theorem assumes that agents must display their *true* money holdings; the two-money theorem assumes that agents can display *less* money than they actually have but cannot display *more* money than they actually have. In both cases, agents have (complete or partial) knowledge about the money holdings of their counter-parties and can condition on it. In our work, agents have *no* knowledge about the money holdings of their counter-parties and so *cannot* condition on it.

Our work is closest to [Zhou \(1999\)](#) and [Berentsen \(2000\)](#). [Zhou \(1999\)](#) assumes money is divisible and the supply of money is given endogenously but the price is determined endogenously.⁵ [Berentsen \(2000\)](#) assumes money is indivisible and the price is given exogenously but the money supply is determined endogenously. In our work, both the money supply *and* the price are chosen exogenously by the designer. Of course, from an economic point of view, all that really matters is the ratio M/p

⁵ We say “the” price because [Zhou \(1999\)](#) considers only single-price equilibria.

of the money supply M to the price p , fixing either the money supply or the price amounts simply to choosing a normalization. So it would be more accurate to say that in both Zhou (1999) and Berentsen (2000), the ratio M/p is determined *endogenously*: fixing M , as Zhou (1999) does, is just choosing a normalization; fixing p , as Berentsen (2000) does, is just choosing a different normalization. In our work, we have chosen a particular price normalization $p = 1$ but the money supply, and hence the ratio M/p , is determined *exogenously* by the designer. Our designer has more control and that control is important because if the designer did not control M/p , the designer could not be sure of designing an optimal equilibrium protocol. Indeed, if the designer did not control M/p , it would seem to make no sense to even talk about designing protocols, much less optimal equilibrium protocols. We also note that neither Zhou (1999) nor Berentsen (2000) prove that equilibrium exist; they both provide sufficient conditions but those conditions are stringent and endogenous—they are not conditions on the primitives of the model (benefit/cost ratio and discount factor).

This work also connects to an Electrical Engineering and Computer Science literature that discusses token exchanges in online communities. Some of that literature assumes that agents are compliant, rather than self-interested, and does not treat incentives and equilibrium (Chandrakumar et al. 2003; Buttyan and Hubaux 2003); some of that literature makes use of very different models than the one offered here Jarvis and Tan (2006) and Figueiredo et al. (2004); and some of the literature is not formal and rigorous, offering simulations rather than theorems (Mohr and Pai 2006). The papers closest to ours are probably Friedman et al. (2006, 2007), which treat somewhat different models. However, these papers seem puzzling in many dimensions and many of the proofs seem mysterious (at least to us).

Another literature to which this work connects is the game-theoretic literature on anonymous interactions. In a context in which interactions were publicly observable, full cooperation (i.e., provision of service) could be achieved at equilibrium by the use of trigger strategies, which deny service in the future to any agent who refuses service in the present. As Kandori (1992) and Ellison (2000) have pointed out, in some contexts, cooperation can be supported even without public observability if agents deny service in the future to *all* agents whenever they have observed an agent who refuses service in the present; in this equilibrium, any failure to provide service results in a contagion, producing wider and wider ripples of defection, until no agent provides service. However, contagion is not likely to sustain cooperation in the systems of interest to us, because the population is so large (typically comprising tens of thousands or even hundreds of thousands of agents) that an agent is unlikely, in a reasonable time frame, to meet any other agent whose network of past associations overlap with his. (When the population is literally a continuum, no agent *ever* meets any other agent whose network of past associations overlap with his.) A more relevant literature, of which Kandori (1992) is again the seminal work, uses reputation and social norms as devices as a means of incentivizing cooperation. The work that is closest to ours is Park et al. (2010), which asks which reputation-based systems can be supported in equilibrium and which of these achieve the greatest social efficiency. Because provision of service in their model depends on the reputations of *both* client and server, some central authority must keep track of and verify reputations; hence, these systems are not distributed in the sense we use here.

3 Model

The population consists of a continuum (mass 1) of infinitely lived agents. Each agent can provide a resource (e.g., a data file, audio file, video file, service) that is of benefit to others but is costly to produce (uploading a file uses bandwidth and time). The benefit of receiving this resource is b and the cost of producing it is c ; we assume $b > c > 0$.⁶ Agents care about current and future benefits/costs and discount future benefits/costs at the constant rate $\beta \in (0, 1)$. Agents are risk neutral so seek to maximize the discounted present value of a stream of benefits and costs.

Time is discrete. In each time period, a fraction $\rho \leq 1/2$ of the population is randomly chosen to be a *client* and matched with a randomly chosen *server*; the fraction $1 - 2\rho$ are unmatched.⁷ (No agent is both a client and a server in the same period.) When a client and server are matched, the client chooses whether or not to request service, the server chooses whether or not to provide service (e.g., transfer the file) if requested.

The parameters b, c, β, ρ completely describe the environment. Because the units of benefit b and cost c are arbitrary (and tokens have no intrinsic value), only the benefit/cost ratio $r = b/c$ is actually relevant. We consider variations in the benefit/cost ratio r and the discount factor β , but view the matching rate ρ as immutable.

3.1 Tokens and strategies

In a single interaction between a server and a client, the server has no incentive to provide services to the client. The mechanism we study for creating incentives to provide service involves the exchange of *tokens*. Tokens are indivisible, have no intrinsic value, cannot be counterfeited, and can be stored and transferred without loss. Each agent can hold an arbitrary non-negative finite number of tokens, but cannot hold a negative number of tokens and cannot borrow. We emphasize that our tokens are purely electronic objects and are transferred electronically.

The designer creates incentives for the agents to provide or share resources by providing a supply of *tokens* and recommending *strategies* (behavior) for agents when they are clients and servers. At the moment, we allow for strategies that depend on histories but we show that optimal strategies (best responses) depend only on current token holdings.

An *event* describes the particulars of a match at a particular time: whether the agent was chosen to be a client or a server or neither, whether the agent was matched with someone who was willing to serve or to buy, whether the agent received a benefit and surrendered a token or provided service and acquired a token or neither, and the change in the token holding. Write ϵ_t for an event at time t . A *history of length T* specifies an initial token holding m and a finite sequence of events $h = (m; \epsilon_0, \epsilon_1, \dots, \epsilon_{T-1})$. Write

⁶ If $b \leq c$ there is no social value to providing service; if $c \leq 0$ agents will always be willing to provide service.

⁷ We assume that the matching procedure is such that the Law of Large Numbers holds exactly; Duffie and Sun (2007), Alós-Ferrer (1999), Podczeck (2010) and Podczeck and Puzello (2012) construct such matching procedures.

H_T for the set of histories of length T , $H = \bigcup_T H_T$ for the set of finite histories. An *infinite history* specifies an initial token holding m and an infinite sequence of events $h = (m; \epsilon_0, \epsilon_1, \dots)$. We insist that finite/infinite histories be feasible in the sense that net token holdings are never negative (i.e., a request for service by an agent holding 0 tokens will not be honored). Given a finite or infinite history h , write $d(h, t)$ for the *change* in token holding at time t and $d^+(h, t), d^-(h, t)$ for the positive and negative parts of $d(h, t)$. Note that $d(h, t) = +1$ if the agent serves, $d(h, t) = -1$ if the agent buys, $d(h, t) = 0$ otherwise. Note also that the token holding at the end of the finite history h is

$$N(h) = m + \sum_{t=0}^{T-1} d(h, t)$$

A *strategy* is a pair $(\sigma, \tau) : H \rightarrow \{0, 1\}$; τ is the *client strategy* and σ is the *server strategy*. Following the history h , $\tau(h) = 1$ means the client requests service and $\tau(h) = 0$ means the client does not request service; $\sigma(h) = 1$ means the server provides service, $\sigma(h) = 0$ means the server does not provide service. (Note that we require individual agents to follow pure strategies, but we will eventually allow for the possibility that different agents follow different pure strategies, so the population strategy might be mixed). If service is requested and provided, a single token is transferred from client to server, so the client's holding of tokens decreases by 1 and the server's holding of tokens increases by 1. Tacitly, we assume that a token is transferred if and only if service is provided; like the transfer of tokens itself, this can be accomplished electronically in a completely distributed way.

3.2 Steady-state payoffs, values and optimal strategies

Because we consider a continuum population, assume that agents are matched randomly and can observe only their own histories, the relevant state of the system from the point of view of a single agent can be completely summarized by the fraction μ of agents who do not request service when they are clients and the fraction ν of agents who do not provide service when they are servers. If the population is in a steady state then μ, ν do not change over time. Given μ, ν , a strategy (τ, σ) determines in the obvious way a probability distribution $P(\tau, \sigma | \mu, \nu)$ over infinite histories H . We define the *discounted expected utility* to an agent whose initial token holding is m and who follows the strategy (τ, σ) to be

$$Eu(m, \tau, \sigma | \mu, \nu) = \sum_{h \in H} P(\tau, \sigma | \mu, \nu)(h) \sum_{t=0}^{\infty} \beta^t [d^+(h, t)b - d^-(h, t)c]$$

(Here and below, when some of the variables $\beta, b, c, \mu, \nu, \tau, \sigma$ are clearly understood we frequently omit all or some of them; this should not cause confusion.)

Given μ, ν, τ, σ and an initial token holding m , we define the *value* to be

$$V(m, \mu, v, \tau, \sigma) = \sup_{(\tau, \sigma)} Eu(m, \tau, \sigma | \mu, v)$$

Discounting implies that the supremum—which is taken over *all* strategy profiles—exists and is at most $b/(1 - \beta)$.

Given μ, v the strategy (τ, σ) is *optimal* or a *best response* for an initial token holding of m if

$$Eu(m, \tau, \sigma | \mu, v) \geq Eu(m, \tau', \sigma' | \mu, v)$$

for all alternative strategies τ', σ' . Because agents discount the future at the constant rate β , the strategy (τ, σ) is optimal if and only if it has the *one-shot deviation property*; that is, there does not exist a finite history h and a profitable deviation (τ', σ') that differs from (τ, σ) following the history h and nowhere else. A familiar and straightforward diagonalization argument establishes that optimal strategies exist and achieve the value; we record this fact below, omitting the proof.

Proposition 1 *For each μ, v and each initial token holding m there is an optimal strategy τ, σ and*

$$Eu(m, \tau, \sigma | \mu, v) = V(m, \mu, v, \tau, \sigma)$$

3.3 Optimal strategies

We want to characterize optimal strategies, but before we do, there is a degeneracy that must be addressed. If $\mu = 1$, then no one ever requests service so the choice of whether to provide service is irrelevant; if $v = 1$, then no one ever provides service so the choice of whether to request service is irrelevant. In what follows, we sometimes ignore or avoid these degenerate cases, but this should not lead to any confusion.

Fix β, b, c, μ, v ; let (τ, σ) be optimal for the initial token holding m . Note that the continuation of (τ, σ) must also be optimal following every history that begins with m . If h is such a history and the token holding at h is n , then (τ, σ) induces a strategy (τ^h, σ^h) from an initial token holding n that simply transposes what follows h back to time 0, and this strategy must be optimal for the initial token holding of n . Conversely, any strategy that is optimal for the initial token holding of n must also be optimal following h . It follows that optimal strategies (τ, σ) (whose existence is guaranteed by Proposition 1) *depend only on the current token holding* but are otherwise independent of history; we frequently say such strategies are *Markov*—but note that they are Markov in *individual* token holdings. Write $\Sigma(\mu, v, \beta)$ for the set of optimal strategies.

Theorem 1 *For all b, c, β, μ, v with $v < 1$, every optimal strategy (τ, σ) has the property that $\tau(n) = 1$ for every $n \geq 1$; i.e. “always request service when possible”.⁸*

⁸ Because a request for service will not be honored when an agent holds 0 tokens, it is irrelevant whether $\tau(0) = 0$ or $\tau(0) = 1$.

In view of Theorem 1, we suppress client strategies τ entirely, assuming that clients always request service whenever possible. We abuse notation and continue to write $\Sigma(\mu, \nu, \beta)$ for the set of optimal strategies.

We now show that optimal (server) strategies also have a simple form. Say that the (server) strategy σ is a *threshold strategy (with threshold K)* if

$$\begin{aligned} \sigma(n) &= 1 & \text{if } n \leq K \\ \sigma(n) &= 0 & \text{if } n > K \end{aligned} \tag{1}$$

We write σ_K for the threshold strategy with threshold K and

$$\Sigma = \{\sigma_K : 0 \leq K < \infty\}$$

for the set of threshold strategies.

Theorem 2 *For each μ, ν, b, c, β with $\mu < 1$ the set of optimal (server) strategies consists of either a single threshold strategy or two threshold strategies with adjacent thresholds.*

(The assumptions in Theorems 1 and 2 that $\nu < 1$ and $\mu < 1$ avoid the degeneracies previously noted.)

