

An Equilibrium Analysis of Scrip Systems

IAN A. KASH, Microsoft Research
ERIC J. FRIEDMAN, University of California, Berkeley
JOSEPH Y. HALPERN, Cornell University

A game-theoretic model of scrip (artificial currency) systems is analyzed. It is shown that relative entropy can be used to characterize the distribution of agent wealth when all agents use *threshold strategies*—that is, they volunteer to do work iff they have below a threshold amount of money. Monotonicity of agents' best-reply functions is used to show that scrip systems have pure strategy equilibria where all agents use threshold strategies. An algorithm is given that can compute such an equilibrium and the resulting distribution of wealth.

Categories and Subject Descriptors: C.2.4 [Computer-Communication Networks]: Distributed Systems; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems; J.4 [Social and Behavioral Sciences]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Artificial Currency, Game Theory, P2P Networks, Scrip Systems

ACM Reference Format:

ACM Trans. Econ. Comp. V, N, Article A (January YYYY), 33 pages.
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Historically, non-governmental organizations have issued their own currencies for a wide variety of purposes. These currencies, known as *scrip*, have been used in company towns where government issued currency was scarce [Timberlake 1987], in Washington DC to reduce the robbery rate of bus drivers [Washington Metropolitan Area Transit Commission 1970], and in Ithaca (New York) to promote fairer pay and improve the local economy [Ithaca Hours Inc. 2005]. Scrip systems have also been proposed for a variety of online systems.

To give some examples, market-based solutions using scrip systems have been suggested for applications such as system-resource allocation [Miller and Drexler 1988], managing replication and query optimization in a distributed database [Stonebraker et al. 1996], and allocating experimental time on a wireless sensor network test bed [Chun et al. 2005]; a number of sophisticated scrip systems have been proposed [Gupta et al. 2003; Ioannidis et al. 2002; Vishnumurthy et al. 2003] to allow agents to pool resources while avoiding what is known as *free riding*, where agents take advantage

Preliminary versions of the material in this paper appeared in the Proceedings of the 7th and 8th ACM Conferences on Electronic Commerce [Friedman et al. 2006; Kash et al. 2007] and the Proceedings of the First Conference on Auctions, Market Mechanisms and Their Applications [Kash et al. 2009]. Authors' addresses: I. A. Kash, Microsoft Research, 21 Station Road, Cambridge, UK CB1 2FB; E. J. Friedman, International Computer Science Institute and Department of Computer Science, University of California, Berkeley; and J. Y. Halpern, Computer Science Department, Cornell University, Gates Hall, Ithaca, NY 14853-7501.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1946-6227/YYYY/01-ARTA \$15.00
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

of the resources the system provides while providing none of their own (as Adar and Huberman [2000] have shown, this behavior certainly takes place in systems such as Gnutella); and Yootles [Reeves et al. 2007] uses a scrip system as a way of helping groups make decisions using economic mechanisms without involving real money.

In this paper, we provide a formal model in which to analyze scrip systems. We describe a simple scrip system and show that, under reasonable assumptions, for each fixed amount of money there is a nontrivial equilibrium involving *threshold strategies*,¹ where an agent accepts a request if he has less than $\$k$ for some threshold k . Although we refer to our unit of scrip as the dollar, these are not real dollars, nor do we view “scrip dollars” as convertible to real dollars. This is a crucial part of our model, as we assume that scrip is only valued for its use in the system.

An interesting aspect of our analysis is that, in equilibrium, the distribution of users with each amount of money is the distribution that minimizes relative entropy to an appropriate distribution (subject to the money supply constraint). This allows us to use techniques from statistical mechanics to explicitly compute the distribution of money and thus agents’ best-reply functions. Using this analysis and results of Tarski [1955], Topkis [1979], and Milgrom and Roberts [1990], we can show that there are pure-strategy equilibria using threshold strategies, and that these can be found using a simple algorithm.

In a companion paper [Kash et al. 2012], we use our analysis of this model to answer questions of interest to system designers. For example, we examine how the quantity of money effects the efficiency of the equilibrium and show that it is maximized by maintaining the appropriate ratio between the total amount of money and the number of agents. This ratio can be found by increasing the money supply up to the point that the system experiences a “monetary crash,” where money is sufficiently devalued that no agent is willing to perform a service. We also incorporate agents altruistically providing service, hoarding money, creating multiple identities, and colluding into the model.

The rest of the paper is organized as follows. In Section 2, we review related work. Then in Section 3, we present the formal model. We analyze the distribution of money in this model when agents are using threshold strategies in Section 4, and show that it is characterized by relative entropy. Using this analysis, we show in Section 5 that, under minimal assumptions, there is a nontrivial equilibrium where all agents use threshold strategies. These results apply to a sufficiently large population of agents after a sufficiently long period of time, so in Section 6 we use simulations to demonstrate that these values are reasonable in practice. We conclude in Section 7.

2. RELATED WORK

Scrip systems have a long history in computer science, with two main thrusts: resource allocation and free-riding prevention. Early applications for resource allocation include allocating time on shared computers (e.g. the PDP-1 at Harvard University in the 1960s [Sutherland 1968]), agoric systems [Miller and Drexler 1988], which envisioned solving problems such as processor scheduling using markets, and Mariposa [Stonebraker et al. 1996], a market-driven query optimizer for distributed databases. More recently, scrip systems have been used to allocate the resources of research testbeds. Examples include Mirage [Chun et al. 2005] for wireless sensor networks, Bellagio [AuYoung et al. 2007] for PlanetLab, and Egg [Brunelle et al. 2006] for grid computing. Virtual markets have been used to coordinate the activity of nodes of a sensor network [Mainland et al. 2004]. Yootles [Reeves et al. 2007] uses a scrip to

¹We do not restrict agents to playing threshold strategies, but our results show it is near optimal for them to do so when other agents do so as well.

help people make everyday decisions, such as where to have lunch, without involving real money.

Systems that use scrip to prevent free riding include KARMA [Vishnumurthy et al. 2003], which provides a general framework for P2P networks. Gupta et al. [2003] propose what they call a “debit-credit reputation computation” for P2P networks, which is essentially a scrip system. Fileteller [Ioannidis et al. 2002] uses payments in a network file storage system. Dandelion [Sirivianos et al. 2007] uses scrip in a content distribution setting. Belenkiy et al. [2007] consider how a BitTorrent-like system can make use of e-cash. Antfarm [Peterson and Siler 2009] uses scrip to optimize content distribution across a number of BitTorrent-like swarms.

Despite this tremendous interest in scrip systems, there has been relatively little work studying how they behave. While there has been extensive work in macroeconomics on the effect of variables such as the amount of money in circulation on the economy (see, for example, the relevant chapters of [Blanchard and Fischer 1989]), this work focuses on goals such as minimizing inflation and maximizing employment that are not directly relevant to a system designer.

There is also a literature in economics that attempts to understand the properties of government-backed currencies that are not tied to a commodity such as gold (*fiat money*). An early model by Kiyotaki and Wright [1989] has some similarity with our approach in terms of analysis technique, but focuses on money as useful because it can be held for free, while there are costs associated with storing a physical good until it can be traded for a different, desired good. A large literature has built on this model. Green and Zhou [1998] relax Kiyotaki and Wright’s assumption that agents can only hold a single unit of money. In subsequent work, Green and Zhou [2002] eliminate storage costs in favor of perishable goods and introduce an analysis of the dynamics of the system. Perhaps most closely related, Berentsen [2002] characterizes the equilibrium distribution of money in a model with perishable goods. Although the model is somewhat different, this characterization corresponds to a simple special case of our characterization. Berentsen and Rocheteau [2002] examine the extent to which inefficiency in this type of model depends on whether money is indivisible.

An important technical difference in these models from our work is that there is assumed to be a continuum of agents. This means that the evolution of the system is deterministic, and the distribution of money does not change in steady state. In contrast, our approach assumes a large but finite population of agents, and explicitly models the extent to which fluctuations in the distribution of money occur. Furthermore, the analysis in this line of work relies on a symmetry in agent preferences, while our model allows different types of agents with distinct preferences. Two other solution concepts have recently been investigated that look at the behavior of systems with an infinite number of agents: *mean field equilibrium* [Adlakha et al. 2010; Adlakha and Johari 2010] and *oblivious equilibrium* [Weintraub et al. 2008; 2010; Weintraub et al. 2011]. These solution concepts are similar in spirit to ours, in that they find equilibria in restricted sets of strategies, and then show that these converge to approximations of the true equilibria. Our results are somewhat stronger; we show that the “restricted equilibrium” is an approximate equilibrium, rather than just converging to one. In fact, we conjecture that our threshold equilibrium is precisely the mean field equilibrium for our model.

Ng et al. [2005] studied user behavior in a deployed scrip system, and observed that users behaved rationally in the presence of a non-strategyproof combinatorial auction. Among other manipulations, they observed users breaking of a single bid into multiple bids to exploit the greedy nature of the clearing algorithm. Their observations suggest that system designers will have to deal with game-theoretic concerns.

Hens et al. [2007] do a theoretical analysis of what can be viewed as a scrip system in a related model. There are a number of significant differences between the models. First, in the Hens et al. model, there is essentially only one type of agent, but an agent's utility for getting service (our γ_t) may change over time. Thus, at any time, there will be agents who differ in their utility. In our model, agents want service only occasionally, so, at each round, we assume that one agent is chosen (by nature) to request a service, while other agents decide whether or not to provide it. (As the number of agents increases, the time between rounds decreases, so as to keep this assumption reasonable.) In the Hens et al. model, every agent always desires service, so, at each round, each agent decides whether to provide service, request service, or opt out, as a function of his utilities and the amount of money he has. They assume that there is no (direct) cost for providing service and everyone is able to do so. However, they do assume that agents cannot simultaneously provide and receive service, an assumption that is typically unreasonable in a peer-to-peer system. Rather than having desire for service arise from explicitly modeled fluctuations in utility, we model it as being exogenously generated for one agent in each round. The decrease in time between rounds captures the ability of agents to provide service at essentially the same time they receive it. Under this assumption, a system with optimal performance is one where half the agents request service and the other half are willing to provide it. Despite these differences, Hens et al. also show that agents will use a threshold strategy.

Aperjis and Johari [2006] examine a model of a P2P filesharing system as an exchange economy. They associate a price (in bandwidth) with each file and find a market equilibrium in the resulting economy. Later work by Aperjis et al. does use a currency, but the focus remains on establishing market prices [2008]. Subsequent to our work, Dandekar et al. [2011] examined networks where each agent can issue his own scrip, Rahman et al. [2010] proposed a protocol to automatically adjust credit policies to maintain stability, Humbert et al. [2011] modeled scrip systems where multiple agents combine to provide service, van der Schaar et al. [2013] studied how to set parameters in a way that is robust to noise, and Johnson et al. [2014] showed that welfare can be improved if the system prefers volunteers with low amounts of scrip.