3.4 Protocols

The designer chooses a *per capita* supply of tokens $\alpha \in (0, \infty)$ and recommends a strategy to each agent; we allow for the possibility that the designer recommends different strategies to different agents. Because self-interested agents will always play a best response, the designer will recommend only strategies in Σ ; in view of anonymity, it does not matter which agents are recommended to play each strategy, but rather only the fraction of agents recommended to play each strategy. Hence, we can identify a recommendation with a *mixed threshold strategy*, which is a probability distribution on Σ ; with the obvious abuse of notation, we view γ as a function $\gamma : \mathbb{N}_+ \rightarrow [0, 1]$ such that

$$\begin{aligned} \gamma(K) &\geq 0 \text{ for each } K \geq 0 \\ \sum_{K=0}^{\infty} \gamma(K) &= 1 \end{aligned}$$

Write $\Delta(\Sigma)$ for the set of mixed threshold strategies. As usual, we identify the threshold strategy σ_K with the mixed strategy that puts mass 1 on σ_K . Assuming that the designer only recommends best responses (because other recommendations would not be followed), we interpret an element $\gamma \in \Delta(\Sigma)$ as a recommendation that the fraction $\gamma(K)$ plays the threshold strategy σ_K .

A *protocol* is a pair $\Pi = (\alpha, \gamma)$ consisting of a per capita supply of tokens $\alpha \in (0, \infty)$ and a mixed strategy recommendation $\gamma \in \Delta(\Sigma)$.

3.5 Invariant distributions

If the designer chooses the protocol $\Pi = (\alpha, \gamma)$ and agents follow the recommendation γ , we can easily describe the evolution of the *token distribution* (the distribution of token holdings). The token distribution must satisfy the two feasibility conditions:

$$\sum_{k=0}^{\infty} \eta(k) = 1 \tag{2}$$

$$\sum_{k=0}^{\infty} k\eta(k) = \alpha \tag{3}$$

Write

$$\mu = \eta(0), \quad \nu = \sum_{\sigma(k)=0} \eta(k)$$

Evidently, with respect to this token distribution, μ is the fraction of agents who have no tokens, hence cannot pay for service, and ν is the fraction of agents who do not serve (assuming they follow the protocol).

To determine the token distribution next period, it is convenient to think backward and ask how an agent could come to have k tokens in the next period. There are three possibilities; the agent could have

- $k - 1$ tokens in the current period, be chosen as a server, meet a client who can pay for service, and provide service (hence acquire a token);
- $k + 1$ tokens in the current period, be chosen as a client, meet a server who provides service, and buy service (hence expend a token);
- k tokens in the current period but neither provide service nor buy service (hence neither acquire nor expend a token).

Given a recommendation γ it is convenient to define $\sigma^\gamma : \mathbb{N}_+ \rightarrow [0, 1]$ by

$$\sigma^\gamma(n) = \sum_{K=0}^{\infty} \gamma(K)\sigma_K(n)$$

Assuming that the Law of Large Numbers holds exactly in our continuum framework and that all agents follow the recommendation γ , $\sigma^\gamma(n)$ is the fraction of agents in the population who serve when they have n tokens, so σ^γ is the population strategy. Keeping in mind that token holdings cannot be negative, it is easy to see that the token distribution next period will be

$$\begin{aligned} \eta_+(k) = & \eta(k - 1)[\rho(1 - \mu)\sigma^\gamma(k - 1)] \\ & + \eta(k + 1)[\rho(1 - \nu)] \\ & + \eta(k)[1 - \rho(1 - \mu)\sigma^\gamma(k) - \rho(1 - \nu)] \end{aligned} \tag{4}$$

where we use the convention $\eta(-1) = 0$.

Given the protocol $\Pi = (\alpha, \gamma)$, the (feasible) token distribution η is *invariant* if $\eta_+ = \eta$; that is, η is stationary when agents comply with the recommendation γ . Invariant distributions always exist and are unique.

Theorem 3 *For each protocol $\Pi = (\alpha, \gamma)$ there is a unique invariant distribution η^Π , which is completely determined by the feasibility conditions (2) and (3) and the recursion relationship*

$$\begin{aligned} \eta^\Pi(k) = & \eta^\Pi(k - 1)[\rho(1 - \mu)\sigma^\gamma(k - 1)] \\ & + \eta^\Pi(k + 1)[\rho(1 - \nu)] \\ & + \eta^\Pi(k)[1 - \rho(1 - \mu)\sigma^\gamma(k) - \rho(1 - \nu)] \end{aligned} \tag{5}$$

3.6 Definition of equilibrium and robust equilibrium

Assuming agents are rational and self-interested, they will comply with a given protocol if and only if compliance is individually optimal, that is, no agent can benefit by deviating from the protocol. To formalize this, fix a protocol $\Pi = (\alpha, \gamma)$, and let η^Π be the unique invariant distribution. Write

$$\mu^\Pi = \eta^\Pi(0), \quad \nu^\Pi = \sum_{\sigma(k)=0} \eta^\Pi(k)$$

for the fraction of agents who have no tokens and the fraction of agents who do not serve (in the invariant distribution induced by Π), respectively. We say $\Pi = (\alpha, \gamma)$ is an *equilibrium protocol* if σ_K is an optimal strategy (given given μ^Π, η^Π) whenever $\gamma(K) > 0$. That is, γ puts positive weight only on threshold strategies that are optimal, given the invariant distribution that Π itself induces.

Using the one step deviation principle, we can provide a useful alternative description of equilibrium in terms of the value function V . As noted before, because optimal strategies exist and are Markov, we may unambiguously write V_k for the value following *any* history at which the agent has k tokens. (The value function depends on the population data μ, ν and on the environmental parameters b, c, β ; but there should be no confusion in suppressing those here.)

Fix any Markov strategy σ . In order for σ to be optimal, it is necessary and sufficient that it achieves the value V_ℓ following *every* token holding ℓ . Expressed in terms of *current* token holdings and *future* values, and taking into account how behavior in a given period affects the token holding in the next period, this means that σ is optimal if and only if it satisfies the following system of equations:

$$\begin{aligned} V_0 = & \rho\sigma(0)[(1 - \mu)(-c + \beta V_1) + \mu\beta V_0] \\ & + \rho[1 - \sigma(0)]\beta V_0 + (1 - 2\rho)\beta V_0 \\ V_k = & \rho[(1 - \nu)(b + \beta V_{k-1}) + \nu\beta V_k] \\ & + \rho\sigma(k)[(1 - \mu)(-c + \beta V_{k+1}) + \mu\beta V_k] \\ & + \rho[1 - \sigma(k)]\beta V_k + (1 - 2\rho)\beta V_k \\ & \text{for each } k > 0 \end{aligned} \tag{6}$$

Applying this observation to the threshold strategy σ_K and carrying out the requisite algebra, we conclude that σ_K is optimal if and only if

$$-c + \beta V_{k+1} \geq \beta V_k \text{ if } k \leq K \tag{7}$$

$$-c + \beta V_{k+1} \leq \beta V_k \text{ if } k > K \tag{8}$$

(If it seems strange that α, γ do not appear in these inequalities, remember that the value depends on the invariant distribution η^Π , which in turn depends on α and on γ .)

Given a benefit/cost ratio $r > 1$ and a discount factor $\beta < 1$, write $EQ(r, \beta)$ for the set of protocols Π that constitute an equilibrium when the benefit/cost ratio is r and the discount factor is β . Conversely, given a protocol Π write $E(\Pi)$ for the set $\{(r, \beta)\}$ of pairs of benefit/cost ratios r and discount factors β such that Π is an equilibrium protocol when the benefit/cost ratio is r and discount factor is β . Note that EQ, E are correspondences (which might have empty values) and are inverse to each other.

Given r, β we say that Π is a *robust equilibrium* if (r, β) belongs to the interior of $E(\Pi)$; i.e., there is some $\varepsilon > 0$ such that $\Pi \in EQ(r', \beta')$ whenever $|r' - r| < \varepsilon$ and $|\beta' - \beta| < \varepsilon$. Write $EQR(r, \beta)$ for the set of robust equilibrium protocols for the benefit/cost ratio r and discount factor β and $ER(\Pi)$ for the set $\{(r, \beta)\}$ of pairs of benefit/cost ratios for which Π is a robust equilibrium. Note that EQR, ER are correspondences (which might have empty values) and are inverse to each other.

4 Equilibrium and robust equilibrium

We first describe the nature of equilibrium and robust equilibrium and then use that description to show that robust equilibria exist. The crucial fact about equilibrium is that the strategy part of an equilibrium protocol can involve mixing over at most two thresholds and that these thresholds must be adjacent; the crucial fact about robust equilibrium is that the strategy cannot involve strict mixing at all but must rather be a pure strategy.

Theorem 4 *For each benefit/cost ratio $r > 1$ and discount factor $\beta < 1$ the set $EQ(r, \beta)$ is either empty or consists of protocols that involve only (possibly degenerate) mixtures of two threshold strategies with adjacent thresholds.*

Theorem 5 *If $\Pi = (\alpha, \sigma)$ is a robust equilibrium then σ is a pure threshold strategy.*

The existence of equilibrium or robust equilibrium does not seem at all obvious (and our proof is not simple). For both intuition and technical convenience, it is convenient to work “backward”: rather than beginning with population parameters r, β and looking for protocols Π that constitute an equilibrium for those parameters, we begin with a protocol Π and look for population parameters r, β for which Π constitutes an equilibrium. That is, we do not study the correspondences $EQ(r, \beta)$ and $EQR(r, \beta)$ directly, but rather the inverse correspondences $E(\Pi)$ and $ER(\Pi)$. This is easier for several reasons, one of which is that the latter correspondences *always* have non-empty values.

To give an intuitive understanding of the difficulty and how we overcome it, fix a protocol $\Pi = (\alpha, \sigma)$ and let η^Π be the invariant distribution. Because we will eventually want to find a robust equilibrium, we assume σ is a threshold strategy: $\sigma = \sigma_K$. To look for population parameters r, β for which Π is an equilibrium, let us fix r and let β vary. (We could fix β and let r vary, or vary both β, r simultaneously, but the intuition is most easily conveyed by fixing r and letting β vary.) As we have already noted, the invariant distribution η^Π , and hence μ^Π, ν^Π , depends only on Π and so does not change as β varies. Given the invariant distribution, if β is close to 0, an agent has little incentive to acquire tokens; however, the incentive to acquire tokens increases as $\beta \rightarrow 1$. It can be shown that there is a smallest discount factor $\beta_L(\Pi)$ with the property that an agent whose discount factor is at least $\beta_L(\Pi)$ will be willing to continue providing service until he has acquired K tokens. This is not enough, because σ_K will only be incentive compatible if the agent is also willing to *stop* providing service *after* he has acquired K tokens. However, it can also be shown that there is a largest discount factor $\beta_H(\Pi)$ for which the agent is willing to stop providing service after he has acquired K tokens, and that $\beta_L(\Pi) < \beta_H(\Pi)$. (Recall that r, Π are fixed.) For every discount factor β in the closed interval $[\beta_L(\Pi), \beta_H(\Pi)]$, the protocol Π is an equilibrium when the population parameters are r, β ; that is, $(r, \beta) \in E(\Pi)$. From this, it can be shown that for every discount factor β in the interval $(\beta_L(\Pi), \beta_H(\Pi))$, the protocol Π is a robust equilibrium when the population parameters are r, β , that is, $(r, \beta) \in ER(\Pi)$. Similarly, we can hold β fixed and let r vary from 1 to ∞ , construct the corresponding intervals $[r_L(\Pi), r_H(\Pi)]$ with $r_L(\Pi) < r_H(\Pi)$ and then show that for every benefit/cost ratio r in the open interval $(r_L(\Pi), r_H(\Pi))$ the protocol Π is a robust equilibrium when the population parameters are r, β ; that is, $(r, \beta) \in ER(\Pi)$. This is the content of Theorem 6 below.

Applying this procedure for every protocol yields a family $\{ER(\Pi)\}$ of non-empty open sets of parameters r, β for which robust equilibria exist. However, our work is not done because we do not know whether a robust equilibrium exists for *given* population parameters r, β . To see that it does, we show that $\{ER(\Pi)\}$ covers a big enough set of population parameters. In particular, for each $r > 1$, there is a $\beta^* < 1$ such that $\{ER(\Pi)\}$ covers the set $\{\Pi(r, \beta) : \beta > \beta^*\}$; this means that for each $r > 1$ and $\beta > \beta^*$ there is a protocol Π that constitutes a robust equilibrium for the population parameters r, β . Similarly, for each $\beta > 0$ there is a $r^* > 0$ such that if $r > r^*$ there is a protocol that constitutes a robust equilibrium for the population parameters β, r . The proof is not easy; to do so, we first establish (Theorem 6) some special properties of protocols of the form $\Pi_K = (K/2, \sigma_K)$; we then apply these special properties (Theorem 7) to obtain the desired result.