The ultimate goal of a scrip system is to promote cooperation. While there is limited theoretical work on scrip systems, there is a large body of work on cooperation. Much of the work involves a large group of agents being randomly matched to play a game such as prisoner's dilemma. Such models were studied in the economics literature [Kandori 1992; Ellison 1994] and first applied to online reputations in [Friedman and Resnick 2001]; Feldman et al. [2004] apply them to P2P systems.

These models fail to capture important asymmetries in the interactions of the agents. When a request is made, there are typically many people in the network who can potentially satisfy it (especially in a large P2P network), but not all can. For example, some people may not have the time or resources to satisfy the request. The random-matching process ignores the fact that some people may not be able to satisfy the request. (Presumably, if the person matched with the requester could not satisfy the match, he would have to defect.) Moreover, it does not capture the fact that the decision as to whether to "volunteer" to satisfy the request should be made before the matching process, not after. That is, the matching process does not capture the fact that if someone is unwilling to satisfy the request, there may well be others who can satisfy it. Finally, the actions and payoffs in the prisoner's dilemma game do not obviously correspond to actual choices that can be made. For example, it is not clear what defection on the part of the requester means. Our model addresses all these issues.

Scrip systems are not the only approach to preventing free riding. Two other approaches often used in P2P networks are barter and reputation systems. The essence of barter for our purposes is that when deciding whether to provide service to others

agents only consider the value created by *their own* current or past interactions with the agent in question. Perhaps the best-known example of a system that uses barter is BitTorrent [Cohen 2003], where clients downloading a file try to find other clients with parts they are missing so that they can trade, thus creating a roughly equal amount of work. Since the barter is restricted to users currently interested in a single file, this works well for popular files, but tends to have problems maintaining availability of less popular ones. An example of a barter-like system built on top of a more traditional file-sharing system is the credit system used by eMule [eMule Project 2009]. Each user tracks his history of interactions with other users and gives priority to those he has downloaded from in the past. However, in a large system, the probability that a pair of randomly-chosen users will have interacted before is quite small, so this interaction history will not be terribly helpful. Anagnostakis and Greenwald [2004] present a more sophisticated version of this approach, but it still seems to suffer from similar problems. More recently, Piatek et al. [2008] have proposed a model based on including intermediaries a single hop away; this model is more liquid than barter but not as liquid as a full scrip system.

A number of attempts have been made at providing general reputation systems (e.g. [Guha et al. 2004; Gupta et al. 2003; Kamvar et al. 2003; Xiong and Liu 2002]). The basic idea is to aggregate each user's experience into a global number for each individual that intuitively represents the system's view of that individual's reputation. However, these attempts tend to suffer from practical problems because they implicitly view users as either "good" or "bad", assume that the "good" users will act according to the specified protocol, and that there are relatively few "bad" users. Unfortunately, if there are easy ways to game the system, once this information becomes widely available, rational users are likely to make use of it. We cannot count on only a few users being "bad" (in the sense of not following the prescribed protocol). For example, Kazaa uses a measure of the ratio of the number of uploads to the number of downloads to identify good and bad users. However, to avoid penalizing new users, they gave new users an average rating. Users discovered that they could use this relatively good rating to free ride for a while and, once it started to get bad, they could delete their stored information and effectively come back as a "new" user, thus circumventing the system (see [Anagnostakis and Greenwald 2004] for a discussion and [Friedman and Resnick 2001] for a formal analysis of this "whitewashing"). Thus, Kazaa's reputation system is ineffective.

3. THE MODEL

Before specifying our model formally, we give an intuitive description of what our model captures. We model a scrip system where, as in a P2P filesharing system, agents provide each other with service. There is a single service (such as file uploading) that agents occasionally want. In practice, at any given time, a number of agents will want service but, to simplify the formal description and analysis, we model the scrip system as proceeding in a series of rounds where, in each round, a single agent wants service (the time between rounds will be adjusted to capture the growth in parallelism as the number of agents grows).² In each round, after an agent requests service, other agents have to decide whether they want to volunteer to provide service. However, not all agents may be able to satisfy the request (not everyone has every file). While, in practice, the ability of agents to provide service at various times may be correlated for a number of reasons (if I don't have the file today I probably still don't have it to-

²For large numbers of agents, our model converges to one in which agents make requests in real time, and the time between an agent's requests is exponentially distributed. In addition, the time between requests served by a single player is also exponentially distributed.

morrow; if one agent does not have a file, it may be because it is rare, so that should increase the probability that other agents do not have it), for simplicity, we assume that the events of an agent being able to provide service in different rounds or two agents being able to provide service in the same or different rounds are independent. If there is at least one volunteer, someone is chosen from among the volunteers (at random) to satisfy the request. Our model allows some agents to be more likely to be chosen (perhaps they have more bandwidth) but does not capture rules that take the amount of scrip volunteers have into account (such rules are considered by Johnson et al. [2014]) or allow an agent to specify which agent is chosen. Allowing agents this type of control would break the symmetries we use to characterize the long-run behavior of the system, and would create new opportunities for strategic behavior. The requester then gains some utility (he got the file) and the volunteer loses some utility (he had to use his bandwidth to upload the file), and the requester pays the volunteer a fee that we fix at one dollar. As is standard in the literature, we assume that agents discount future payoffs. This captures the intuition that a reward today is worth more than a reward tomorrow, and allows us to compute the total utility derived by an agent in an infinite game. The amount of utility gained by having a service performed and the amount lost by performing it, as well as many other parameters may depend on the agent.

More formally, we assume that agents have a *type* t drawn from some finite set T of types. We can describe the entire population of agents using the pair (T, \vec{f}) , where \vec{f} is a vector of length $|T|$ and f_t is the fraction with type t . In this paper, we consider only what we call *standard agents*. These are agents who derive no pleasure from performing a service, and for whom money is valued only for its use in obtaining service. Thus, for a standard agent, there is no direct connection between money (measured in dollars) and utility (measured in utils). We can characterize the type of a standard agent by a tuple $t = (\alpha_t, \beta_t, \gamma_t, \delta_t, \rho_t, \chi_t)$, where

- $\alpha_t > 0$ is the undiscounted cost in utils for an agent of type t to satisfy a request;
- $0 < \beta_t < 1$ is the probability that an agent of type t can satisfy a request;
- $\gamma_t > \alpha_t$ is the utility that an agent of type t gains for having a request satisfied;
- $0 < \delta_t < 1$ is the rate at which an agent of type t discounts utility;
- $\rho_t > 0$ represents the (relative) request rate (some people want files more often than others). For example, if there are two types of agents with $\rho_{t_1} = 2$ and $\rho_{t_2} = 1$ then agents of the first type will make requests twice as often as agents of the second type. Since these request rates are relative, we can multiply them all by a constant to normalize them. To simplify later notation, we assume the ρ_t are normalized so that $\sum_{t \in T} \rho_t f_t = 1$;
- $\chi_t > 0$ represents the (relative) likelihood of an agent to be chosen when he volunteers (some uploaders may be more popular than others). In particular, this means the relative probability of two given agents being chosen is independent of which other agents volunteer; and
- $\omega_t = \beta_t \chi_t / \rho_t$ is not part of the tuple, but is an important derived parameter that, as we will see in Section 4, helps determine how much money an agent will have.

We occasionally omit the subscript t on some of these parameters when it is clear from context or irrelevant.

Representing the population of agents in a system as (T, \vec{f}) captures the essential features of a scrip system we want to model: there are a large number of agents who may have different types. Note that some tuples (T, \vec{f}) may be incompatible with there being some number N of agents. For example, if there are two types, and \vec{f} says that half of the agents are of each type, then there cannot be 101 agents. Similar issues

arise when we want to talk about the amount of money in a system. We could deal with this problem in a number of ways (for example, by having each agent determine his type at random according to the distribution \vec{f}). For convenience, we simply do not consider population sizes that are incompatible with \vec{f} . This is the approach used in the literature on *N-replica economies* [Mas-Colell et al. 1995].

Formally, we consider games specified by a tuple (T, \vec{f}, h, m, n) , where T and \vec{f} are as defined above, $h \in \mathbb{N}$ is the *base* number of agents of each type, $n \in \mathbb{N}$ is number of *replicas* of these agents and $m \in \mathbb{R}^+$ is the average amount of money. The total number of agents is thus hn . We ensure that the fraction of agents of type t is exactly f_t and that the average amount of money is exactly m by requiring that $f_t h \in \mathbb{N}$ and $mh \in \mathbb{N}$. Having created a base population satisfying these constraints, we can make an arbitrary number of copies of it. More precisely, we assume that agents $0 \dots f_{t_1}h - 1$ have type t_1 , agents $f_{t_1}h \dots (f_{t_1} + f_{t_2})h - 1$ have type t_2 , and so on through agent $h - 1$. These base agents determine the types of all other agents. Each agent $j \in \{h, \dots, hn - 1\}$ has the same type as $j \bmod h$; that is, all the agents of the form $j + kh$ for $k = 1, \dots, n - 1$ are replicas of agent j .

We also need to specify how money is initially allocated to agents. Our results are based on the long-run behavior of the system and so they turn out to hold for any initial allocation of money. For simplicity, at the start of the game we allocate each of the hmn dollars in the system to an agent chosen uniformly at random, but all our results would hold if we chose any other initial distribution of money.

To make precise our earlier informal description, we describe (T, \vec{f}, h, m, n) as an infinite extensive-form game. A non-root node in the game tree is associated with a round number (how many requests have been made so far), a phase number, either 1, 2, 3, or 4 (which describes how far along we are in determining the results of the current request), a vector \vec{x} where x_i is the current amount of money agent i has, and $\sum_i x_i = mhn$, and, for some nodes, some additional information whose role will be made clear below. We use $\tau(i)$ to denote the type of agent i .

- The game starts at a special root node, denoted Λ , where nature moves. Intuitively, at Λ , nature allocates money uniformly at random, so it transitions to a node of the form $(0, 1, \vec{x})$: round zero, phase one, and allocation of money \vec{x} , and each possible transition is equally likely.
- At a node of the form $(r, 1, \vec{x})$, nature selects an agent to make a request in the current round. Agent i is chosen with probability $\rho_{\tau(i)}/hn$. If i is chosen, a transition is made to $(r, 2, \vec{x}, i)$.
- At a node of the form $(r, 2, \vec{x}, i)$, nature selects the set V of agents (not including i) able to satisfy the request. Each agent $j \neq i$ is included in V with probability $\beta_{\tau(j)}$. If V is chosen, a transition is made to $(r, 3, \vec{x}, i, V)$.
- At a node of the form $(r, 3, \vec{x}, i, V)$, each agent in V chooses whether to volunteer. If V' is the set of agents who choose to volunteer, a transition is made to $(r, 4, \vec{x}, i, V')$.
- At a node of the form $(r, 4, \vec{x}, i, V')$, if $V' \neq \emptyset$, nature chooses a single agent in V' to satisfy the request. Each agent j is chosen with probability $\chi_{\tau(j)} / \sum_{j' \in V'} \chi_{\tau(j')}$. If j is chosen, a transition is made to $(r + 1, 1, \vec{x}')$, where

$$x'_j = \begin{cases} x_j - 1 & \text{if } i = j \text{ and } x_j > 0, \\ x_j + 1 & \text{if } j \text{ is chosen by nature and } x_i > 0, \\ x_j & \text{otherwise.} \end{cases}$$

If $V' = \emptyset$, nature has no choice; a transition is made to $(r + 1, 1, \vec{x})$ with probability 1.