It is natural to ask why our proof seems (and is) so much more complicated than existence proofs in the literature, such as in Berentsen (2002). The answer is that the literature establishes the existence of equilibrium only under the assumption that there is an *exogenous upper bound* K^* on the number of tokens any agent can hold. As discussed above, this assumption makes it relatively easy to show that equilibrium exists: Fix the benefit/cost ratio $r > 1$ and an arbitrary $\alpha > 0$ and consider the protocol (α, σ_{K^*}) . As above, an agent whose discount factor β is at least $\beta_L(\alpha, \sigma_{K^*})$ will provide service until he has acquired K^* tokens; under the assumption that K^*




Fig. 1 Pure and Mixed Equilibrium: $\alpha = 1/4$ (blue (thick)—pure equilibrium; red (thin)—mixed equilibrium) (color figure online)

is an upper bound on the number of tokens any agent can hold, the agent will *stop* providing service *after* he has acquired K^* tokens because, by assumption, he *cannot hold more than K^* tokens* so providing service incurs a present cost with no future benefit. Hence, (α, σ_{K^*}) is an equilibrium protocol for *every* $\beta \geq \beta_L(\Pi)$ and is a robust equilibrium protocol for every $\beta > \beta_L(\alpha, \sigma_{K^*})$. Thus, *any* protocol can be supported in equilibrium so long as agents are sufficiently patient. As we have noted in the Introduction, assuming an exogenous upper bound on token holdings does not seem realistic in the environments we consider.

Theorem 6 Fix a protocol $\Pi = (\alpha, \sigma_K)$.

- (i) For each benefit/cost ratio $r > 1$, the set $\{\beta : \Pi \in EQ(r, \beta)\}$ is a non-degenerate closed interval $[\beta_L(\Pi), \beta_H(\Pi)]$ whose endpoints are continuous functions of r .
- (ii) For each discount factor $\beta < 1$, the set $\{r : \Pi \in EQ(r, \beta)\}$ is a non-degenerate closed interval $[r_L(\Pi), r_H(\Pi)]$ whose endpoints are continuous functions of β .

These results are illustrated for $\alpha = 1/4$ in Figs. 1 and 2. (Fig. 1 may give the impression that the intervals for successive values of K do not overlap, but as Fig. 2 illustrates, they actually *do* overlap; the overlap is masked by the granularity of the Figure. However, as we have already said, we do not assert that overlapping of intervals for successive values of K is a general property).

For the special protocols $\Pi_K = (K/2, \sigma_K)$, in which the supply of tokens is exactly half the selling threshold, we prove in Theorem 7 below that the intervals corresponding to successive values of the threshold overlap but are not nested. This is exactly what we need to guarantee that (*non-degenerate*) equilibria *always exist* provided that β, r are sufficiently large. Theorem 10 provides estimates on how big β, r must be.)

Theorem 7 Robust equilibria exist whenever β, r are sufficiently large. More precisely:




Fig. 2 Threshold Equilibrium: $\alpha = 1/4$

(i) For each fixed threshold K and benefit/cost ratio $r > 1$, successive β -intervals overlap but are not nested:

$$\beta_L(\Pi_{K-1}) < \beta_L(\Pi_K) < \beta_H(\Pi_{K-1}) < \beta_H(\Pi_K)$$

Moreover,

$$\lim_{K \rightarrow \infty} \beta_L(\Pi_K) = 1$$

In particular, there is some $\beta^* < 1$ such that $EQR(r, \beta) \neq \emptyset$ for all $\beta > \beta^*$.

(ii) For each fixed threshold K and discount factor $\beta < 1$, successive r -intervals overlap but are not nested:

$$r_L(\Pi_{K-1}) < r_L(\Pi_K) < r_H(\Pi_{K-1}) < r_H(\Pi_K)$$

Moreover,

$$\lim_{K \rightarrow \infty} r_L(\Pi_K) = \infty$$

In particular, there is some $r^* > 1$ such that $EQR(r, \beta) \neq \emptyset$ for all $r > r^*$.

It follows from Theorem 7 that, as $K \rightarrow \infty$, the left-hand end points $\beta_L(\Pi_K) \rightarrow 1$, so *a fortiori* the lengths of β -intervals shrink to 0. It is natural to guess that the lengths of these intervals shrink *monotonically* to 0, and simulations suggest that this guess is correct, but we have neither a proof nor a good intuition that this is actually true. We also guess that the lengths of r -intervals shrink *monotonically*, but again we have neither a proof nor a good intuition that this is actually true.

5 Efficiency

If agents were compliant (rather than self-interested), the designer could simply instruct them to provide service at every meeting and they would comply, so the per capita social gain in each period would be $\rho(b - c)$. If agents follow the protocol $\Pi = (\alpha, \sigma_K)$, then service will be provided only in those meetings where the client can buy service and the server is willing to provide service, so the per capita social gain in each period will be $\rho(b - c)(1 - \mu^\Pi)(1 - v^\Pi)$. Hence, we define the *efficiency* of the protocol Π to be

$$\text{Eff}(\Pi) = (1 - \mu^\Pi)(1 - v^\Pi)$$

In general, it seems hard to determine the efficiency of a given protocol or to compare the efficiency of different protocols. However, we can provide efficiency bounds for protocols that utilize a given threshold strategy σ_K and compute the precise efficiency of the protocols Π_K .⁹

Theorem 8 *For each $\alpha \in (0, \infty)$, each threshold K and all values of the population parameters we have:*

- (i) $\text{Eff}(\alpha, \sigma_K) \leq 1 - \frac{1}{2\lfloor \alpha \rfloor + 1}$
- (ii) $\text{Eff}(\alpha, \sigma_K) \leq \text{Eff}(\Pi_K)$
- (iii) $\text{Eff}(\Pi_K) = \left(1 - \frac{1}{K+1}\right)^2 = \left(\frac{K}{K+1}\right)^2$

Two implications of Theorem 8 are immediate. The first is that, in order that a (threshold) protocol achieve efficiency near 1, it is necessary that it provide a large number of tokens *and also* that it prescribe a high selling threshold. Put differently: to yield full efficiency in the limit it is *not* enough to increase the number of tokens without bound or to increase the threshold without bound—*both* must be increased without bound. The second is that the protocols Π_K that provide $K/2$ tokens per capita are the most efficient protocols that utilize a given threshold strategy σ_K .

We caution the reader, however, that the protocols Π_K *need not be* equilibrium protocols, and it is (robust) equilibrium protocols that we seek. However, it follows immediately from Theorem 7 that whenever agents are sufficiently patient or the benefit/cost ratio is sufficiently large (or both), then *some* protocol Π_K is an equilibrium for large K , and hence that nearly efficient equilibrium protocols always exist.

Theorem 9 (i) *for each fixed discount factor $\beta < 1$*

$$\liminf_{r \rightarrow \infty} \sup \{ \text{Eff}(\Pi_K) : \Pi_K \in EQR(\beta, r) \} = 1$$

(ii) *for each fixed benefit/cost ratio $r > 1$*

$$\liminf_{\beta \rightarrow 1} \sup \{ \text{Eff}(\Pi_K) : \Pi_K \in EQR(\beta, r) \} = 1$$

⁹ Berentsen (2002) derives similar results in a different model, with Poisson arrival rates.

In words, as agents become arbitrarily patient or the benefit/cost ratio becomes arbitrarily large, it is possible to choose *robust equilibrium protocols* that achieve efficiency arbitrarily close to first best. Some intuition might be useful. Consider the protocols Π_K and the corresponding invariant distributions. As K increases, the fraction of agents who cannot purchase service and the fraction of agents who will do not provide both decrease—so efficiency increases. However, if r, β are fixed and K increases, then the protocols Π_K will eventually cease to be equilibrium protocols so equilibrium efficiency is bounded. On the other hand, if we fix r and let $\beta \rightarrow 1$ or fix β and let $r \rightarrow \infty$, then the thresholds K for which the protocols Π_K are equilibrium protocols blow up, and hence, efficiency tends to 1. Put differently: *high discount factors or high benefit/cost ratios make the use of high thresholds consistent with equilibrium.*

Theorem 9 provides asymptotic efficiency results; the following result presents an explicit lower bound (in terms of the population parameters r, β) for the efficiency obtainable by a robust equilibrium protocol.

Theorem 10 *Given the benefit/cost ratio $r > 1$ and the discount factor $\beta < 1$, define¹⁰*

$$K^L = \max \left\{ \log \frac{\rho\beta}{2(1-\beta)+2\rho\beta} \left(\frac{1}{1+r} \right) - 1, 0 \right\}$$

$$K^H = \log \frac{\rho\beta}{1-\beta+\rho\beta} \left(\frac{1}{2r} \right)$$

Then:

- (i) *all the thresholds K for which Π_K is a robust equilibrium protocol lie in the interval $[K^L, K^H]$;*
- (ii) *the efficiency of the optimal robust equilibrium protocol is at least $\left(1 - \frac{1}{K^L+1}\right)^2 = \left(\frac{K^L}{K^L+1}\right)^2$.*

Theorem 10 yields a lower bound on efficiency because the optimal robust equilibrium protocol is at least as efficient as any protocol Π_K that is a robust equilibrium, but does not yield an upper bound on efficiency because the optimal robust equilibrium protocol might be more efficient than any protocol Π_K that is a robust equilibrium.

Theorem 10 also yields the designer an effective procedure for finding a robust equilibrium whose efficiency is *good*, if not optimal, since all that is necessary is to check protocols Π_K with thresholds K in the (finite) interval $[K^L, K^H]$. Moreover, it is not necessary to conduct an exhaustive search. Rather the designer can begin by checking the protocol Π_K , where K is the midpoint of the interval $[K^L, K^H]$. If $M_{\sigma_K}(K - 1) \geq c/\beta$ and $M_{\sigma_K}(K) \leq c/\beta$, then Π_K is an equilibrium protocol and the search can stop. If $M_{\sigma_K}(K - 1) < c/\beta$, then for all $K' > K$, $M_{\sigma_{K'}}(K' - 1) < c/\beta$ (because $\beta^L(\Pi_{K'}) > \beta^L(\Pi_K)$). Therefore, threshold protocols for which $K' > K$

¹⁰ Note that both the basis of the logarithms and the arguments are less than 1.

cannot be an equilibrium and the designer can restrict search to the left half interval $[K^L, K]$. If $M_{\sigma_K}(K) > c/\beta$, then for all $K' < K$, $M_{\sigma_{K'}}(K') > c/\beta$ (because $\beta^H(\Pi_{K'}) > \beta^H(\Pi_K)$). Therefore, threshold protocols for which $K' < K$ cannot be an equilibrium and the designer can restrict search to the right half interval $[K, K^H]$. Continuing to bisect in this way, the designer can find an equilibrium threshold protocol in at most $\log_2(K^H - K^L)$ iterations.

5.1 The optimal quantity of money

The question naturally arises: “Which equilibrium protocols are most efficient?” Because all robust equilibrium protocols are threshold protocols, this amounts to asking for which values of α , K is (α, σ_K) the most efficient equilibrium protocol. If we focus on α we are asking a familiar question: “What is the optimal quantity of money?” Kiyotaki and Wright (1989) constrain agents to hold no more than 1 unit of money and show that the optimal quantity of money is $1/2$. Berentsen (2002) relaxes the constraint on money holdings to K and shows that (with certain assumptions) the optimal quantity of money is $K/2$. However, this conclusion is an artifact of the constraint that agents can hold no more than K units of money. In our framework, which does not place an exogenous constraint on money holdings, $K/2$ may not be—and often will not be—the optimal quantity of money. Figure 3 illustrates this point in a simulation, but it is in fact quite a robust phenomenon.