A strategy for agent ℓ describes whether or not he will volunteer at every node of the form $(r, 3, \vec{x}, i, V)$ such that $\ell \in V$. (These are the only nodes where ℓ can move.)

We also need to specify what agents know when they make their decisions. To make our results as strong as possible, we allow an agent to base his strategy on the entire history of the game, which includes, for example, the current wealth of every other agent. As we show, even with this unrealistic amount of information, available, it would still be approximately optimal to adopt a simple strategy that requires little information—specifically, agents need to know only their current wealth. That means that our results would continue to hold as long as agents knew at least this information. A strategy profile \vec{S} consists of one strategy per agent. A strategy profile \vec{S} determines a probability distribution over paths $\Pr_{\vec{S}}$ in the game tree. Each path determines the value of the following random two variables:

- x_i^r , the amount of money agent i has during round r , defined as the value of x_i at the nodes with round number r and
- u_i^r , the utility of agent i for round r . If i is a standard agent, then

$$u_i^r = \begin{cases} \gamma_{\tau(i)} & \text{if a node } (r, 4, \vec{x}, i, V') \text{ is on the path with } V' \neq 0 \\ -\alpha_{\tau(i)} & \text{if } i \text{ is chosen by nature at node } (r, 4, \vec{x}, j, V') \\ 0 & \text{otherwise.} \end{cases}$$

$U_i(\vec{S})$, the total expected utility of agent i if strategy profile \vec{S} is played, is the discounted sum of his per round utilities u_i^r , but the exact form of the discounting requires some explanation. In our model, only one agent makes a request each round. As the number of agents increases, an agent has to wait a larger number of rounds to make requests, so naively discounting utility would mean his utility decreases as the number of agents increases, even if all of his requests are satisfied. This is an artifact of our model breaking time into discrete rounds where a single agent makes a request. In reality, many agents make requests in parallel, and how often an agent desires service typically does not depend on the number of agents. It would be counterintuitive to have a model that says that if agents make requests at a fixed rate and they are all satisfied, then their expected utility depends on the number of other agents. As the following lemma shows, there is a unique discount rate that removes this dependency.³

LEMMA 3.1. *With a discount rate of $(1 - (1 - \delta_t)/n)$, an agent of type t 's expected discounted utility for having all his requests satisfied is independent of the number of replicas n . Furthermore, this is the unique such rate such that the discount rate is δ_t when $n = 1$.*

PROOF. The agent makes a request each round with probability ρ_t/hn , so his expected discounted utility for having all his requests satisfied is

$$\begin{aligned} \sum_{r=0}^{\infty} (1 - (1 - \delta_t)/n)^r (\rho_t \gamma_t / (hn)) &= (\rho_t \gamma_t / (hn)) / (1 - (1 - (1 - \delta_t)/n)) \\ &= (\rho_t \gamma_t / h) / (1 - \delta_t) \end{aligned}$$

This is independent of n , and satisfies $(1 - (1 - \delta_t)/1) = \delta_t$, as desired. It is unique because choosing any other discount rate for some n will cause the value of the sum to differ from $(\rho_t \gamma_t / h) / (1 - \delta_t)$ for that n . \square

³In preliminary versions of this work we used the discount rate of $\delta_t^{1/n}$. This rate captures the intuitive idea of making the time between rounds $1/n$, but results in an agent's utility depending on the number of other agents, even if all the agent's requests are satisfied. However, in the limit as δ_t goes to 1, agents' normalized expected utilities (multiplied by $1 - \delta_t$, as in Equation 1) are the same with either discount rate, so our main results hold with the discount rate $\delta_t^{1/n}$ as well.

As is standard in economics, for example in the folk theorem for repeated games [Fudenberg and Tirole 1991], we multiply an agent's utility by $(1 - \delta_t)$ so that his expected utility is independent of his discount rate as well. With these considerations in mind, the total expected utility of agent i given the vector of strategies \vec{S} is

$$U_i(\vec{S}) = (1 - \delta_{\tau(i)}) \sum_{r=0}^{\infty} (1 - (1 - \delta_{\tau(i)})/n)^r E_{\vec{S}}[u_i^r], \quad (1)$$

In modeling the game this way, we have implicitly made a number of assumptions. For example, we have assumed that all of agent i 's requests that are satisfied give agent i the same utility, and that prices are fixed. We discuss the implications of these assumptions in Section 7.

Our solution concept is the standard notion of an approximate Nash equilibrium. As usual, given a strategy profile \vec{S} and agent i , we use (S'_i, \vec{S}_{-i}) to denote the strategy profile that is identical to \vec{S} except that agent i uses S'_i .

Definition 3.2. A strategy S'_i for agent i is an ϵ -best reply to a strategy profile \vec{S}_{-i} for the agents other than i in the game (T, \vec{f}, h, m, n) if, for all strategies S''_i ,

$$U_i(S''_i, \vec{S}_{-i}) \leq U_i(S'_i, \vec{S}_{-i}) + \epsilon.$$

Definition 3.3. A strategy profile \vec{S} for the game (T, \vec{f}, h, m, n) is an ϵ -Nash equilibrium if for all agents i , S_i is an ϵ -best reply to \vec{S}_{-i} . A Nash equilibrium is an epsilon-Nash equilibrium with $\epsilon = 0$.

As we show in Section 5, (T, \vec{f}, h, m, n) has equilibria where agents use a particularly simple type of strategy, called a *threshold strategy*. Intuitively, an agent with “too little” money will want to work, to minimize the likelihood of running out due to making a long sequence of requests before being able to earn more money. On the other hand, a rational agent with plenty of money will think it is better to delay working, thanks to discounting. These intuitions suggest that the agent should volunteer if and only if he has less than a certain amount of money. Let s_k be the strategy where an agent volunteers if and only if the requester has at least 1 dollar and the agent has less than k dollars. Note that s_0 is the strategy where the agent never volunteers. While everyone playing s_0 is a Nash equilibrium (nobody can do better by volunteering if no one else is willing to), it is an uninteresting one.

We frequently consider the situation where each agent of type t uses the same threshold s_{k_t} . In this case, a single vector \vec{k} suffices to indicate the threshold of each type, and we can associate with this vector the strategy $\vec{S}(\vec{k})$ where $\vec{S}(\vec{k})_i = s_{k_{\tau(i)}}$ (i.e., agent i of type $\tau(i)$ uses threshold $k_{\tau(i)}$).

For the rest of this paper, we focus on threshold strategies (and show why it is reasonable to do so). In particular, we show that, if all other agents use threshold strategies, it is approximately optimal for an agent to use one as well. Furthermore there exist Nash equilibria where agents do so.⁴ While there are potentially other equilibria that use different strategies, if a system designer has agents use threshold strategies by default (e.g., through the standard behavior of the client software), no agent will have an incentive to change. Since threshold strategies have such low information requirements, they are a particularly attractive choice for a system designer as well for the agents, since they are so easy to play.

⁴These equilibria actually satisfy the stronger condition of (approximate) subgame perfection.

For threshold strategy $\vec{S}(\vec{k})$, if $mh_n \geq \sum_t f_t k_t h_n$ the system will quickly reach a state where each agent has k_t dollars, so no agent will volunteer. This is equivalent to all agents using s_0 , and similarly uninteresting. Therefore, in our analysis we assume that $mh_n < \sum_t f_t k_t h_n$.

4. ANALYZING THE DISTRIBUTION OF WEALTH

Our main goal is to show that there exists an approximate equilibrium where all agents play threshold strategies. In this section, we examine a more basic question: if all agents play a threshold strategy, what happens? We show that there is some distribution over money (i.e., a distribution that describes what fraction of people have each amount of money) such that the system “converges” to this distribution in a sense to be made precise shortly.

To motivate our interest in the distribution, consider an agent i who is trying to decide on a best response in a setting where all other agents are playing a threshold strategy, and all agents of a particular type play the same strategy. Specifically, suppose that agent i has $\$k$, and is trying to decide whether to volunteer. Assume that the system is sufficiently large that one agent’s decision does not affect the distribution. To figure out whether to volunteer, agent i must compute how likely he is to run out of money before he gets a chance to make another dollar. Clearly this depends on how much money he has. But it also depends on how likely he is to be chosen when he volunteers, which, in turn, depends on how many other volunteers there are (and thus on how many agents are not at their threshold). Our results show that there is a distribution of money d^* such that, with extremely high probability, the actual distribution is almost always extremely close to d^* . By knowing d^* , the agent will know what fraction of agents of each type t have each amount of money. If all the agents of type t use the same threshold strategy, he will also know how many agents of type t volunteer. Moreover, this number will be essentially the same at every round. This will enable him to figure out when he should volunteer.

We remark that, in addition to providing an understanding of system behavior that underpins our later results, this result also provides a strong guarantee about the stability of the economy. It shows that we do not have wild swings of behavior; in particular, the fraction of agents volunteering is essentially constant.⁵

Suppose that all agents of each type t use the same threshold k_t , so we can write the vector of thresholds as \vec{k} . For simplicity, assume that each agent has at most k_t dollars. We can make this assumption with essentially no loss of generality, since if someone has more than k_t dollars, he will just spend money until he has at most k_t dollars. After this point he will never acquire more than k_t . Thus, eventually the system will be in a state where, for all types t , no agent of type t has more than k_t dollars.