To see what this is so, fix $r \geq 1$ and $K \geq 1$. Theorem 7 guarantees that there is an open interval of discount factors for which Π_K is an equilibrium and an open interval of discount factors for which Π_{K+1} is an equilibrium and these intervals overlap:

$$\beta_L(\Pi_K) < \beta_L(\Pi_{K+1}) < \beta_H(\Pi_K) < \beta_H(\Pi_{K+1})$$




Fig. 3 Optimal Equilibrium Protocols (red (thick)— Π_K ; blue (thin)—other protocols) (color figure online)

Consider a discount factor β with $\beta_L(\Pi_K) < \beta < \beta_L(\Pi_{K+1})$. By construction, $\Pi_K = (K/2, \sigma_K)$ is an equilibrium and $\Pi_{K+1} = ((K + 1)/2, \sigma_{K+1})$ is not, so Π_K is the most efficient equilibrium protocol among the protocols $\Pi_{K'}$. However, these are not the only protocols: if we seek the most efficient among all equilibrium protocols, we must also consider protocols $(\alpha, \sigma_{K'})$ for values of α other than $\alpha = K'/2$. However, for discount factors $\beta < \beta_L(\Pi_{K+1})$ for which $|\beta_L(\Pi_{K+1}) - \beta|$ is sufficiently small, there will be token supplies $\alpha < (K + 1)/2$ for which $|(K + 1)/2 - \alpha|$ is as small as we like and for which (α, σ_{K+1}) is an equilibrium protocol. If $|(K + 1)/2 - \alpha|$ is small, then the invariant distributions for (α, σ_{K+1}) and for $((K + 1)/2, \sigma_{K+1})$ will be close, and hence the efficiency of (α, σ_{K+1}) will be almost equal to the efficiency of $((K + 1)/2, \sigma_{K+1}) = \Pi_{K+1}$. Since the efficiency of Π_{K+1} is strictly greater than the efficiency of Π_K this means that (α, σ_{K+1}) is an equilibrium protocol that is more efficient than Π_K . In other words, for discount factors less than but very close to β_{K+1} , $K/2$ is not the optimal quantity of money.

As this discussion illustrates, it is crucial to the design problem that the designer be able to choose the quantity of money α , since it is through α that the designer controls efficiency (social welfare).

5.2 Choosing the right protocol

The reader may wonder why we have put so much emphasis on choosing the right protocol. As Fig. 3 already shows, the reason is simple: choosing the wrong protocol can result in an enormous efficiency loss. Figure 4, which compares efficiency of the most efficient protocol with efficiency of a protocol for which the strategic threshold is constrained to be $K = 3$, makes this point in an even starker way: as the reader will see, except for a small range of discount factors, the efficiency loss is enormous.




Fig. 4 Inefficient Equilibrium Protocols (red (thick)— Π_3 ; blue (thin)—optimal equilibrium protocols) (color figure online)

6 Conclusion

In this paper, we have analyzed in some detail a simple, practicable, and distributed method to incentivize trade in online environments through the use of (electronic) tokens. We have shown that when agents are patient, the method we offer can achieve outcomes that are nearly efficient, provided the right protocol (supply of tokens and recommended threshold) is chosen, but that equilibrium and efficiency are both sensitive to the precise choice of protocol. Surprisingly, the “optimal” supply of tokens need *not* be half the recommended threshold; this conclusion, and others, and much of the difficulty of our arguments are a consequence of our allowing agents to accumulate as many tokens as they wish, rather than imposing an exogenous bound on token holdings (which is common in the literature).

Our analysis is silent about convergence to the steady state. In particular, we do not know whether the recommended strategies would lead to convergence to the invariant distribution for all initial token distributions or for some particular token distributions. Berentsen (2002) proves convergence under some conditions, but in a continuous time model in which token holdings are subject to an exogenous bound; we have already noted that the latter is a strong (and, in our view, unrealistic) assumption. Another point is worth making as well. By definition, the recommended strategy is a best reply when the system is in the steady state, but the recommended strategy need not be a best reply—and very likely is *not* a best reply—when the system is *not* in the steady state—so why should agents follow it?

We have assumed that service and tokens are both indivisible. This seems a natural assumption given the environment in which we are interested because a partial file is usually worthless by itself and because there is no (extant) technology for online exchange of fractional tokens. The assumption that service and tokens are exchanged one-for-one is a genuine restriction. It is conceivable that there would exist an equilibrium in which different quantities of tokens sometimes change hands, and such equilibria (if they do exist) might be more efficient than the ones we consider here. Determining whether such equilibria exist and characterizing them (if they do exist) seems a daunting task that none of the literature seems to have addressed.¹¹

We have considered the simplest setting, in which agents are identical, all files are equally valuable, and no errors occur. In a more realistic setting, we would need to take account of heterogeneous agents and files and allow for the possibilities of errors (in transmission of files or exchange of tokens or both). We have followed here the well-known adage “one has to start somewhere”—but we are keenly aware that there is much more work to be done.

Appendix: Proofs

Proof of Theorem 1 We first estimate $V(n + 1) - V(n)$ (for $n \geq 0$) which is the loss from having one less token. To this end, fix an optimal Markov strategy (τ, σ) . We

¹¹ Zhou (1999) considers equilibria for various prices, but all the equilibria she studies are assumed to have a single price and agents holding in these equilibria are only in integral multiples of the single price.

define a history-dependent strategy (τ', σ') and estimate the expected utility to an agent who begins with n tokens and follows (τ', σ') ; this is a lower bound on $V(n)$. The strategy (τ', σ') is most easily described in the following way: Begin by following the behavior prescribed by the strategy (σ, τ) but for an agent who holds one more token than is actually held, i.e., $(\tau', \sigma')(h) = (\tau, \sigma)(N(h) + 1)$. If it never happens that the agent holds 0 tokens, requests service, and is matched with an agent who is willing to provide service, then continue in this way forever. If it does happen that the agent holds 0 tokens, requests service, and is matched with an agent who is willing to provide service, then service is not provided in that period (because the agent cannot pay) and after that period $(\tau', \sigma') = (\tau, \sigma)$. In other words, the agent behaves “as if” he held one more token than actually held until the first time such behavior results in requesting service, being offered service, and being unable to pay for service; after that point, revert to (τ, σ) . The point to keep in mind is that if a moment of deviation occurs, then an agent with one more token would hold exactly 1 token, would request and receive service, and in the next period would have 0 tokens—so that reverting to (τ, σ) is possible. Beginning with n tokens and following the strategy (τ', σ') yields the same string of payoffs as beginning with $n + 1$ tokens and following the strategy (τ, σ) except in the single period in which deviation occurs; in that period, the expected loss of utility is at most $b\rho$. Hence, the expected utility from beginning with n tokens and following the strategy (τ', σ') yields utility at least $V(n + 1) - b\rho$. Hence, $V(n + 1) - V(n) \leq b\rho < b < b/\beta$. However, this is the incentive compatibility condition that guarantees that an agent strictly prefers to request service when holding $n + 1$ tokens, so the proof is complete. \square

At this point, it is convenient to collect some notation and isolate two technical results. Fix ρ, b, c, μ, v and consider a Markov strategy σ . For each k , let $V_\sigma(k, \beta)$ be the value of following σ when the initial token holding is k and the discount factor is β . As with the optimal value function V defined in the text, the value function V_σ can be defined by a recursive system of equations:

$$\begin{aligned}
 V_\sigma(0, \beta) &= \rho\sigma(0)[(1 - \mu)(-c + \beta V_\sigma(1, \beta)) \\
 &\quad + \rho[1 - \sigma(0)]\beta V_\sigma(0, \beta) + (1 - 2\rho)\beta V_\sigma(0, \beta)] \\
 V_\sigma(k, \beta) &= \rho[(1 - v)(b + \beta V_\sigma(k - 1, \beta) + v\beta V_\sigma(k, \beta))] \\
 &\quad + \rho\sigma(k)[(1 - \mu)(-c + \beta V_\sigma(k + 1, \beta)) + \mu\beta V_\sigma(k, \beta)] \\
 &\quad + \rho[1 - \sigma(k)]\beta V_\sigma(k, \beta) + (1 - 2\rho)\beta V_\sigma(k, \beta) \\
 &\quad \text{for } k > 0
 \end{aligned} \tag{9}$$

From the value function, we define the marginal utilities

$$M_\sigma(k, \beta) = V_\sigma(k + 1, \beta) - V_\sigma(k, \beta) \tag{10}$$

If β is fixed/understood, we simplify notation by writing $V_\sigma(k) = V_\sigma(k, \beta)$ and $M_\sigma(k) = M_\sigma(k, \beta)$.

It is also convenient to introduce some auxiliary parameters:

$$\begin{aligned}
 \phi_l &= -(1 - \nu)\rho\beta \\
 \phi_c &= 1 - \beta + ((1 - \nu) + (1 - \mu))\rho\beta \\
 \phi_r &= -(1 - \mu)\rho\beta
 \end{aligned}
 \tag{11}$$

We note the signs of these parameters and various combinations:

$$\begin{aligned}
 \phi_l < 0, \quad \phi_c > 0, \quad \phi_r < 0 \\
 \phi_l + \phi_c + \phi_r > 0, \quad \phi_l + \phi_c > 0, \quad \phi_r + \phi_c > 0
 \end{aligned}
 \tag{12}$$

Using these auxiliary parameters and the recursion relations for V_σ and performing some simple algebraic manipulations yields a useful matrix representation involving marginals that we will use frequently:

$$\begin{bmatrix} \phi_c & \phi_r & 0 & \cdots & 0 \\ \phi_l & \phi_c & \phi_r & 0 & \vdots \\ 0 & \phi_l & \phi_c & \phi_r & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ 0 & \cdots & 0 & \phi_l & \phi_c \end{bmatrix}_{K \times K} \begin{bmatrix} M_\sigma(0) \\ M_\sigma(1) \\ \vdots \\ M_\sigma(K_1 - 1) \end{bmatrix} = \begin{bmatrix} (1 - \nu)\rho b \\ 0 \\ \vdots \\ 0 \\ (1 - \mu)\rho c \end{bmatrix}
 \tag{13}$$

In short form, write this matrix representation as

$$\Phi \mathbf{M} = \mathbf{u}
 \tag{14}$$

Lemma 1 Fix ρ, b, c, μ, ν and β . Let σ be a Markov strategy with the property that

$$\sigma(k) = \begin{cases} 1 & \text{if } 0 \leq k < K_1 \\ 0 & \text{if } K_1 \leq k < K_2 \end{cases}$$

Then,

- (i) if $0 \leq k < K_2$ then $M_\sigma(k) > 0$
- (ii) in the range $0 \leq k < K_1$, M_σ is either increasing, decreasing or decreasing then increasing
- (iii) if $M_\sigma(K_1 - 1) \geq c/\beta$ then

$$M_\sigma(0) > M_\sigma(1) > \cdots > M_\sigma(K_1 - 2) \geq M_\sigma(K_1 - 1) \geq c/\beta$$

Proof We first consider the token holding levels $0 \leq k < K_1$. We make use of the matrix representation (13).

To prove (i), we first show that $M_\sigma(k) > 0$ for $0 \leq k < K_1$. If $K_1 < 3$, this follows by simply solving the matrix representation, so we henceforward assume $K_1 \geq 3$. If there exists a token holding level k^* with $0 \leq k^* < K_1$ such that $M_\sigma(k^*) \leq 0$, then one of the following must hold: either (a) there two consecutive such token holding

levels, or (b) the marginal payoffs of the neighboring token holding levels are both positive. We consider these cases separately.

- (a) In this case, there exists k^* such that $M_\sigma(k^*), M_\sigma(k^* + 1)$ are both non-positive. Of these, one is at least as big; say $M_\sigma(k^*) \geq M_\sigma(k^* + 1)$. From the identities above we see that

$$\begin{aligned} M_\sigma(k^* + 2) &= \frac{\phi_l M_\sigma(k^*) + \phi_c M_\sigma(k^* + 1)}{-\phi_r} \\ &\leq \frac{(\phi_l + \phi_c) M_\sigma(k^* + 1)}{-\phi_r} \\ &\leq M_\sigma(k^* + 1) \end{aligned}$$

Proceeding inductively, it follows that

$$0 \geq M_\sigma(k^*) \geq M_\sigma(k^* + 1) \geq \dots \geq M_\sigma(K_1 - 1)$$

Moreover,

$$\begin{aligned} \phi_c M_\sigma(K_1 - 1) &= (1 - \mu)\rho c - \phi_l M_\sigma(K_1 - 2) \\ &> -\phi_l M_\sigma(K_1 - 2) \\ &> -\phi_l M(K - 1) \end{aligned}$$

This requires $\phi_c < -\phi_l$ which contradicts the sign relations (12). The argument when $M_\sigma(k^* + 1) \geq M_\sigma(k^*)$ is similar and is left for the reader.

- (b) In this case, there exists k^* such that $M_\sigma(k^* - 1) > 0, M_\sigma(k^*) \leq 0, M_\sigma(k^* + 1) > 0$. This entails

$$\phi_l M_\sigma(k^* - 1) + \phi_c M_\sigma(k^*) + \phi_r M_\sigma(k^* + 1) < 0 \tag{15}$$

which again contradicts the sign relations (12).