We are interested in the vectors \vec{x}^r that can be observed in round r (recall that x_i^r is the amount of money that agent i has at round r). By assumption, if agent i has type $\tau(i)$, then $x_i^r \in \{0, \dots, k_{\tau(i)}\}$. In addition, since the total amount of money is hmn ,

$$\vec{x}^r \in X_{T, \vec{f}, h, m, n, \vec{k}} = \{\vec{x} \in \mathbb{N}^{hn} \mid \forall i. x_i \leq k_{\tau(i)}, \sum_i x_i = hmn\}.$$

The evolution of \vec{x}^r can be described by a Markov chain $\mathcal{M}_{T, \vec{f}, h, m, n, \vec{k}}$ over the state space $X_{T, \vec{f}, h, m, n, \vec{k}}$. For brevity, we refer to the Markov chain and state space as \mathcal{M}

⁵Our results show that such swings occur with extremely small probability. While this guarantees they will eventually occur, the expected time this takes is so large that it will effectively never happen in practice for even a moderately sized system. Even if such an unlikely event did occur, our results guarantee that it is transient and that the system will converge back toward the steady-state distribution.

and X , respectively, when the subscripts are clear from context. It is possible to move from state s to state s' in a single round if, by choosing a particular agent i to make a request and another agent j to satisfy it, i 's amount of money in s' is 1 more than in s ; j 's amount of money in s' is 1 less than in s , and all other agents have the same amount of money in s and s' . Therefore, the probability of a transition from a state \vec{x} to \vec{y} is 0 unless there exist two agents i and j such that $\vec{y}_{i'} = \vec{x}_{i'}$ for all $i' \notin \{i, j\}$, $\vec{y}_i = \vec{x}_i + 1$, and $\vec{y}_j = \vec{x}_j - 1$. In this case, the probability of transitioning from \vec{x} to \vec{y} is the probability of j being chosen to make a request and i being chosen to satisfy it. Let $\Delta_{\vec{f}, m, \vec{k}}$ denote the set of probability distributions d on $\cup_{t \in T} \{t\} \times \prod_t \{0, \dots, k_t\}$ such that for all types t , $\sum_{l=0}^{k_t} d(t, l) = f_t$ and $\sum_t \sum_{l=0}^{k_t} ld(t, l) = m$. We can think of $d(t, l)$ as the fraction of agents that are of type t and have l dollars. We can associate each state \vec{x} with its corresponding distribution $d^{\vec{x}}$. This is a useful way of looking at the system, since we typically just care about the fraction of people with each amount of money, not the amount that each particular agent has. We show that, if n is large, then there is a distribution $d^* \in \Delta_{\vec{f}, m, \vec{k}}$ such that, after a sufficient amount of time, the Markov chain \mathcal{M} is almost always in a state \vec{x} such that $d^{\vec{x}}$ is close to d^* . Thus, agents can base their decisions about what strategy to use on the assumption that they will be in a state where the distribution of money is essentially d^* . Note that, since agents discount future utility, the transient behavior of the Markov chain does matter, but by making δ_t sufficiently large (i.e., if agents are sufficiently patient) the effect on utility can be made arbitrarily small. Similarly, for sufficiently large n , the effect on utility due to extremely rare deviations from d^* becomes arbitrarily small.

We can in fact completely characterize the distribution d^* . Given two distributions $d, q \in \Delta_{\vec{f}, m, \vec{k}}$, let

$$H(d||q) = \sum_{(t,j) \text{ s.t. } q(t,j) \neq 0} d(t,j) \log \left(\frac{d(t,j)}{q(t,j)} \right)$$

denote the *relative entropy* of d relative to q ($H(d||q) = \infty$ if $d(t,j) = 0$ and $q(t,j) \neq 0$ or vice versa); this is also known as the *Kullback-Leibler divergence of q from d* [Cover and Thomas 1991]. If Δ is a closed convex set of distributions, then it is well known that, for each q , there is a unique distribution in Δ that minimizes the relative entropy to q . Since $\Delta_{\vec{f}, m, \vec{k}}$ is easily seen to be a closed convex set of distributions, in particular, this is the case for $\Delta_{\vec{f}, m, \vec{k}}$. We now show that there exists a q such that, for n sufficiently large, the Markov chain \mathcal{M} is almost always in a state \vec{x} such that $d^{\vec{x}}$ is close to the distribution $d_{q, \vec{f}, m}^* \in \Delta_{\vec{f}, m, \vec{k}}$ that minimizes entropy relative to q . (We omit some or all of the subscripts on d^* when they are not relevant.) The statement is correct under a number of senses of “close”. For definiteness, we consider the Euclidean distance. Given $\varepsilon > 0$ and q , let $X_{T, \vec{f}, h, m, n, \vec{k}, \varepsilon, q}$ (or $X_{\varepsilon, q}$, for brevity) denote the set of states $\vec{x} \in X_{T, \vec{f}, h, m, n, \vec{k}}$ such that $\sum_{(t,j)} |d^{\vec{x}}(t, j) - d_q^*|^2 < \varepsilon$.

Let $I_{q, n, \varepsilon}^r$ be the random variable that is 1 if $d^{\vec{x}^r} \in X_{\varepsilon, q}$, and 0 otherwise.

THEOREM 4.1. *For all games $(T, \vec{f}, h, m, 1)$, all vectors \vec{k} of thresholds, and all $\varepsilon > 0$, there exist $q \in \Delta_{\vec{f}, m, \vec{k}}$ and n_ε such that, for all $n > n_\varepsilon$, there exists a round r^* such that, for all $r > r^*$, we have $\Pr(I_{q, n, \varepsilon}^r = 1) > 1 - \varepsilon$.*

The proof of Theorem 4.1 can be found in Appendix A. One interesting special case of the theorem is when there exist β , χ , and ρ such that for all types t , $\beta_t = \beta$, $\chi_t = \chi$, and $\rho_t = \rho$. In this case q is the distribution $q(t, j) = f_t / (k_t + 1)$ (i.e., q is uniform within each type t). We sketch the proof for this special case here.

PROOF. (Sketch) Using standard techniques, we can show that our Markov Chain has a *limit distribution* π such that for all \vec{y} , $\lim_{r \rightarrow \infty} \Pr(\vec{x}^r = \vec{y}) = \pi(\vec{y})$. Let $T_{\vec{x}\vec{y}}$ denote the probability of transitioning from (recurrent) state \vec{x} to (recurrent) state \vec{y} . It is easily verified by an explicit computation of the transition probabilities that (in this special case) $T_{\vec{x}\vec{y}} = T_{\vec{y}\vec{x}}$. It is well known that this symmetry implies that π is the uniform distribution [Resnick 1992]. Thus, after a sufficient amount of time, the distribution of \vec{x}^r will be arbitrarily close to uniform.

Since, for large r , $\Pr(\vec{x}^r = \vec{y})$ is approximately $1 / |X|$, the probability of \vec{x}^r being in a set of states is the size of the set divided by the total number of states. Using a straightforward combinatorial argument, it can be shown that the fraction of states not in $X_{\varepsilon, q}$ is bounded by $p(n)/e^{cn}$, where p is a polynomial. This fraction goes to 0 as n gets large. Thus, for sufficiently large n , $\Pr(I_{q, n, \varepsilon}^r = 1) > 1 - \varepsilon$. \square

The last portion of the proof sketch is actually a standard technique from statistical mechanics that involves showing that there is a *concentration phenomenon* around the maximum entropy distribution [Jaynes 1978]. To illustrate what we mean by a concentration phenomenon, consider a system with only two dollars. With n agents, there are $O(n^2)$ ways to assign the dollars to different agents and $O(n)$ ways to assign them to the same agent. If each way of assigning the two dollars to agents is equally likely, we are far more likely to see a distribution of money where two agents have one dollar each than one where a single agent has two dollars. In the special case considered in the proof sketch, when π is the uniform distribution, the number of states corresponding to a particular distribution d is proportional to $e^{nH(d)}$ (where H here is the standard entropy function). In general, each state is not equally likely, which is why the general proof in Appendix A uses relative entropy.⁶

Theorem 4.1 tells us that, after enough time, the distribution of money is almost always close to some d^* , where d^* can be characterized as a distribution that minimizes relative entropy subject to some constraints. In Appendix A, we show that the appropriate distribution q is $q(t, i) = (\omega_t)^i / (\sum_t \sum_{j=0}^{k_t} (\omega_t)^j)$. The following lemma shows how we can compute d^* from q .

LEMMA 4.2.

$$d^*(t, i) = \frac{f_t \lambda^i q(t, i)}{\sum_{j=0}^{k_t} \lambda^j q(t, j)}, \quad (2)$$

where λ is the unique value such that

$$\sum_t \sum_i i d^*(t, i) = m. \quad (3)$$

The proof of Lemma 4.2 is omitted because it can be easily checked using Lagrange multipliers in the manner of [Jaynes 1978] where the function to be minimized is the relative entropy of d^* relative to q and the constraints are that an f_t fraction of the agents are of type t and the average amount of money is m .

To give some intuition for the form of d^* , suppose that there is only a single agent who randomly receives opportunities to earn and spend money, but receives opportunities to earn ω_t times as often as opportunities to spend, is unwilling to earn more than k_t dollars, and is unable to spend when he has zero dollars. Up to a normalizing constant, his steady state probability of having i dollars is $q(t, i)$. In the full system,

⁶Note that in generalizing to relative entropy we switch from maximizing to minimizing; maximizing entropy is equivalent to minimizing relative entropy relative to the uniform distribution.

each agent is not walking independently, but instead is limited by the common amount of money. This acts as a common bias on their walks, a factor which is captured by λ .

5. EXISTENCE OF EQUILIBRIA

We have seen that the system is well behaved if the agents all follow a threshold strategy; we now want to show that, if the discount factor δ is sufficiently large for all agents, there is a nontrivial approximate Nash equilibrium where they do so (that is, an approximate Nash equilibrium where all the agents use s_k for some $k > 0$.) To understand why we need δ to be sufficiently large, note that if δ is small, then agents have no incentive to work. Intuitively, if future utility is sufficiently discounted, then all that matters is the present, and there is no point in volunteering to work. Thus, for sufficiently small δ , s_0 is the only equilibrium. To show that there is a nontrivial equilibrium if the discount factor is sufficiently large, we first show that, if every other agent is playing a threshold strategy, then there is an approximate best reply that is also a threshold strategy. Furthermore, we show that the best-reply function is monotone; that is, if some agents change their strategy to one with a higher threshold, no other agent can do better by lowering his threshold. This makes our game one with what Milgrom and Roberts [1990] call *strategic complementarities*. Using results of Tarski [1955], Topkis [1979] showed that there are pure strategy equilibria in such games, since the process of starting with a strategy profile where everyone always volunteers (i.e., the threshold is ∞) and then iteratively computing the best-reply profile to it converges to a Nash equilibrium in pure strategies. This procedure also provides an efficient algorithm for explicitly computing equilibria.

To see that threshold strategies are approximately optimal, consider a single agent i of type t and fix the vector \vec{k} of thresholds used by the other agents. If we assume that the number of agents is large, what an agent i does has essentially no effect on the behavior of the system (although it will, of course, affect that agent's payoffs). In particular, this means that the distribution q of Theorem 4.1 characterizes the distribution of money in the system with high probability. (This applies only after some period of time, but for patient agents the importance of this initial period is negligible.) This distribution, together with the vector \vec{k} of thresholds, determines what fraction of agents volunteers at each step. This, in turn, means that from the perspective of agent i , the problem of finding an optimal response to the strategies of the other agents reduces to finding an optimal policy in a Markov decision process (MDP) $\mathcal{P}_{G, \vec{s}(\vec{k}), t}$. The behavior of the MDP $\mathcal{P}_{G, \vec{s}(\vec{k}), t}$ depends on two probabilities: p_u and p_d . Informally, p_u is the probability of i earning a dollar during each round if is willing to volunteer, and p_d is the probability that i will be chosen to make a request during each round. Note that p_u depends on m , \vec{k} , and t (although it turns out that p_d depends only on n , the number of agents in the system); if the dependence of p_u on m , \vec{k} , and/or t is important, we add the relevant parameters to the superscript, writing, for example, $p_u^{m, \vec{k}}$. We show that the optimal policy for i in $\mathcal{P}_{G, \vec{s}(\vec{k}), t}$ is a threshold policy, and that this policy is an ε -optimal strategy for G . Importantly, the same policy is optimal independent of the value of n (as long as n is sufficiently large). This allows us to ignore the exact size of the system in our later analysis.

For our results, it will be important to understand how p_u , p_d , and t affect the optimal policy for $\mathcal{P}_{G, \vec{s}(\vec{k}), t}$, and thus the ε -optimal strategies in the game. We use this understanding in this section to show that there exist nontrivial equilibria in Lemma 5.4 and for a number of results in our companion paper [Kash et al. 2012].

In the following lemma, whose proof (and the relevant formal definitions) are deferred to Appendix B, Equation (4), quantifies the effects of these parameters. When

choosing whether he should volunteer with his current amount of money, an agent faces a choice of whether to pay a utility cost of α_t now in exchange for a discounted payoff of γ_t when he eventually spends the resulting dollar. His choice will depend on how much time he expects to pass before he spends that dollar (captured by the random variable J in equation 4), which in turn depends on his current amount of money k and the probabilities p_u and p_d . The following lemma quantifies this calculation.

LEMMA 5.1. *Consider the games $G_n = (T, \vec{f}, h, m, n)$ (where T , \vec{f} , h , and m are fixed, but n may vary). There exists a k such that for all n , s_k is an optimal policy for $\mathcal{P}_{G_n, \vec{S}(\vec{k}), t}$. The threshold k is the maximum value of κ such that*

$$\alpha_t \leq E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u, p_d)}] \gamma_t, \quad (4)$$

where $J(\kappa, p_u, p_d)$ is a random variable whose value is the first round in which an agent starting with κ dollars, using strategy s_κ , and with probabilities p_u and p_d of earning a dollar and of being chosen given that he volunteers, respectively, runs out of money.

Note that the uniqueness of the optimal value of k holds generically. In particular, it holds unless (4) is satisfied with equality.

The following theorem shows that an optimal threshold policy for $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ is an ε -optimal strategy for G . In particular, this means that Equation (4) allows us to understand how changing parameters affect an ε -optimal strategy for G , not just for $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$.

THEOREM 5.2. *For all games $G = (T, \vec{f}, h, m, n)$, all vectors \vec{k} of thresholds, and all $\varepsilon > 0$, there exist n_ε^* and $\delta_{\varepsilon, n}^*$ such that for all $n > n_\varepsilon^*$, types $t \in T$, and $\delta_t > \delta_{\varepsilon, n}^*$, an optimal threshold policy for $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ is an ε -best reply to the strategy profile $\vec{S}(\vec{k})_{-i}$ for every agent i of type t .*

We defer the proof of Theorem 5.2 to Appendix B. While, in this and later theorems, the acceptable values of $\delta_{\varepsilon, n}^*$ depend on n , they are independent if, as we suggest in Section 6, the Markov Chain from Section 4 is rapidly mixing.

Given a game $G = (T, \vec{f}, h, m, n)$ and a vector \vec{k} of thresholds, Lemma 5.1 gives an optimal threshold k'_t for each type t . Theorem 5.2 guarantees that $s_{k'_t}$ is an ε -best reply to $\vec{S}_{-i}(\vec{k})$, but does not rule out the possibility of other best replies. However, for ease of exposition, we will call k'_t the best reply to \vec{S}_{-i} and call $BR_G(\vec{k}) = \vec{k}'$ the best-reply function. The following lemma shows that this function is monotone (non-decreasing). Along the way, we prove that several other quantities are monotone. First, we show that $\lambda_{m, \vec{k}}$, the value of λ from Lemma 4.2 given m and \vec{k} , is non-decreasing in m and non-increasing in \vec{k} . We use this to show that $p_u^{m, \vec{k}}$ is non-increasing in \vec{k} , which is needed to show the monotonicity of BR_G . We defer the proof to Appendix B.

LEMMA 5.3. *Consider the family of games $G_m = (T, \vec{f}, h, m, n)$ and the strategies $\vec{S}(\vec{k})$, for $mhn < \sum_t f_t k_t hn$. For this family of games, $\lambda_{m, \vec{k}}$ is non-decreasing in m and non-increasing in \vec{k} ; $p_u^{m, \vec{k}}$ is non-decreasing in m and non-increasing in \vec{k} ; and the function BR_G is non-decreasing in \vec{k} and non-increasing in m .*

The intuition behind monotonicity is easy to explain: if the agents other than agent i use a higher threshold, then they will volunteer more often. Thus, agent i is less likely to be chosen when he volunteers, and thus he will need to volunteer more often (and so use a higher threshold himself). Monotonicity is enough to guarantee the existence of

an equilibrium. We actually know that there is an equilibrium even without the use of monotonicity. If all agents choose a threshold of \$0, so no agent ever volunteers, then clearly i 's best response is also never to volunteer; getting a dollar is useless if it can never be spent. Fortunately, we can use monotonicity to show that there is a nontrivial equilibrium in threshold strategies as well. Indeed, to guarantee the existence of a nontrivial equilibrium, it suffices to show there is some vector \vec{k} of thresholds such that $BR_G(\vec{k}) > \vec{k}$.⁷ The following lemma, whose proof is again deferred to Appendix B, shows that we can always find such a point for sufficiently large δ_t .

LEMMA 5.4. *For all games $G = (T, \vec{f}, h, m, n)$, there exists a $\delta^* < 1$ such that if $\delta_t > \delta^*$ for all t , there is a vector \vec{k} of thresholds such that $BR_G(\vec{k}) > \vec{k}$.*

We are now ready to prove our main theorem: there exists a non-trivial equilibrium where all agents play threshold strategies greater than zero.

THEOREM 5.5. *For all games $G = (T, \vec{f}, h, m, 1)$ and all ϵ , there exist n_ϵ^* and $\delta_{\epsilon, n}^*$ such that, if $n > n_\epsilon^*$ and $\delta_t > \delta_{\epsilon, n}^*$ for all t , then there exists a nontrivial vector \vec{k} of thresholds that is an ϵ -Nash equilibrium. Moreover, there exists a greatest such vector (i.e. an equilibrium where the threshold used by each type is weakly higher than in every other equilibrium).*

PROOF. By Lemma 5.3, BR_G is a non-decreasing function on a complete lattice, so Tarski's fixed point theorem [Tarski 1955] guarantees the existence of a greatest and least fixed point; these fixed points are equilibria. The least fixed point is the trivial equilibrium. We can compute the greatest fixed point by starting with the strategy profile (∞, \dots, ∞) (where each agent uses the strategy S_∞ of always volunteering) and considering ϵ -best-reply dynamics, that is, iteratively computing the ϵ -best-reply strategy profile. Monotonicity guarantees this process converges to the greatest fixed point, which is an equilibrium (and is bound to be an equilibrium in pure strategies, since the best reply is always a pure strategy). Since there is a finite amount of money, this process needs to be repeated only a finite number of times. By Lemma 5.4, there exists a \vec{k} such that $BR_G(\vec{k}) > \vec{k}$. Monotonicity then guarantees that $BR_G(BR_G(\vec{k})) \geq BR_G(\vec{k})$ and similarly for any number of applications of BR_G . If \vec{k}^* is the greatest fixed point of BR_G , then $\vec{k}^* > \vec{k}$. Thus, the greatest fixed point is a nontrivial equilibrium. \square

The use of Tarski's fixed point theorem in the proof of Theorem 5.5 also provides an algorithm for finding equilibria that seems efficient in practice: start with the strategy profile (∞, \dots, ∞) and iterate the best-reply dynamics until an equilibrium is reached.

There is a subtlety in our results. In general, there may be many equilibria. From the perspective of social welfare, some will be better than others. As we show in our companion paper, strategies that use smaller (but nonzero) thresholds increase social welfare (intuitively, when agents are less willing to accumulate money, there are fewer agents with zero dollars and so fewer missed opportunities). Consider the best-reply function shown in Figure 1. In the game G in the example, there is only one type of agent, so $BR_G : \mathbb{N} \rightarrow \mathbb{N}$. In equilibrium, we must have $BR(k) = k$; that is, an equilibrium is characterized by a point on the line $y = x$. This example has three equilibria, where all agents play s_0 , s_5 , and s_{10} respectively. The strategy profile where all agents play s_5 is the equilibrium that maximizes social welfare, while s_{10} is the greatest equilibrium.

⁷We write $\vec{k}' \geq \vec{k}$ to denote $k'_t \geq k_t$ for all t , with strictness if the inequality for at least one t is strict.




Fig. 1. A hypothetical best-reply function with one type of agent.

In our companion paper, we focus on the greatest equilibrium in all our applications (although a number of our results hold for all nontrivial equilibria). This equilibrium has several desirable properties. First, it is guaranteed to be stable; best-reply dynamics from nearby points converge to it. By way of contrast, best-reply dynamics moves the system away from the equilibrium S_5 in Figure 1. Unstable equilibria are difficult to find in practice, and seem unlikely to be maintained for any length of time. Second, the “greatest” equilibrium is the one found by the natural algorithm given in Theorem 5.5. The proof of the theorem shows that it is also the outcome that will occur if agents adopt the reasonable initial strategy of starting with a large threshold and then using best-reply dynamics. Finally, by focusing on the worst nontrivial equilibrium, our results provide guarantees on social welfare, in the same way that results on *price of anarchy* [Roughgarden and Tardos 2002] provide guarantees (since price of anarchy considers the social welfare of the Nash equilibrium with the worst social welfare).

6. SIMULATIONS

Theorem 4.1 proves that, for a sufficiently large number n of agents, and after a sufficiently large number r of rounds, the distribution of wealth will almost always be close to the distribution that minimizes relative entropy. In this section, we simulate the game to gain an understanding of how large n and r need to be in practice. The simulations show that our theoretical results apply even to relatively small systems; we get tight convergence with a few thousand agents, and weaker convergence for smaller numbers, in very few rounds, indeed, a constant number per agent.

The first simulation explores the tightness of convergence to the distribution that minimizes relative entropy for various values of n . We used a single type of agent, with $\beta = \rho = \chi = 1$, $m = 2$, and $k = 5$. For each value of n , the simulation was started with a distribution of money as close as possible to the distribution d^* . (Recall that d^* is the distribution that minimizes relative entropy to the distribution q defined in Theorem 4.1, and that d^* characterizes the distribution of money in equilibrium when the threshold strategy 5 is used.)

We then computed the maximum Euclidean distance between d^* and the observed distribution over 10^6 rounds. As Figure 2 shows, the system does not move far from d^* once it is there. For example, if $n = 5000$, the system is never more than distance .001 from d^* . If $n = 25,000$, it is never more than .0002 from d^* .




Fig. 2. Maximum Euclidean distance from minimum relative entropy distribution d^* over 10^6 timesteps.




Fig. 3. Distance from minimum relative entropy distribution with 1000 agents.