From the above, we conclude $M_\sigma(k) > 0$ for $0 < k < K_1$. To see that $M_\sigma(0) > 0$ note that

$$-\phi_r M_\sigma(1) = \phi_c M_\sigma(0) - (1 - \nu)\rho b < -\phi_r M_\sigma(0) \tag{16}$$

Therefore, $M_\sigma(1) < M_\sigma(0)$, so $M_\sigma(0) > 0$, as desired.

Finally, to see that $M_\sigma(k) > 0$ for $K_1 \leq k < K_2$, apply the recursion equations (9) to obtain

$$\phi_l M_\sigma(k - 1) + (\phi_c + \phi_r) M_\sigma(k) = 0 \tag{17}$$

We know that $M_\sigma(K_1 - 1) > 0$ so the sign relations (12) imply that $M_\sigma(K_1) > 0$ as well. Now it follows inductively that $M_\sigma(k) > 0$ for $K_1 \leq k < K_2$. This completes the proof of (i).

To prove (ii), it is enough to show that M_σ has no local maximum for $0 < k < K_1$. If M had a local maximum k^* in this range, we would have $M_\sigma(k^*) \geq M_\sigma(k^* - 1)$ and $M_\sigma(k^*) \geq M_\sigma(k^* + 1)$. However, algebraic manipulation yields the inequalities

$$\begin{aligned} M_\sigma(k^*) &= \frac{-\phi_l M_\sigma(k^* - 1) - \phi_r M_\sigma(k^* + 1)}{\phi_c} \\ &\leq \frac{-\phi_l - \phi_r}{\phi_c} M_\sigma(k^*) \\ &< M_\sigma(k^*) \end{aligned}$$

which is a contradiction. This establishes (ii)

To prove (iii), first manipulate the matrix identity (13) to obtain:

$$\begin{aligned} &(1 - v)\rho\beta M_\sigma(K_1 - 2) \\ &= (1 - \beta + ((1 - v) + (1 - \mu))\rho\beta)M_\sigma(K - 1) - (1 - \mu)\rho c \\ &\geq (1 - \beta + (1 - v)\rho\beta)M_\sigma(K_1 - 1) \geq (1 - v)\rho\beta M_\sigma(K_1 - 1) \end{aligned} \tag{18}$$

In view of (ii), the marginal payoffs are decreasing, so this establishes (iii).

Lemma 2 Fix ρ, b, c and a threshold protocol $\Pi = (\alpha, \sigma_K)$ with corresponding μ^Π, v^Π . The marginal utility $M_{\sigma_K}(k, \beta)$ is strictly increasing in the discount factor β , i.e., if $0 \leq \beta_1 < \beta_2 < 1$, then,

$$M_{\sigma_K}(k, \beta_1) < M_{\sigma_K}(k, \beta_2) \text{ for all } k \tag{19}$$

Proof To economize slightly on notation, we write $\sigma = \sigma_K$. We present the proof in three steps.

In Step 1, we prove that if there exist $0 < K_1 \leq K_2 < K - 1$ such that $\forall k \in [K_1, K_2], M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$, then at least one of the following is true, $M_\sigma(K_1 - 1, \beta_1) \geq M_\sigma(K_1 - 1, \beta_2)$ or $M_\sigma(K_2 + 1, \beta_1) \geq M_\sigma(K_2 + 1, \beta_2)$.

In Step 2, we prove that if there exists a $k^* \in [0, K - 1]$ such that $M_\sigma(k^*, \beta_1) \geq M_\sigma(k^*, \beta_2)$, then for all $k \in [0, K - 1], M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$. Step 2 uses the result of Step 1.

In Step 3, we disprove the possibility that $k \in [0, K - 1], M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$.

Step 2 and Step 3 together show a contradiction and therefore, $k \in [0, K - 1], M_\sigma(k, \beta_1) < M_\sigma(k, \beta_2)$.

Step 1 We assert that if there are indices $0 < K_1 \leq K_2 < K - 1$ such that $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for all $K_1 \leq k \leq K_2$ then at least one of the following must hold:

- (A) $M_\sigma(K_1 - 1, \beta_1) \geq M_\sigma(K_1 - 1, \beta_2)$
- (B) or $M_\sigma(K_2 + 1, \beta_1) \geq M_\sigma(K_2 + 1, \beta_2)$.

To see this, note that simple manipulations of the matrix representation (13) yield

– if $K_2 = K_1$ then

$$\begin{aligned} &(1 - \nu)\rho M_\sigma(K_1 - 1, \beta) + (1 - \mu)\rho M_\sigma(K_2 + 1, \beta) \\ &= (1/\beta - 1 + ((1 - \nu) + (1 - \mu))\rho)M_\sigma(K_1, \beta) \end{aligned}$$

– if $K_2 > K_1$ then

$$\begin{aligned} &(1 - \nu)\rho M_\sigma(K_1 - 1, \beta) + (1 - \mu)\rho M_\sigma(K_2 + 1, \beta) \\ &= (1/\beta - 1 + (1 - \mu)\rho)M_\sigma(K_1, \beta) \\ &= +(1/\beta - 1)[M_\sigma(K_1 + 1, \beta) + \dots + M_\sigma(K_2 - 1, \beta)] \\ &= +(1/\beta - 1 + (1 - \nu)\rho)M_\sigma(K_2, \beta) \end{aligned}$$

Since $\beta_1 < \beta_2$ and we have assumed $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for $0 < K_1 \leq K_2 < K - 1$, in each of the cases above the right-hand side is larger when $\beta = \beta_1$ than when $\beta = \beta_2$. Because the terms in the left-hand sides are positive, it follows that at least one of (A), (B) must hold, as asserted.

Step 2 We assert first that if there is a k^* , $0 \leq k^* \leq K_1$ such that $M_\sigma(k^*, \beta_1) \geq M_\sigma(k^*, \beta_2)$, then at least one of the following must hold:

- (C) there exists some K_3 , $0 \leq K_3 \leq K_1$, such that $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for all k , $0 \leq k \leq K_3$
- (D) there exists some K_4 , $0 \leq K_4 \leq K_1$, such that $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for all k , $K_4 \leq k \leq K - 1$

To see this, note first that if $k^* = 0$ satisfies the hypothesis, then (C) holds with $K_3 = 0$ and that if $k^* = K - 1$ satisfies the hypothesis, then (D) holds with $K_4 = K - 1$. Hence, it suffices to consider a k^* , $0 < k^* < K - 1$, that satisfies the hypothesis. We now make use of Step 1. Set $K_1 = K_2 = k^*$. Applying Step 1 once increases the token holding interval where $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ by 1. Let K_1 and K_2 be the new end points of the interval and apply Step 1 again. Continuing in this way, we come eventually to a point where either $K_1 = 0$ or $K_2 = K - 1$. If $K_1 = 0$, set $K_3 = K_2$ and note that (C) holds. If $K_2 = K - 1$, set $K_4 = K - 1$ and note that (D) holds

We now show that either (C) or (D) leads to the desired conclusion. Consider (C) first. Using the matrix representation (13), we obtain

$$\begin{aligned} &(1 - \nu)\rho\beta M_\sigma(K_1 + 1, \beta) + (1 - \nu)\rho b \\ &= [1 - (1 - (1 - \mu)\rho)\beta]M_\sigma(0, \beta) \\ &\quad + (1 - \beta)[M_\sigma(1, \beta) + \dots + M_\sigma(K_1 - 1, \beta)] \\ &\quad + [1 - (1 - (1 - \nu)\rho)\beta]M_\sigma(K_1, \beta) \end{aligned}$$

The right-hand side is bigger when $\beta = \beta_1$ than when $\beta = \beta_2$. Therefore, $M_\sigma(K_1 + 1, \beta_1) \geq M_\sigma(K_1 + 1, \beta_2)$. By induction, $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for all k , $0 \leq k \leq K - 1$.

Now consider (D). Using the matrix representation (13), we obtain

$$\begin{aligned} & (1 - \mu)\rho\beta M_\sigma(K_2 - 1, \beta) + (1 - \mu)\rho c \\ &= [1 - (1 - (1 - \nu)\rho)\beta]M_\sigma(K - 1, \beta) \\ & \quad + (1 - \beta)[M_\sigma(K - 2, \beta) + \dots + M_\sigma(K_2 + 1, \beta)] \\ & \quad + [1 - (1 - (1 - \mu)\rho)\beta]M_\sigma(K_2, \beta) \end{aligned}$$

The right-hand side is bigger when $\beta = \beta_1$ than when $\beta = \beta_2$. Therefore, $M_\sigma(K_2 - 1, \beta_1) \geq M_\sigma(K_2 - 1, \beta_2)$. By induction, $M_\sigma(k, \beta_1) \geq M_\sigma(k, \beta_2)$ for all $k, 0 \leq k \leq K - 1$.

Taking (C) and (D) together completes Step 2.

Step 3 Using the matrix representation (13), we obtain

$$\begin{aligned} & [1 - (1 - (1 - \mu)\rho)\beta]M_\sigma(0, \beta) \\ & \quad + (1 - \beta)[M_\sigma(1, \beta) + \dots + M_\sigma(K_1 - 1, \beta)] \\ & \quad + [1 - (1 - (1 - \nu)\rho)\beta]M_\sigma(K - 1, \beta) \\ &= (1 - \nu)\rho b + (1 - \mu)\rho c \end{aligned}$$

In view of Step 2, the left-hand side is bigger when $\beta = \beta_1$ than when $\beta = \beta_2$. However, the right-hand side is independent of β , so this is a contradiction. We conclude that $M_\sigma(k, \beta_1) < M_\sigma(k, \beta_2)$ for every $k, 0 \leq k \leq K - 1$. \square

Proof of Theorem 2 Fix β . The Markov strategy σ is optimal if and only if it satisfies the Bellman optimality conditions:

$$\beta(V_\sigma(k + 1) - V_\sigma(k)) \geq c, \text{ if } \sigma(k) = 1 \tag{20}$$

$$\beta(V_\sigma(k + 1) - V_\sigma(k)) \leq c, \text{ if } \sigma(k) = 0 \tag{21}$$

If σ is not a threshold strategy, there must exist integers $K_1 < K_2$ such that

$$\begin{aligned} \sigma(k) &= 1, \quad 0 \leq k < K_1 \\ \sigma(k) &= 0, \quad K_1 \leq k < K_2 \\ \sigma(k) &= 1, \quad k = K_2 \end{aligned} \tag{22}$$

We will show that the Bellman optimality conditions are violated at K_2 and $K_2 - 1$. To this end, let K_3 be the smallest integer greater than K_2 for which $\sigma(K_3) = 0$. (Such an integer exists because it cannot be optimal to serve when the token holding is sufficiently high.) Thus, $\sigma(k) = 1$, for $K_2 \leq k < K_3$ and $M_\sigma(K_3 - 1) \geq c/\beta$. Following σ ,

$$M_\sigma(K_3 - 2) = [(1 - \mu)\rho c - \phi_c M_\sigma(K_3 - 1)]/\phi_l > M_\sigma(K_3 - 1) \geq c/\beta \tag{23}$$

An inductive argument shows that $M_\sigma(K_2) > M_\sigma(K_2 + 1) \geq c/\beta$. According to the recursion equations (9), we have

$$M_{\sigma}(K_2 - 1) = (\phi_c M_{\sigma}(K_2) + \phi_r M_{\sigma}(K_2 + 1)) / (-\phi_l) > c/\beta$$

which is a contradiction. We conclude that a non-threshold strategy cannot be optimal; equivalently, only threshold strategies can be optimal strategies.