Figure 2 does show a larger distance for $n = 1000$, although in absolute terms it is still small. The next simulation shows that, while the system may occasionally move away from d^* , it quickly converges back to it. We averaged 10 runs of the Markov chain, starting from an extreme distribution (every agent has either \$0 or \$5), and considered the average time needed to come within various distances of d^* . As Figure 3 shows, after 2 rounds per agent, on average, the Euclidean distance from the average distribution of money to d^* is .008; after 3 rounds per agent, the distance is down to .001.

Finally, we considered more carefully how quickly the system converges to d^* for various values of n . There are approximately k^n possible states, so the convergence time could in principle be quite large. However, we suspect that the Markov chain that arises here is *rapidly mixing*, which means that it will converge significantly faster (see [Lovasz and Winkler 1995] for more details about rapid mixing). We believe that the actually time needed is $O(n)$. This behavior is illustrated in Figure 4, which shows that for our example chain (again averaged over 10 runs), after approximately $3n$ steps, the Euclidean distance between the actual distribution of money in the system and d^* is less than .001. This suggests that we should expect the system to converge in a constant number of rounds per agent.

Note that this analysis assumes that agents follow the same strategy even when the system has a distribution of money far from d^* . If agents are aware that the system is




Fig. 4. Average time to get within .001 of the minimum relative entropy distribution.

in such a state and can find a better strategy, then convergence would not necessarily be so rapid. However, we would never expect to observe such a deviation in practice in even a moderately large system, and therefore these results help give confidence that the deviations we would expect to see (for which the equilibrium strategies will still be approximately optimal) will be short-lived. Even if such a large deviation were to somehow occur (e.g., if a bug led to unusual behavior) this rapid convergence means that the loss due to suboptimal play during that time would be small. As there is not an obvious profitable deviation, it seems reasonable to believe that players would continue to follow equilibrium behavior.

7. DISCUSSION

We have given a formal analysis of a scrip system and have shown that approximate equilibria exist in threshold strategies and that the distribution of money in these equilibria is given by relative entropy. As part of our equilibrium argument, we have shown that the best-reply function is monotone. This proves the existence of equilibria in pure strategies and permits efficient algorithms to compute these equilibria.

Our model makes a number of assumptions that are worthy of further discussion. Some of the simplifying assumptions can be relaxed without significant changes to our results (albeit at the cost of greater strategic and analytic complexity). At a high level, our results show the system converges to a steady state when agents follow simple threshold strategies and that there is in fact an equilibrium in these strategies. If, for example, rather than all requests having the same value to agent (γ_t), the value of a request is stochastic, agents might wish to have thresholds for each type of request. This would allow an agent to forgo a low-valued request if he is low on money. This makes the space of agent strategies larger and significantly complicates the proofs in the appendix, but this high-level characterization still holds.

The most significant assumption we make is that prices are fixed. However, our results provide insight even if we relax this assumption. With variable prices, the behavior of the system depends on the value of β , the probability that an agent can satisfy a request. For large β , where there are a large number of agents who can satisfy each request, we expect the resulting competition to effectively produce a fixed price, so our analysis applies directly. For small β , where there are few volunteers for each request, variable prices can have a significant impact.

However, allowing prices to be set endogenously, by bidding, has a number of negative consequences. For one thing, it removes the ability of the system designer to optimize the system using monetary policy. In addition, for small β , it is possible for

colluding agents to form a cartel to fix prices on a resource they control. It also greatly increases the strategic complexity of using the system: rather than choosing a single threshold, agents need an entire pricing scheme. Finally, the search costs and costs of executing a transaction are likely to be higher with floating prices. Thus, we believe that adopting a fixed price or a small set of fixed prices is often a reasonable design decision.

A. PROOF OF THEOREM 4.1

Given a Markov chain \mathcal{M} over a state space X and state $s \in S$, let $I_{\vec{x}, \vec{y}}^r$ be the random variable that is 1 if \mathcal{M} is in state \vec{y} at time r and the chain started in state \vec{x} and 0 otherwise. Then $\lim_{r \rightarrow \infty} \Pr(I_{\vec{x}, \vec{y}}^r = 1)$ is the limit probability of being in state \vec{y} given that the Markov chain starts in state \vec{x} . In general, this limit does not exist. However, there are well-known conditions under which the limit exists, and is independent of the initial state \vec{x} . A Markov chain is said to be *irreducible* if every state is reachable from every other state; it is *aperiodic* if, for every state \vec{x} , there exist two cycles from \vec{x} to itself such that the gcd of their lengths is 1.

THEOREM A.1. [Resnick 1992] *If \mathcal{M} is a finite, irreducible, and aperiodic Markov chain over state space X , then there exists a $d : X \rightarrow \mathbb{R}$ such that, for all \vec{x} and $\vec{y} \in X$, $\lim_{r \rightarrow \infty} \Pr(I_{\vec{x}, \vec{y}}^r = 1) = d(\vec{y})$.*

Thus, if we can show that \mathcal{M} is finite, irreducible, and aperiodic, then the limit distribution exists and is independent of the start state \vec{x} . This is shown in the following lemma.

LEMMA A.2. *If there are at least three agents, then \mathcal{M} is finite, irreducible, and aperiodic and therefore has a limit distribution π .*

PROOF. \mathcal{M} is clearly finite since X is finite. We prove that it is irreducible by showing that state \vec{y} is reachable from state \vec{x} by induction on the distance $w = \sum_{i=1}^n |x_i - y_i|$ (i.e., the sum of the absolute differences in the amount of money each person has in states \vec{x} and \vec{y}). If $w = 0$, then $\vec{x} = \vec{y}$ so we are done. Suppose that $w > 0$ and all pairs of states that are less than w apart are reachable from each other. Consider a pair of states \vec{x} and \vec{y} such that the distance from \vec{x} to \vec{y} is w . Since $w > 0$ and the total amount of money is the same in all states, there must exist i_1 and i_2 such that $x_{i_1} > y_{i_1}$ and $x_{i_2} < y_{i_2}$. Thus, in state \vec{y} , i_1 is willing to work (since he has strictly less than the threshold amount of money) and i_2 has money to pay him (since i_2 has a strictly positive amount of money). The state \vec{z} that results from i_1 doing work for i_2 in state \vec{y} is of distance $w - 2$ from \vec{x} . By the induction hypothesis, \vec{z} is reachable from \vec{x} . Since \vec{y} is clearly reachable from \vec{z} , \vec{y} is reachable from \vec{x} .

Finally, we must show that \mathcal{M} is aperiodic. Suppose \vec{x} is a state such that there exist three agents i_1, i_2 , and i_3 where i_1 has more than 0 dollars and i_2 and i_3 have less than their threshold amount of money. There must be such a state by our assumption that $mhn < \sum_t f_t k_t hn$. Clearly there is a cycle of length 2 from \vec{x} to itself: i_2 does work for i_1 and then i_2 does work for i_1 . There is also a cycle of length 3: i_2 does work for i_1 , i_3 does work for i_2 , then i_1 does work for i_3 . By irreducibility, identifying a single state with this property is sufficient. \square

We next give an explicit formula for the limit distribution. Recall that in the special case discussed in the main text, β_t, χ_t , and ρ_t were the same for all types, so the transition probabilities were symmetric and the limit distribution was uniform. While with more general values they are no longer symmetric, they still have significant structure that allows us to give a concise description of the limit distribution.

LEMMA A.3. *For all states \vec{x} of \mathcal{M} , let $w_{\vec{x}} = \prod_i (\beta_{\tau(i)} \chi_{\tau(i)} / \rho_{\tau(i)})^{x_i}$, and let $Z = \sum_{\vec{y}} w_{\vec{y}}$. Then the limit distribution of \mathcal{M} is $\pi(\vec{x}) = w_{\vec{x}} / Z$.*

PROOF. Define π by taking $\pi(\vec{x}) = w_{\vec{x}} / Z$, where $w_{\vec{x}}$ and Z are as in the statement of the lemma. If $T_{\vec{x}\vec{y}}$ is the probability of transitioning from state \vec{x} to state \vec{y} , it is well known that it suffices to show that π satisfies the *detailed balance condition* [Resnick 1992], i.e., $\pi(\vec{x}) T_{\vec{x}\vec{y}} = \pi(\vec{y}) T_{\vec{y}\vec{x}}$ for all states \vec{x} and \vec{y} and π is a probability measure. The fact that π is a probability measure is immediate from its definition. To check the first condition, let \vec{x} and \vec{y} be adjacent states such that \vec{y} is reached from \vec{x} by i spending a dollar and j earning a dollar. This means that for the transition from \vec{x} to \vec{y} to happen, i must be chosen to spend a dollar and j must be able to work and chosen to earn the dollar. Similarly for the reverse transition to happen, j must be chosen to spend a dollar and i must be able to work and chosen to earn the dollar. All other agents have the same amount of money in each state, and so will make the same decision in each state. Thus the probabilities associated with each transition differ only in the relative likelihoods of i and j being chosen at each point. These may differ for three reasons: one might be more likely to be able to satisfy requests (β), to want to make requests (ρ), or to be chosen to satisfy requests (χ). Thus, for some p , which captures the effect of other agents volunteering on the likelihood of i and j being chosen, we can write the transition probabilities as $T_{\vec{x}\vec{y}} = \rho_{\tau(i)} \beta_{\tau(j)} \chi_{\tau(j)} p$ and $T_{\vec{y}\vec{x}} = \rho_{\tau(j)} \beta_{\tau(i)} \chi_{\tau(i)} p$. From the definition of π , we have that

$$\frac{\pi(\vec{x})}{\pi(\vec{y})} = \frac{\beta_{\tau(i)} \chi_{\tau(i)} \rho_{\tau(j)}}{\rho_{\tau(i)} \beta_{\tau(j)} \chi_{\tau(j)}} = \frac{T_{\vec{y}\vec{x}}}{T_{\vec{x}\vec{y}}}.$$

Thus, $\pi(\vec{x}) T_{\vec{x}\vec{y}} = \pi(\vec{y}) T_{\vec{y}\vec{x}}$, as desired. \square

Note that for the special case considered in the main text, Lemma A.3 shows that the limit distribution is the uniform distribution.

The limit distribution tells us the long run probability of being in a given state. Theorem 4.1 does not mention states directly, but rather the distributions of money associated with a state. In order to prove the theorem, we need to know the probability of being in some state associated with a given distribution. This is established in the following lemma.

LEMMA A.4. *Let π be the limit distribution from Lemma A.3, and let $V(d) = H(d) - H(\vec{f}) - \log Z + \sum_t \sum_{i=0}^{k_t} id(t, i) \log \omega_t$ (where H is the standard entropy function; that is, $H(d) = \sum_{t,i} d(t, i) \log d(t, i)$). For all $d \in \Delta_{\vec{f}, m, \vec{k}}$, either $\pi(\{\vec{x} \mid d^{\vec{x}} = d\}) = 0$ or $F(hn) e^{hnV(d)} \leq \pi(\{\vec{x} \mid d^{\vec{x}} = d\}) \leq G(hn) e^{hnV(d)}$, where F and G are polynomials.*

PROOF. Before computing the probability of being in such a state, we first compute the number of states. It is possible that there is no state \vec{x} such that $d = d^{\vec{x}}$ (e.g., if hn is odd and d has half the agents with 0 dollars). If there is such a state \vec{x} , each such state has $hnd(t, i)$ agents of type t with i dollars. Thus, the number of states \vec{x} with $d = d^{\vec{x}}$ is the number of ways to divide the agents into groups of these sizes. Since there are hnf_t agents of type t , the number of such states is

$$\prod_t \binom{hnf_t}{hnd(t, 0), \dots, hnd(t, k_t)}.$$

To complete the proof, we use the fact (shown in the proof of Lemma 3.11 of [Grove et al. 1994]) that

$$\frac{1}{F(hn)} e^{hnf_t H(d_t)} \leq \binom{hnf_t}{hnd(t, 0), \dots, hnd(t, k_t)} \leq G(hn) e^{hnf_t H(d_t)},$$

where F and G are polynomial in hn , and d_t is the distribution restricted to a single type t (i.e., $d_t(i) = d(t, i) / \sum_i d(t, i)$). The (*generalized*) *grouping property* [Cover and Thomas 1991] of entropy allows us to express $H(d)$ in terms of the entropy of the distributions for each fixed t , or the $H(d_t)$. Because $f_t = \sum_i d(t, i)$, this has the particularly simple form $H(d) = H(\vec{f}) + \sum_t f_t H(d_t)$. Thus, up to a polynomial factor, the number of such states is

$$\prod_t e^{hn f_t H(d_t)} = e^{hn(\sum_t f_t H(d_t))} = e^{hn(H(d) - H(\vec{f}))}.$$

By Lemma A.3, each of these states has the same probability $\pi(\vec{x})$. Thus, dropping the superscript \vec{x} on $d^{\vec{x}}$ for brevity, the probability of being in such a state is:

$$\begin{aligned} e^{hn(H(d) - H(\vec{f}))} \pi(\vec{x}) &= e^{hn(H(d) - H(\vec{f}))} \prod_i (\beta_i \chi_t / \rho_i)^{x_i} / Z \\ &= e^{hn(H(d) - H(\vec{f}))} Z^{-hn} \prod_i (\omega_t)^{x_i} \\ &= e^{hn(H(d) - H(\vec{f}))} Z^{-hn} \prod_t \prod_{i=0}^{k_t} (\omega_t)^{h n i d(t, i)} \\ &= e^{hn(H(d) - H(\vec{f}))} Z^{-hn} \prod_t \prod_{i=0}^{k_t} e^{h n i d(t, i) \log \omega_t} \\ &= e^{hn(H(d) - H(\vec{f}) - \log Z + \sum_t \sum_{i=0}^{k_t} i d(t, i) \log \omega_t)} \\ &= e^{hn V(d)} \end{aligned}$$

□

Theorem 4.1 says that there exists a $q \in \Delta_{\vec{f}, m, \vec{k}}$ (i.e., a probability distribution on agent types t and amounts of money i) with certain properties. We now define the appropriate q . Let

$$q(t, i) = (\omega_t)^i / \left(\sum_t \sum_{j=0}^{k_t} (\omega_t)^j \right). \quad (5)$$

It is not immediately clear why this is the right choice of q . As the following lemma shows, this definition allows us to characterize the distribution that maximizes the probability of being in a state corresponding to that distribution (as given by Lemma A.4) in terms of relative entropy.

LEMMA A.5. *The unique maximum of $V(d) = H(d) - H(\vec{f}) - \log Z + \sum_t \sum_{i=0}^{k_t} i d_i^t \log \omega_t$ on $\Delta_{\vec{f}, m, \vec{k}}$ occurs at d_q^* .*

PROOF. For brevity, we drop the superscript \vec{x} on d and let $Y = \sum_t \sum_j (\omega_t)^j$.

$$\begin{aligned}
\operatorname{argmax}_d V(d) &= \operatorname{argmax}_d (H(d) - H(\vec{f}) - \log Z + \sum_t \sum_{i=0}^{k_t} id(t, i) \log \omega_t) \\
&= \operatorname{argmax}_d (H(d) + \sum_t \sum_{i=0}^{k_t} id(t, i) \log \omega_t) \\
&= \operatorname{argmax}_d \sum_t \sum_{i=0}^{k_t} [-d(t, i) \log d(t, i) + id(t, i) \log \omega_t] \\
&= \operatorname{argmax}_d \sum_t \sum_{i=0}^{k_t} [-d(t, i) \log d(t, i) + d(t, i) \log(q(t, i)Y)] \\
&= \operatorname{argmax}_d \sum_t \sum_{i=0}^{k_t} [-d(t, i) \log d(t, i) + d(t, i) \log q(t, i) + d(t, i) \log Y] \\
&= \operatorname{argmax}_d \log Y + \sum_t \sum_{i=0}^{k_t} [-d(t, i) \log d(t, i) + d(t, i) \log q(t, i)] \\
&= \operatorname{argmin}_d \sum_t \sum_{i=0}^{k_t} [d(t, i) \log d(t, i) - d(t, i) \log q(t, i)] \\
&= \operatorname{argmin}_d \sum_t \sum_{i=0}^{k_t} d(t, i) \log \frac{d(t, i)}{q(t, i)} \\
&= \operatorname{argmin}_d H(d||q).
\end{aligned}$$

By definition, d_q^* minimizes $H(d||q)$. It is unique because H (and thus V) is a strictly concave function on a closed convex set. \square

Lemma A.5 tells us that the most likely distributions of money to be observed are those with low relative entropy to q . Among all distributions in $\Delta_{\vec{f}, m, k}$, relative entropy is minimized by d_q^* . However, given n , it is quite possible that d_q^* is not $d^{\vec{x}}$ for any \vec{x} . For example, if $d_q^*(t, i) = 1/3$ for some t and i , but $f_t hn = 16$, then $d^{\vec{x}}(t, i) = d_q^*(t, i)$ only if exactly $16/3$ agents of type t to have i dollars, which cannot be the case. However, as the following lemma shows, for sufficiently large n , we can always find a $d^{\vec{x}}$ that is arbitrarily close to d_q^* . For convenience, we use the 1-norm as our notion of distance.

LEMMA A.6. *For all ϵ , there exists n_ϵ such that, if $n > n_\epsilon$, then for some state \vec{x} , $\|d^{\vec{x}} - d_q^*\| < \epsilon$.*

PROOF. Given n , we construct $d \in \Delta_{\vec{f}, m, k}$ that is of the form $d^{\vec{x}}$ and is close to d_q^* in a number of steps. As a first step, for all t and i , let $d_1(t, i)$ be the result of rounding $d_q^*(t, i)$ to the nearest $1/hn$ (where ties are broken arbitrarily). The function d_1 may not be in $\Delta_{\vec{f}, m, k}$; we make minor adjustments to it to get a function in $\Delta_{\vec{f}, m, k}$. First, note that we may have $\sum_i d_1(t, i) \neq f_t$. Since f_t is a multiple of $1/hn$, can get a function d_2 that satisfies these constraints by modifying each term $d_1(t, i)$ by either adding $1/hn$ to it, subtracting $1/hn$ from it, or leaving it alone. Such a function d_2 may still violate the final constraint that $\sum_{t, i} id_2(t, i) = m$. We construct a function d_3 that satisfies this constraint (while continuing to satisfy the constraint that $\sum_i d_3(t, i) = f_t$) as follows.

Note that if we increase $d_2(t, i)$ by $1/hn$ and decrease $d_2(t, j)$ by $1/hn$, then we keep the keep $\sum_i d_2(t, i) = f_t$, and change $\sum_i id_2(t, i)$ by $(i - j)/hn$. Since each term $d_2(t, i)$ is a multiple of $1/hn$ and m is a multiple of $1/h$, we can perform these adjustments until all the constraints are satisfied.

The rounding to create d_1 changed each $d_1(t, i)$ by at most $1/hn$, so $\|d_q^* - d_1\|_1 \leq (\sum_t k_t + 1)/hn$. Since, each term $d_1(t, i)$ was changed by at most $1/hn$ to obtain $d_2(t, i)$, we have $\|d_1 - d_2\|_1 \leq (\sum_t k_t + 1)/hn$. Let $c = \max_t(\max(k_t - m, m))$. Each movement of up to $1/hn$ in the creation of d_1 and d_2 altered m by at most c/hn . Thus at most $2c$ movements are needed in the creation of d_3 for each pair (t, i) . Therefore, $\|d_2 - d_3\|_1 \leq (\sum_t k_t + 1)2c/hn$. By the triangle inequality, $\|d_q^* - d_3\|_1 \leq (\sum_t k_t + 1)(2c + 2)/hn$, which is $O(1/n)$. Hence, for n_ϵ sufficiently large, the resulting d_3 will always be within distance ϵ of d_q^* .

Finally, we need to show that $d_3 = d^{\vec{x}}$ for some \vec{x} . Each $d_3(t, i)$ is a multiple of $1/hn$. There are hn agents in total, so we can find such an \vec{x} by taking any allocation of money such that $d_3(t, i)hn$ agents of type t have i dollars. \square

We are now ready to prove Theorem 4.1. We repeat the statement here for the reader's convenience.

THEOREM 4.1. *For all games $(T, \vec{f}, h, m, 1)$, all vectors \vec{k} of thresholds, and all $\epsilon > 0$, there exist $q \in \Delta_{\vec{f}, m, \vec{k}}$ and n_ϵ such that, for all $n > n_\epsilon$, there exists a round r^* such that, for all $r > r^*$, we have $\Pr(I_{q, n, \epsilon}^r = 1) > 1 - \epsilon$.*

PROOF. From Lemma A.3, we know that, after a sufficient amount of time, the probability of being in state \vec{x} will be close to $\pi_{\vec{x}} = w_{\vec{x}}/Z$. Since \mathcal{M} converges to a limit distribution, it is sufficient to show that the theorem holds in the limit as $r \rightarrow \infty$. If the theorem holds in the limit for some $\epsilon' < \epsilon$, then we can take r large enough that the L1 distance between the distribution of the chain at time r and the limit distribution (i.e. treating the distributions as vectors and computing the sum of the absolute values of their differences) is at most $\epsilon - \epsilon'$.

The remainder of the proof is essentially that of Theorem 3.13 in [Grove et al. 1994] (applied in a very different setting). Let $V(d) = H(d) - H(\vec{f}) - \log Z + \sum_t \sum_{i=0}^{k_t} id(t, i) \log \omega_t$. We show there exists a value v_L such that, for all states \vec{x} such that $d^{\vec{x}}$ is not within ϵ of d_q^* , we have $V(d^{\vec{x}}) \leq v_L$, and a value $v_H > v_L$ such that $v_H = V(d^{\vec{y}})$ for some point \vec{y} such that $d^{\vec{y}}$ is within distance ϵ of d_q^* . Lemma A.4 then shows that it is exponentially more likely that $d^{\vec{x}^r} = d^{\vec{y}}$ than any distribution d such that $V(d) \leq v_L$. If $\vec{x}^r = \vec{y}$ then $I_{q, n, \epsilon}^r = 1$, and if $I_{q, n, \epsilon}^r = 0$ then $V(d^{\vec{x}^r}) \leq v_L$, so this suffices to establish the theorem.