It remains to show that the only possible optimal threshold strategies have adjacent thresholds. Consider first two threshold strategies with consecutive thresholds K and $K + 1$. We assert that

$$M_{\sigma_K}(K) < c/\beta \Leftrightarrow M_{\sigma_{K+1}}(K) < c/\beta \tag{24}$$

We prove direction “ \Rightarrow ”; the “ \Leftarrow ” direction is similar and left to the reader. Suppose instead that $M_{\sigma_{K+1}}(K) \geq c/\beta$. It follows that $-\phi_r M_{\sigma_{K+1}}(K) \geq (1 - \mu)\rho c$. If we delete the last line in the matrix equation (13) for σ_{K+1} and move $M_{\sigma_{K+1}}(K)$ to the right-hand side, we get another matrix equation

$$\Phi_{K \times K} \mathbf{M}_{\sigma_{K+1}} = \tilde{\mathbf{u}}$$

where $\tilde{\mathbf{u}} = ((1 - \nu)\rho b, 0, \dots, 0, -\phi_r M_{\sigma_{K+1}}(K))^T$. For the threshold K , $\Phi_{K \times K} \mathbf{M}_{\sigma_K} = \mathbf{u}$. Therefore,

$$\Phi_{K \times K} (\mathbf{M}_{\sigma_{K+1}} - \mathbf{M}_{\sigma_K}) = \tilde{\mathbf{u}} - \mathbf{u} \tag{25}$$

Lemma 1 guarantees that $\tilde{\mathbf{u}} - \mathbf{u} \geq 0$, so $\mathbf{M}_{\sigma_{K+1}} \geq \mathbf{M}_{\sigma_K}$. That is, $M_{\sigma_{K+1}}(k) \geq M_{\sigma_K}(k)$ for $0 \leq k \leq K - 1$. Because $M_{\sigma_{K+1}}(K) \geq c/\beta > M_{\sigma_K}(K)$, it follows that $M_{\sigma_{K+1}}(k) \geq M_{\sigma_K}(k)$ for $0 \leq k \leq K$. According to the matrix equation, the following identity holds for both $\sigma = \sigma_K$ and $\sigma = \sigma_{K+1}$:

$$\begin{aligned} (1 - \nu)\rho b + (1 - \mu)\rho c &= (1 - \beta + (1 - \mu)\rho\beta)M_{\sigma}(0) \\ &+ (1 - \beta) \sum_{k=1}^{K-1} M_{\sigma}(k) + (1 - \beta + (1 - \nu)\rho\beta)M_{\sigma}(K) \end{aligned} \tag{26}$$

This is a contradiction so we have established the direction \Rightarrow , as desired.

It follows directly from the matrix identity that

$$M_{\sigma_K}(K) = c/\beta \Leftrightarrow M_{\sigma_{K+1}}(K) = c/\beta$$

Hence

$$M_{\sigma_K}(K) > c/\beta \Leftrightarrow M_{\sigma_{K+1}}(K) > c/\beta \tag{27}$$

We now assert that if $\tilde{K} > K$ then

$$M_{\sigma_K}(K) < c/\beta \Rightarrow M_{\sigma_{\tilde{K}}}(K - 1) < c/\beta \tag{28}$$

We have already shown that this is true when $\tilde{K} = K + 1$; i.e. $M_{\sigma_{K+1}}(K) < c/\beta$. Consider $\tilde{K} = K + 2$. Of $M_{\sigma_{K+2}}(K + 1) \geq c/\beta$, then (27) implies that $M_{\sigma_{K+1}}(K + 1) \geq c/\beta$. Therefore, $M_{\sigma_{K+1}}(K + 1) > M_{\sigma_{K+1}}(K)$. This is a contradiction to $M_{\sigma_{K+1}}(K + 1) < M_{\sigma_{K+1}}(K)$. Following inductively we obtain the assertion (28).

A similar argument (which we omit) shows that:

$$M_{\sigma_K}(K - 1) > c/\beta \Rightarrow M_{\sigma_{\tilde{K}}}(\tilde{K}) > c/\beta, \forall \tilde{K} < K \tag{29}$$

Finally, suppose σ_K is an optimal threshold strategy. Then, $M_{\sigma_K}(K - 1) \geq c/\beta$ and $M_{\sigma_K}(K) \leq c/\beta$. If the equalities hold strictly, (28) and (29) guarantee that σ_K is the only optimal threshold strategy. If $M_{\sigma_K}(K - 1) = c/\beta$ (and hence, $M_{\sigma_K}(K) < c/\beta$), only σ_K and σ_{K-1} are optimal threshold strategies. If $M_{\sigma_K}(K) = c/\beta$ (and hence, $M_{\sigma_K}(K - 1) > c/\beta$), only σ_K and σ_{K+1} are optimal threshold strategies. This completes the proof. \square

Proof of Theorem 3 This follows immediately from the representation of η_+ and the definition of invariance. \square

Proof of Theorem 4 Given a protocol $\Pi = (\alpha, \sigma)$, let η^Π be the unique invariant distribution; let μ^Π be the fraction of agents who have no tokens and ν^Π the fraction of agents who do not provide service; these depend only on Π and not on the population parameters. If $\sigma = \sum \gamma(K)\sigma_K$ is a best response given the population parameters and μ^Π, ν^Π, γ must put strictly positive weight only on threshold strategies σ_K that are pure best responses. In view of Theorem 2, there are at most two threshold strategies that are pure best responses and they are at adjacent thresholds. That is, σ is either a pure threshold strategy or a mixture of two adjacent threshold strategies, as asserted. \square

Proof of Theorem 5 Suppose to the contrary that $\Pi = (\alpha, \sigma)$ is a robust equilibrium protocol and that $\sigma = \sum \gamma(K)\sigma_K$ is a proper mixed strategy, so that $\gamma(K) > 0$ for at least two values of the threshold K . Let μ^Π be the fraction of agents who have no tokens and ν^Π the fraction of agents who do not provide service; these depend only on Π and not on the population parameters. In view of Theorem 4, σ must assign positive probability only to two adjacent threshold strategies; say $\sigma = \gamma(K)\sigma_K + \gamma(K + 1)\sigma_{K+1}$ with $\gamma(K) > 0$ and $\gamma(K + 1) > 0$, and both σ_K, σ_{K+1} must be best responses. Because $\sigma_K(K + 1) = 0$ and $\sigma_{K+1}(K + 1) = 1$, equations (8), (9) (which provide necessary and sufficient conditions for optimality in terms of the true value function) entail that

$$\begin{aligned} -c + \beta V_{K+1} &\leq \beta V_K \\ -c + \beta V_{K+1} &\geq \beta V_K \end{aligned}$$

Hence, $-c + \beta V_{K+1} = \beta V_K$. Because σ_K is a best response, the value functions V_{σ_K} must coincide with the true value function V . Hence, an agent following σ_K must be indifferent to providing service when holding K tokens. However, if β increases slightly M_{σ_K} also increases, whence an agent following σ_K must strictly prefer to

provide service. In other words, when β increases slightly, σ_K can no longer be a best response and σ_K can no longer be an equilibrium protocol. This is a contradiction, so we conclude that a robust equilibrium protocol Π cannot involve proper mixed strategies, as asserted. \square

Proof of Theorem 6 We divide the proof of (i) into several steps.

Step 1 We first prove there exists $\beta^L \in [0, 1)$ such that

$$\begin{aligned} M_\sigma(K - 1, \beta) &< \frac{c}{\beta} \text{ for } \beta < \beta^L \\ M_\sigma(K - 1, \beta^L) &= \frac{c}{\beta} \\ M_\sigma(K - 1, \beta) &> \frac{c}{\beta} \text{ for } \beta > \beta^L \end{aligned}$$

To see this, define the auxiliary function

$$F(\beta) = M_\sigma(K - 1, \beta) - \frac{c}{\beta}$$

F is evidently continuous. Lemma 2 guarantees that $M_\sigma(K - 1, \beta)$ is strictly increasing in β , so $F(\beta)$ is also strictly increasing in β as well. We show that $F(1) > 0$ and $\lim_{\beta \rightarrow 0} F(\beta) < 0$ and then apply the intermediate value theorem to find β^L .

To see that $F(1) > 0$, note first that the coefficients in the left-hand matrix of (13) are simply $\phi_l = -\rho(1 - \nu)$, $\phi_c = \rho(1 - \nu + 1 - \mu)$ and $\phi_r = \rho(1 - \mu)$. We split the matrix \mathbf{M}_{σ_K} in two parts. To do this, write

$$\begin{aligned} \mathbf{u}' &= (\rho(1 - \nu)c \quad 0 \quad \dots \quad 0 \quad \rho(1 - \mu)c)^T \\ \mathbf{u}'' &= (\rho(1 - \nu)(b - c) \quad 0 \quad \dots \quad 0 \quad 0)^T \end{aligned} \tag{30}$$

and define $\mathbf{M}'_{\sigma_K}, \mathbf{M}''_{\sigma_K}$ to be the solutions to the equations

$$\Phi \mathbf{M}'_{\sigma_K} = \mathbf{u}', \quad \Phi \mathbf{M}''_{\sigma_K} = \mathbf{u}'' \tag{31}$$

Note that $\mathbf{M}_{\sigma_K} = \mathbf{M}'_{\sigma_K} + \mathbf{M}''_{\sigma_K}$ and \mathbf{M}_{σ_K} is the solution to (13). It is easy to check that \mathbf{M}'_{σ_K} is a constant matrix: $M'_{\sigma_K}(k) = c$ for $0 \leq k < K - 1$. Lemma 1 guarantees that the entries of \mathbf{M}''_{σ_K} are strictly positive: $M''_{\sigma_K}(k) > 0$ for $0 \leq k < K - 1$. Hence the entries of \mathbf{M}_{σ_K} are strictly greater than c : $M_\sigma(k) > c$ for $0 \leq k < K - 1$. In particular, $F(1) > 0$.

To see that $\lim_{\beta \rightarrow 0} F(\beta) < 0$, suppose not. Because F is strictly increasing, this means $F(\beta) \geq 0$ for every $\beta \in (0, 1]$, which entails that $M_\sigma(k) \geq \frac{c}{\beta}$ for $0 \leq k < K - 1$. Summing the rows in (13) yields:

$$\rho(1 - \nu)b + \rho(1 - \mu)c > K(1 - \beta)\frac{c}{\beta} = \frac{Kc}{\beta} - Kc \tag{32}$$

Note that Kc/β flows up as $\beta \rightarrow 0$, so this is impossible. We conclude that $\lim_{\beta \rightarrow 0} F(\beta) < 0$, as asserted.

Because F is strictly increasing, the intermediate value theorem guarantees that we can find a unique β^L such that

$$\begin{aligned} F(\beta) &< 0 && \text{for } \beta < \beta^L \\ F(\beta^L) &= 0 \\ F(\beta) &> 0 && \text{for } \beta > \beta^L \end{aligned}$$

The definition of F yields the desired property of β^L

Step 2 Next we prove there exists $\beta^H \in (\beta^L, 1)$ such that if $\beta \in [0, \beta^H]$ then

$$\begin{aligned} M_{\sigma_K, \beta}(K - 1) &< \frac{\phi_c + \phi_r}{-\phi_l} \frac{c}{\beta} && \text{for } \beta < \beta^H \\ M_{\sigma_K, \beta^H}(K - 1) &= \frac{\phi_c + \phi_r}{-\phi_l} \frac{c}{\beta} \\ M_{\sigma_K, \beta}(K - 1) &> \frac{\phi_c + \phi_r}{-\phi_l} \frac{c}{\beta} && \text{for } \beta > \beta^H \end{aligned}$$

To see this, note first that $\frac{\phi_c + \phi_r}{-\phi_l} \frac{c}{\beta} = \left[1 - \frac{1}{\rho(1-v)} + \frac{1}{\rho(1-v)\beta} \right] \frac{c}{\beta}$ and define another auxiliary function:

$$G(\beta) = M_{\Pi}(K - 1, \beta) - \left(1 - \frac{1}{\rho(1-v)} + \frac{1}{\rho(1-v)\beta} \right) \frac{c}{\beta}$$

G is continuous and increasing. From Step 1 it follows that $M_{\sigma_K}(K - 1, 1) > c$ so $G(1) = M_{\sigma_K}(K - 1, 1) - c > 0$. It also follows that $M_{\sigma_K}(K - 1, \beta^L) = \frac{c}{\beta^L}$; because $(1 - \frac{1}{\rho(1-v)} + \frac{1}{\rho(1-v)\beta^L}) \frac{c}{\beta^L} > \frac{1}{\beta^L}$, we conclude that $G(\beta^L) < 0$. Because G is continuous and increasing, there is a unique $\beta^H \in (\beta^L, 1)$ such that

$$\begin{aligned} G(\beta) &< 0 && \text{for } \beta < \beta^H \\ G(\beta^H) &= 0 \\ G(\beta) &> 0 && \text{for } \beta > \beta^L \end{aligned}$$

Step 3 The definitions of F, G imply that in order for Π to be an equilibrium protocol when the discount factor is β it is the necessary and sufficient condition that $F(\beta) \geq 0$ and $G(\beta) \leq 0$. Hence, Π is an equilibrium protocol when the discount factor is β exactly for $\beta \in [\beta^L, \beta^H]$.

Because F, G are continuous in all their arguments and strictly increasing, β^L, β^H , which are the zeroes of F, G , are continuous functions of the parameters as well. This completes the proof of (i).

The proof of (ii) is similar and left to the reader.

Proof of Theorem 7 We first consider (i). Fix r . Consider the two protocols $\Pi_K = (K/2, \sigma_K)$ and $\Pi_{K+1} = ((K+1)/2, \sigma_{K+1})$ and the corresponding intervals $[\beta_1^L, \beta_1^H]$ and $[\beta_2^L, \beta_2^H]$ of discount factors that sustain equilibrium. We need to show that

$$\beta_1^L < \beta_2^L < \beta_1^H < \beta_2^H$$

(The sustainable ranges for two consecutive threshold protocols overlap but are not nested). There are three inequalities to be established; we carry out the analyses in (A), (B), and (C) below.

(A) To prove $\beta_2^L > \beta_1^L$, write $\beta = \beta_1^L$. We show that $M_{\sigma_{K+1}}(K) < \frac{c}{\beta}$. To see this, suppose not; i.e. $M_{\sigma_{K+1}}(K) \geq \frac{c}{\beta}$. The construction of β_1^L guarantees that $M_{\sigma_K}(K-1) = c/\beta$. We will use this inequality and equality to show that *all* marginal payoffs of Π_{K+1} so large that they violate the restrictions imposed by the bounded benefit b and cost c .

To simplify the notation, let $\omega_X = \frac{X+1}{X}(\frac{1}{\beta} - 1)\frac{1}{\rho}$. Note $\omega_{K+1} < \omega_K$. Then the matrix identity (13) becomes:

$$\begin{bmatrix} \omega_X + 2 & -1 & 0 & \cdots & 0 \\ -1 & \omega_X + 2 & -1 & 0 & \vdots \\ 0 & -1 & \omega_X + 2 & -1 & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ 0 & \cdots & 0 & -1 & \omega_X + 2 \end{bmatrix}_{X \times X} \begin{bmatrix} M_{\sigma_X}(0) \\ M_{\sigma_X}(1) \\ \vdots \\ M_{\sigma_X}(X-1) \end{bmatrix} = \begin{bmatrix} b/\beta \\ 0 \\ \vdots \\ 0 \\ c/\beta \end{bmatrix} \quad (33)$$

Suppose $M_{\sigma_{K+1}}(K) \geq M_{\sigma_K}(K-1) = \frac{c}{\beta}$. We investigate the relation between $M_{\sigma_{K+1}}(K-1)$ and $M_{\sigma_K}(K-2)$. Using the matrix identity,

$$\begin{aligned} \frac{M_{\sigma_{K+1}}(K-1)}{M_{\sigma_K}(K-2)} &= \frac{(\omega_{K+1} + 2)M_{\sigma_{K+1}}(K) - \frac{c}{\beta}}{(\omega_K + 2)M_{\sigma_K}(K-1) - \frac{c}{\beta}} \\ &> \frac{(\omega_{K+1} + 2)M_{\sigma_{K+1}}(K)}{(\omega_K + 2)M_{\sigma_K}(K-1)} > \frac{\omega_{K+1} + 1}{\omega_K + 1} \end{aligned}$$

Moreover, if $2 \leq k \leq K-1$ then

$$\frac{M_{\sigma_{K+1}}(K-k)}{M_{\sigma_K}(K-k-1)} = \frac{(\omega_{K+1} + 1)[M_{\sigma_{K+1}}(K) + M_{\sigma_{K+1}}(K-k+1)] - \frac{c}{\beta}}{(\omega_K + 1)[M_{\sigma_K}(K-1) + M_{\sigma_K}(K-k)] - \frac{c}{\beta}}$$

By induction,

$$\begin{aligned} \frac{M_{\sigma_{K+1}}(K-k)}{M_{\sigma_K}(K-k-1)} &> \left(\frac{\omega_{K+1} + 1}{\omega_K + 1}\right)^k > \left(\frac{\omega_{K+1}}{\omega_K}\right)^k > \left(1 - \frac{1}{(K+1)^2}\right)^k \\ &> 1 - \frac{k}{(K+1)^2} > \frac{K+1}{K+2}, \forall 0 \leq k \leq K-1 \end{aligned}$$

Next we prove $M_{\sigma_{K+1}}(0) \geq M_{\sigma_K}(0)$. This is relatively easy since, if $M_{\sigma_{K+1}}(0) < M_{\sigma_K}(0)$, then using the marginal payoff matrix and by induction, $M_{\sigma_{K+1}}(K - 1) < M_{\sigma_K}(K - 1) = \frac{c}{\beta}$. This is a contradiction to $M_{\sigma_{K+1}}(K - 1) > M_{\sigma_{K+1}}(K) = \frac{c}{\beta}$. Therefore, $M_{\sigma_{K+1}}(0) \geq M_{\sigma_K}(0)$.

The marginal payoffs are bounded as follows,

$$(M_{\sigma_X}(0) + M_{\sigma_X}(X - 1)) + \omega_X \sum_{k=0}^{X-1} M_{\sigma_X}(k) = b/\beta + c/\beta \tag{34}$$

However, since

$$\begin{aligned} \omega_{K+1} \sum_{k=0}^K M_{\sigma_{K+1}}(k) &> \frac{K + 1}{K} \omega_{K+1} \sum_{k=1}^K M_{\sigma_{K+1}}(k) \\ &> \frac{K + 1}{K} \frac{K(K + 2)}{(K + 1)^2} \frac{K + 1}{K + 2} \omega_K \sum_{k=0}^{K-1} M_{\sigma_K}(k) \\ &= \omega_K \sum_{k=0}^{K-1} M_{\sigma_K}(k) \end{aligned}$$

and $M_{\sigma_{K+1}}(0) + M_{\sigma_{K+1}}(K) > M_{\sigma_K}(0) + M_{\sigma_K}(K - 1)$, a contradiction occurs. Therefore, for $\beta = \beta_1^L$, $M_{\sigma_{K+1}}(K) < \frac{c}{\beta}$. This means $\beta_2^L > \beta_1^L$. This completes (A).

(B) To prove $\beta_2^H > \beta_1^H$, let $\beta = \beta_1^H$, we need to show that the protocol Π_{K+1} must have $M_{\sigma_{K+1}}(K + 1) < c/\beta$. We use contradiction to prove this. The idea is: Suppose $M_{\sigma_{K+1}}(K + 1) \geq c/\beta$, then we show that all the marginal payoffs of Π_{K+1} are large enough such that they violate the restriction imposed by the bounded benefit b and cost c .

Suppose $M_{\sigma_{K+1}}(K + 1) \geq M_{\sigma_K}(K) = c/\beta$. According to the matrix equation, similar to part (A), by induction we can get,

$$\frac{M_{\sigma_{K+1}}(K + 1 - k)}{M_{\sigma_K}(K - k)} > \left(\frac{\omega_{K+1} + 1}{\omega_K + 1} \right)^k > \frac{(K + 1)^3}{K(K + 2)^2}, \quad \forall 0 \leq k \leq K$$

Also $M_{\sigma_{K+1}}(0) \geq M_{\sigma_K}(0)$. The marginal payoffs are bounded as follows,

$$(M_{\sigma_X}(0) + M_{\sigma_X}(X)) + \omega_X \sum_{k=0}^X M_{\sigma_X}(k) = b/\beta + c/\beta \tag{35}$$

However, since

$$\begin{aligned} \omega_{K+1} \sum_{k=0}^{K+1} M_{\sigma_{K+1}}(k) &> \frac{K+2}{K+1} \omega_{K+1} \sum_{k=1}^{K+1} M_{\sigma_{K+1}}(k) \\ &> \frac{K+2}{K+1} \frac{K(K+2)}{(K+1)^2} \frac{(K+1)^3}{K(K+2)^2} \omega_K \sum_{k=0}^K M_{\sigma_K}(k) \\ &= \omega_K \sum_{k=0}^K M_{\sigma_K}(k) \end{aligned}$$

and $M_{\sigma_{K+1}}(0) + M_{\sigma_{K+1}}(K+1) > M_{\sigma_K}(0) + M_{\sigma_K}(K)$, a contradiction occurs. Therefore, for $\beta = \beta_1^H$, $M_{\sigma_{K+1}}(K+1) < \frac{c}{\beta}$. This means $\beta_2^H > \beta_1^H$. This completes part (B).

- (C) To prove $\beta_2^L < \beta_1^H$, write $\beta = \beta_1^H$. We show that $M_{\sigma_{K+1}}(K) > M_{\sigma_K}(K) = \frac{c}{\beta}$. If not, then as in (A) we must have $M_{\sigma_{K+1}}(K) \leq M_{\sigma_K}(K) = \frac{c}{\beta}$; in that case we show $M_{\sigma_{K+1}}(k) \leq M_{\sigma_K}(k)$ for $0 \leq k \leq K$. This will again violate the restrictions imposed by b and c .

We extend the marginal payoff matrix in (33) from $K \times K$ to $(K+1) \times (K+1)$ and incorporate $M_{\sigma_K}(K)$. If $M_{\sigma_K}(K) = \frac{c}{\beta}$, such extension does not change the solution of the marginal payoffs $M_{\sigma_K}(k), \forall k \in [0, K]$. Note the new coefficient matrix has the same size of the coefficient matrix for σ_{K+1} . Suppose $M_{\sigma_{K+1}}(K) < M_{\sigma_K}(K) = \frac{c}{\beta}$. According to the matrix equation,

$$\frac{M_{\sigma_{K+1}}(K-1)}{M_{\sigma_K}(K-1)} = \frac{(\omega_{K+1} + 2)M_{\sigma_{K+1}}(K) - c/\beta}{(\omega_K + 2)M_{\sigma_K}(K) - c/\beta} < 1$$

Moreover, for $0 \leq k \leq K$, we have

$$\frac{M_{\sigma_{K+1}}(K-k)}{M_{\sigma_K}(K-k)} = \frac{(\omega_{K+1} + 1)[M_{\sigma_{K+1}}(K) + M_{\sigma_{K+1}}(K-k+1)] - c/\beta}{(\omega_K + 1)[M_{\sigma_K}(K) + M_{\sigma_K}(K-k+1)] - c/\beta}$$

By induction, $M_{\sigma_{K+1}}(k) < M_{\sigma_K}(k) 0 \leq k \leq K$. However, since

$$(M_{\sigma_X}(0) + M_{\sigma_X}(X-1)) + \omega_X \sum_{k=0}^{X-1} M_{\sigma_X}(k) = b/\beta + c/\beta \tag{36}$$

Again, the left-hand side is bigger when $X = K$ than when $X = K+1$, which is a contradiction. This completes part (C).

Combining (A), (B) and (C) establishes the desired string of inequalities. The remaining conclusions of (i) follow immediately.

The argument for (ii) is very similar and left to the reader.

Proof of Theorem 8 Fix a protocol $\Pi = (\alpha, \sigma_K)$ and let η^Π be the corresponding invariant distribution. We first find a closed form expression for η^Π . To do this, plug the strategy σ_K into the characterization of the invariant distribution given in Theorem 3. A little algebra provides an identity involving $\eta^\Pi(0)$, $\eta^\Pi(1)$, $\eta^\Pi(K)$ and a simpler recursion relationship.

$$\begin{aligned} \eta^\Pi(1) &= \left[\frac{1 - \eta^\Pi(0)}{1 - \eta^\Pi(K)} \right] \eta^\Pi(0) \\ \eta^\Pi(k) &= \left[\frac{2 - \eta^\Pi(0) - \eta^\Pi(K)}{1 - \eta^\Pi(K)} \right] \eta^\Pi(k - 1) + \left[\frac{1 - \eta^\Pi(0)}{1 - \eta^\Pi(K)} \right] \eta^\Pi(0) \eta^\Pi(k - 2) \end{aligned}$$

From this we can solve recursively, obtaining

$$\eta^\Pi(k) = \left[\frac{1 - \eta^\Pi(0)}{1 - \eta^\Pi(K)} \right]^k \tag{37}$$

for all $k = 0, 1, \dots, K$. Note that the one remaining degree of freedom is pinned down by the requirement that the total token holding be equal to α .

We next solve the following simple maximization problem:

$$\begin{aligned} \text{maximize } & E^*(x_1, x_2) = 1 - x_1 - x_2 + x_1x_2 \\ & 0 \leq x_1, x_2 \leq 1 \\ \text{subject to } & x_1(1 - x_1)^K = x_2(1 - x_2)^K \end{aligned} \tag{38}$$

To solve this problem, set $f(x) = x(1 - x)^K$. A straightforward calculus exercise shows that if $0 \leq x_1 \leq \frac{1}{K+1} \leq x_2 \leq 1$ and $f(x_1) = f(x_2)$ then:

- (a) $x_1 + x_2 \geq \frac{2}{K+1}$, with equality achieved at $x_1 = x_2 = \frac{1}{K+1}$.
- (b) $x_1x_2 \leq \frac{1}{K+1}$, with equality achieved at $x_1 = x_2 = \frac{1}{K+1}$.

Putting (a) and (b) together shows that the optimal solution to the maximization problem (38) is to have $x_1 = x_2 = \frac{1}{K+1}$ and $\max E^* = \left(1 - \frac{1}{K+1}\right)^2$.

If we take $x_1 = \mu^\Pi$, $x_2 = \nu^\Pi$ and apply the closed form solution (37) for the invariant distribution, we see that $f(x_1) = f(x_2)$. By definition, $\text{Eff}(\Pi) = E^*(x_1, x_2)$ so

$$\text{Eff}(\Pi) \leq \max E^* = \left(1 - \frac{1}{K + 1}\right)^2$$

On the other hand, if $\alpha = K/2$ then the invariant distribution has $\eta^\Pi(k) = \frac{1}{K+1}$ for all k and

$$\text{Eff}(K/2, \sigma_K) = \left(1 - \frac{1}{K + 1}\right)^2 = [K/(K + 1)]^2$$

Taken together, part (ii) and (iii) are proved.

Next fix a protocol (α, σ_K) . Let $\lceil \alpha \rceil$ be the least integer greater than or equal to α and set $K^* = 2\lceil \alpha \rceil$. There are two cases to consider.

In the first case, $K \leq K^*$.

$$\text{Eff}(\alpha, \sigma_K) \leq \left(1 - \frac{1}{K+1}\right)^2 \leq \left(1 - \frac{1}{K^*+1}\right)^2 = \left(1 - \frac{1}{2\lceil \alpha \rceil + 1}\right)^2$$

which is the desired result in the first case.

In the second case, $K > K^*$. Define the protocol $\Pi' = (\lceil \alpha \rceil, \sigma_{K^*})$; let η' be the invariant token distribution for Π' . Let $\Pi^* = (\lceil \alpha \rceil, \sigma_{K^*})$; note that the invariant token distribution η^* is uniform ($\eta^*(k) = \frac{1}{K^*+1} = \frac{1}{2\lceil \alpha \rceil + 1}$ for all $k = 0, 1, \dots, K^*$). Note that Π' and Π have the same strategy component but that the token supply for Π' is larger than for Π , and that Π' and Π^* have the same token supply but that the strategy component of Π' has a higher threshold.

We assert that $\eta'(0) \geq \frac{1}{2\lceil \alpha \rceil + 1}$. If not then $\eta'(0) < \frac{1}{2\lceil \alpha \rceil + 1} = \frac{1}{K^*+1}$. It follows that for all $k \in \{0, 1, \dots, K\}$ we have $\eta'(k) < \frac{1}{K^*+1} = \eta^*(k)$. Hence

$$\begin{aligned} \lceil \alpha \rceil &= \sum_{k=0}^{K^*} k\eta^*(k) = \sum_{k=0}^{K^*} k(\eta^*(k) - \eta'(k)) + \sum_{k=0}^{K^*} k\eta'(k) \\ &\leq K^* \sum_{k=0}^{K^*} (\eta^*(k) - \eta'(k)) + \sum_{k=0}^{K^*} k\eta'(k) = K^* \left(1 - \sum_{k=0}^{K^*} \eta'(k)\right) + \sum_{k=0}^{K^*} k\eta'(k) \\ &= K^* \sum_{k=K^*}^K \eta'(k) + \sum_{k=0}^{K^*} k\eta'(k) \leq \sum_{k=K^*}^K k\eta'(k) + \sum_{k=0}^{K^*} k\eta'(k) = \lceil \alpha \rceil \end{aligned}$$

This is a contradiction. Hence, $\eta'(0) \geq \frac{1}{2\lceil \alpha \rceil + 1}$.

Because the token supply for Π is less than Π' , the number of agents with no tokens is larger, so $\eta(0) > \eta'(0) \geq \frac{1}{2\lceil \alpha \rceil + 1}$. Hence,

$$\text{Eff}(\Pi) = (1 - \eta(0))(1 - \eta(K)) < (1 - \eta(0)) < \left(1 - \frac{1}{2\lceil \alpha \rceil + 1}\right)$$

which is the desired result in the second case. This complete the proof for part (i).

Proof of Theorem 9 Both assertions follow immediately by combining Theorems 7 and 8. □

Proof of Theorem 10 We first derive the lower bound K^L . If $\Pi_K = (K/2, \sigma_K)$ is an equilibrium protocol then consecutive marginal utilities bear the relationship

$$\phi_l M_{\sigma_K}(k-1) + \phi_c M_{\sigma_K}(k) = -\phi_r M_{\sigma_K}(k+1) > 0$$

(Because β is fixed, we suppress it in the notation.) Therefore, $M_{\sigma_K}(k) > \frac{-\phi_l}{\phi_c} M_{\sigma_K}(k-1)$. By induction,

$$M_{\sigma_K}(k) > \left(\frac{-\phi_l}{\phi_c}\right)^k M_{\sigma_K}(0) > \left(\frac{\rho\beta}{2(1-\beta) + 2\rho\beta}\right)^k M_{\sigma_K}(0)$$

Because $\phi_c M_{\sigma_K}(0) = (1-v)\rho b - \phi_r M_{\sigma_K}(1) > (1-v)\rho b + (1-v)\rho c$, we have

$$M_{\sigma_K}(0) > \frac{(1-v)\rho b}{\phi_c} = \frac{\rho\beta}{2(1-\beta) + 2\rho\beta} \frac{b+c}{\beta}$$

Therefore,

$$M_{\sigma_K}(k) > \left(\frac{\rho\beta}{2(1-\beta) + 2\rho\beta}\right)^{k+1} \frac{b+c}{\beta} \tag{39}$$

Because Π_K is assumed to be an equilibrium protocol, we must have $M_{\sigma_K}(K) \leq c/\beta$. Moreover, we must also have

$$\left(\frac{\rho\beta}{2(1-\beta) + 2\rho\beta}\right)^{K+1} \frac{b+c}{\beta} \leq \frac{c}{\beta}$$

because otherwise $M_{\sigma_K}(K) > c/\beta$. Therefore,

$$K \geq \max \left\{ \log \frac{\rho\beta}{2(1-\beta) + 2\rho\beta} \frac{c}{b+c} - 1, 0 \right\} \tag{40}$$

This provides the lower bound K^L .

We now derive the upper bound K^H . Rewriting the relation between consecutive marginal utilities, we obtain

$$\begin{aligned} 0 &= \phi_l M_{\sigma_K}(k-1) + \phi_c M_{\sigma_K}(k) + \phi_r M_{\sigma_K}(k+1) \\ &> \phi_l M_{\sigma_K}(k-1) + (\phi_c + \phi_r) M_{\sigma_K}(k) \end{aligned}$$

Therefore, $M_{\sigma_K}(k) < \frac{-\phi_l}{\phi_c + \phi_r} M_{\sigma_K}(k-1)$. By induction,

$$M_{\sigma_K}(k) < \left(\frac{-\phi_l}{\phi_c + \phi_r}\right)^k M_{\sigma_K}(0) < \left(\frac{\rho\beta}{1-\beta + \rho\beta}\right)^k M_{\sigma_K}(0)$$

Because $\phi_c M_{\sigma_K}(0) = (1-v)\rho b - \phi_r M_{\sigma_K}(1) < (1-v)\rho b - \phi_r b/\beta = 2(1-v)\rho b$, we have,

$$M_{\sigma_K}(0) < \frac{\rho\beta}{1-\beta + \rho\beta} \frac{2b}{\beta}$$

Therefore,

$$M_{\sigma_K}(k) < \left(\frac{\rho\beta}{1-\beta+\rho\beta} \right)^{k+1} \frac{2b}{\beta} \quad (41)$$

Because Π_K is assumed to be an equilibrium protocol, we must have $M_{\sigma_K}(K-1) \geq c/\beta$. Moreover,

$$\left(\frac{\rho\beta}{1-\beta+\rho\beta} \right)^K \frac{2b}{\beta} \geq \frac{c}{\beta}$$

because otherwise $M_{\sigma_K}(K-1) < c/\beta$. Therefore,

$$K \leq \log_{\frac{\rho\beta}{1-\beta+\rho\beta}} \frac{c}{2b} \quad (42)$$

This provides the upper bound K^H .

Combining the two estimates yields the range containing all integers K for which Π_K is an equilibrium protocol. The estimate for efficiency follows immediately since $Eff(\Pi_K) \geq Eff(\Pi_{K^L})$ if $K \geq K^L$, so the proof is complete. \square

References

- Adar, E., Huberman, B.: Free riding on Gnutella. *First Monday* **5**(10), 2 (2000)
- Alós-Ferrer, C.: Dynamical systems with a continuum of randomly matched agents. *J. Econ. Theory* **86**(2), 245–267 (1999)
- Berentsen, A.: Money inventories in search equilibrium. *J. Money Credit Bank.* **32**(2), 168–178 (2000)
- Berentsen, A.: On the distribution of money holdings in a random matching model. *Int. Econ. Rev.* **43**(3), 945–954 (2002)
- Berentsen, A., Camera, G., Waller, C.: Money, credit and banking. *J. Econ. Theory* **135**(1), 171–195 (2007)
- Buttayan, L., Hubaux, J.-P.: Nuglets: a virtual currency to stimulate cooperation in self-organized mobile ad hoc networks. EPFL Technical Report (2001)
- Buttayan, L., Hubaux, J.-P.: Stimulating cooperation in self-organizing mobile ad hoc Networks. *Mobile Netw. Appl.* **8**(5), 579–592 (2003)
- Camera, G., Corbae, D.: Money and price dispersion. *Int. Econ. Rev.* **40**(4), 985–1008 (1999)
- Cavalcanti, R.O., Wallace, N.: Inside and outside money as alternative media of exchange. *J. Money Credit Bank.* **31**(3), 443–457 (1999)
- Cavalcanti, R.O., Wallace, N.: A model of private bank-note issue. *Rev. Econ. Dyn.* **2**(1), 104–136 (1999)
- Chandrakumar, S., Siner, E.G., Vishnumurthy, V.: KARMA: a secure economic framework for peer-to-peer resource sharing. Workshop on the Economics of Peer-to-Peer Systems (2003)
- Ciuffoletti, A.: Secure token passing at application level. *Futur. Gener. Comput. Syst.* **26**(7), 1026–1031 (2010)
- Duffie, D., Sun, Y.: Existence of independent random matching. *Ann. Appl. Probab.* **17**(1), 386–419 (2007)
- Ellison, G.: Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Rev. Econ. Stud.* **67**(1), 17–45 (2000)
- Figueiredo, D.R., Shapiro, J.K., Towsley, D.: Payment-based incentives for anonymous peer-to-peer systems. UMass CMPSCI Technical, Report 04–62 July (2004)
- Friedman, E., Halpern, J., Kash, I.: Efficiency and nash equilibria in a scrip system for P2P networks. In: Proceedings of the Seventh ACM Conference on Electronic Commerce, pp. 140–149 (2006)
- Friedman, E., Halpern, J., Kash, I.: Optimizing scrip systems: efficiency, crashes, hoarders, and altruists. In: Proceedings of the Eighth ACM Conference on Electronic Commerce, pp. 305–315 (2007)

- Hu, T., Kennan, J., Wallace, N.: Coalition-proof trade and the Friedman rule in the Lagos-Wright Model. *J. Polit. Econ.* **117**(1), 116–137 (2002)
- Jarvis, S.A., Tan, G.: A payment-based incentive and service differentiation mechanism for peer-to-peer streaming broadcast. In: 14th IEEE International Workshop on Quality of Service (2006)
- Kandori, M.: Social norms and community enforcement. *Rev. Econ. Stud.* **59**(1), 63–80 (1992)
- Kiyotaki, N., Wright, R.: On money as a medium of exchange. *J. Political Econ.* **97**(4), 927–954 (1989)
- Kocherlakota, N.R.: The two money theorem. *Int. Econ. Rev.* **43**(2), 333–346 (2002)
- Mohr, A.E., Pai, V.: Improving robustness of peer-to-peer streaming with incentives. Workshop on the Economics of Networks, Systems and Computation (2006)
- Park, J., van der Schaar, M., Zhang, Y.: Rating protocols for online communities. *ACM Transactions on Economics and Computation* (to appear) (2010). <http://arxiv.org/abs/1101.0272>
- Podczeck, K.: On existence of rich Fubini extensions. *Econ. Theory* **45**(1–2), 1–22 (2010)
- Podczeck, K., Puzzello, D.: Independent random matching. *Econ. Theory* **50**(1), 1–29 (2012)
- Zhou, R.: Individual and aggregate real balances in a random-matching model. *Int. Econ. Rev.* **40**(4), 1009–1038 (1999)