By Lemma A.5, the unique maximum of V on $\Delta_{\vec{f}, m, \vec{k}}$ occurs at d_q^* . The set $\{d \in \Delta_{\vec{f}, m, \vec{k}} \mid \|d_q^* - d\|_2 \geq \epsilon\}$ is closed. V is a continuous function, so it takes some maximum v_L on this set. Pick some v_H such that $v_L < v_H < V(d_q^*)$. By the continuity of V , there exists an ϵ such that if $\|d_q^* - d\|_1 < \epsilon$ then $V(d) \geq v_H$. By Lemma A.6, for sufficiently large n , there is always some \vec{x} such that $\|d_q^* - d^{\vec{x}}\|_1 < \epsilon$. Thus, for some $\vec{x} \in X_{\epsilon, q}$, $V(d^{\vec{x}}) \geq v_H$.

$\Pr(I_{q, n, \epsilon}^r = 1) \geq \Pr(\vec{x}^r \in \{\vec{y} \mid d^{\vec{y}} = d^{\vec{x}}\})$. By Lemma A.4, $\Pr(I_{q, n, \epsilon}^r = 1)$ is at least $1/F(hn)e^{hnV(d^{\vec{x}})} \geq 1/F(hn)e^{hnv_H}$. Now consider a \vec{y} such that $I_{q, n, \epsilon}(\vec{y}) = 0$. By Lemma A.4, the probability that $d^{\vec{x}^r} = d^{\vec{y}}$ is at most $G(hn)e^{hnV(d^{\vec{y}})} \leq G(hn)e^{hnv_L}$. There are at most $(hn + 1)^{\sum_t (k_t + 1)}$ such points, a number which is polynomial in hn . Thus, for $G'(hn) = G(hn)(hn + 1)^{\sum_t (k_t + 1)}$, the probability that $I_{q, n, \epsilon}^r = 0$ is at most

$G'(hn)e^{hvv_L}$. The ratio of these probabilities is at most

$$\frac{G'(hn)e^{hvv_L}}{\frac{1}{F(hn)}e^{hvv_H}} = \frac{G'(hn)F(hn)}{e^{hn(v_H-v_L)}}.$$

This is the ratio of a polynomial to an exponential, so the probability of seeing a distribution of distance greater than ε from d_q^* goes to zero as n goes to infinity. \square

B. PROOFS FROM SECTION 5

In this appendix, we provide the omitted proofs from Section 5.

The proof of Theorem 5.2 relies on modeling the game from the perspective of a single agent. Consider a vector \vec{k} of thresholds and the corresponding strategy profile $\vec{S}(\vec{k})$. Fix an agent i of type t . Assume that all the agents other than i continue play their part of $\vec{S}(\vec{k})$. What is i 's best response? Since the set of agents is large, i 's choice of strategy will have (essentially) no impact on the distribution of money. By Theorem 4.1, the distribution of money will almost always be close to a distribution d^* . Suppose, the distribution were exactly d^* . Since we know the exact distribution of money and the thresholds used by the other agents, we can calculate the number of each type of agent that wish to volunteer and thus the probabilities that our single agent will be able to earn or spend a dollar. Thus, by assuming the distribution of money is always exactly d^* , we can model the game from the perspective of agent i as a Markov Decision Process (MDP). We show in Lemma B.2 that this MDP has an optimal threshold policy. (Threshold policies are known as *monotone* policies in the more general setting where there are more than two actions.) We then prove that any optimal policy for the MDP is an ε -best reply to the strategies of the other agents in the actual game.

Taking notation from Puterman [1994], we formally define the MDP $\mathcal{P}_{G, \vec{S}(\vec{k}), t} = (S, A, p(\cdot | s, a), r(s, a))$ that describes the game where all the agents other than i are playing $\vec{S}(\vec{k})_{-i}$ and i has type t .

- $S = \{0, \dots, mhn\}$ is the set of possible states for the MDP (i.e., the possible amounts of money compatible with the distribution d^*).
- $A = \{0, 1\}$ is the set of possible actions for the agent, where 0 denotes not volunteering and 1 denotes volunteering iff another agent who has at least one dollar makes a request.
- p_u is the probability of earning a dollar, assuming the agent volunteers (given that all other agents have fixed their thresholds according to \vec{k} and the distribution of money is exactly d^*). Each agent of type t' who wishes to volunteer can do so with probability $\beta_{t'}$. Assuming exactly the expected number of agents are able to volunteer, $v_{t'} = \beta_{t'}(f_{t'} - d^*(t', k_{t'}))n$ agents of type t' volunteer. Note that we are disregarding the effect of i in computing the $v_{t'}$, since this will have a negligible effect for large n . Using the $v_{t'}$ s, we can express p_u as the product of two probabilities: that some agent other than i who has a dollar is chosen to make a request and that i is the agent chosen to satisfy it. Thus,

$$p_u = \left(\sum_{t'} \rho_{t'}(f_{t'} - d^*(t', 0)) \right) \left(\frac{\chi_t \beta_t}{\sum_{t'} \chi_{t'} v_{t'}} \right). \quad (6)$$

- p_d is the probability of agent i having a request satisfied, given that agent i has a dollar. Given that all agents are playing a threshold strategy, if the total number n of agents is sufficiently large, then it is almost certainly the case that some agent will always be willing and able to volunteer. Thus, we can take p_d to be the probability

that agent i will be chosen to make a request; that is,

$$p_d = \frac{\rho_t}{n} \quad (7)$$

- $r(s, a)$ is the (immediate) expected reward for performing action a in state s . Thus, $r(s, 0) = \gamma_t p_d$ if $s > 0$; $r(0, 0) = 0$; $r(s, 1) = \gamma_t p_d - \alpha_t p_u$ if $s > 0$; and $r(0, 1) = -\alpha_t p_u$.
- $p(s' | s, a)$ is the probability of being in state s' after performing action a in state s ; $p(s' | s, a)$ is determined by p_u and p_d ; specifically, $p(s+1 | s, 1) = p_u$, $p(s-1 | s, a) = p_d$ if $s > 0$, and the remainder of the probability is on $p(s | s, a)$ (i.e., $p(s | s, a) = 1 - (p(s+1 | s, 1) + p(s-1 | s, a))$).
- $u^*(s)$ is the expected utility of being in state s if agent i uses the optimal policy for the MDP $\mathcal{P}_{G, \vec{s}(\bar{k}), t}$.
- $u(s, a)$ is the expected utility for performing action a in state s , given that the optimal strategy is followed after this action;

$$u(s, a) = r(s, a) + \delta \sum_{s'=0}^{mhn} p(s' | s, a) u^*(s').$$

To prove Theorem 5.2, we need two preliminary lemmas about the MDP $\mathcal{P}_{G, \vec{s}(\bar{k}), t}$.

LEMMA B.1. *For the MDP $\mathcal{P}_{G, \vec{s}(\bar{k}), t}$, $u^*(s+2) + u^*(s) \leq 2u^*(s+1)$.*

PROOF. The MDP $\mathcal{P}_{G, \vec{s}(\bar{k}), t}$ has an optimal stationary policy [Puterman 1994, Theorem 6.2.10] (a policy where the chosen action depends only on the current state). Let π be such a policy. Consider the policy π' starting in state $s+1$ that “pretends” it actually started in state s and is following π . More precisely, if $s_0 = s+1$ and $s_j > 0$ for $j = 0, \dots, k$, define $\pi'(s_0, s_1, \dots, s_k) = \pi(s_k - 1)$; otherwise, if $j \leq k$ is the least index such that $s_j = 0$, define $\pi'(s_0, \dots, s_k) = \pi(s_k)$. Given a history (s_0, \dots, s_k) , j is the random variable whose value is the minimum i such that $s_i = 0$ or ∞ if no such value exists. The definition of π' from π creates a bijection between histories that start in state $s+1$ and histories that start in state s , such that if h' corresponds to h , the probability of history h' with policy π' is the same as the probability of h with policy π . Technically, making the mapping a bijection requires the introduction of a new state $0'$, which intuitively represents the state where the agent has zero dollars and missed an opportunity to have a request satisfied last round because of it. More formally, we let $p(0' | 0, a) = p_d$ and $p(s | 0', a) = p(s | 0, a)$. With this change, the probabilities of corresponding histories are the same because the probability of transitioning from a state to the one “immediately below” it (where $s-1$ is immediately below s , $0'$ is immediately below 0 , and $0'$ is immediately below itself) is always p_d , and the probability of transitioning from a state to the one “immediately above” it (where $s+1$ is immediately above s , and 1 is immediately above $0'$) is always p_u .⁸

This argument shows that an agent starting with $s+1$ dollars “pretending” to start with s will have the same expected reward each round as an agent who actually started with s dollars, except during the first round j in a history such that $s_j = 0$. Thus (treating j as a random variable), we have

$$u^*(s+1) \geq u^*(s) + E[\delta^j \gamma_t].$$

⁸Note that this means that $0'$ is immediately below 0 but 1 is immediately above $0'$. This is intended, because $0'$ intuitively represents the state where the agent has 0 dollars and had a request go unsatisfied due to a lack of money in the previous round, so if he then earns a dollar he will have 1 dollar regardless of whether or not his request of two rounds previous was satisfied.

Similarly, we can use π starting from state $s + 2$ to define a policy π'' starting from state $s + 1$, where i “pretends” he has one more dollar and is using π , up to the first round j' that he is chosen to make a request with π in a state where he has no money (in which case he can make the request with π started from $s + 2$, but cannot make it with π'' started from $s + 1$); from that point on, he uses π . For corresponding histories, the utilities of an agent starting with $s + 1$ dollars and following π'' and an agent starting with $s + 2$ dollars and following π will be the same, except during round j' the agent following π will have a request satisfied but the agent following π'' will not. Thus,

$$u^*(s + 1) \geq u^*(s + 2) - E[\delta^{j'} \gamma_t].$$

Since, if i uses π , he will run out of money sooner if he starts with s dollars than if he starts with $s + 2$ dollars,

$$E[\delta^j \gamma_t] > E[\delta^{j'} \gamma_t].$$

Thus, $u^*(s + 2) + u^*(s) \leq 2u^*(s + 1)$. \square

LEMMA B.2. $\mathcal{P}_{G, \vec{s}(\bar{k}), t}$ has an optimal threshold policy.

PROOF. As shown by Puterman [1994, Lemma 4.7.1], it suffices to prove that $u(s, a)$ is subadditive. That is, we need to prove that, for all states s ,

$$u(s + 1, 1) + u(s, 0) \leq u(s + 1, 0) + u(s, 1). \quad (8)$$

We consider here only the case that $s > 0$ (the argument is essentially the same if $s = 0$). Because $s > 0$, $r(s + 1, a) = r(s, a)$, so (8) is equivalent to

$$\begin{aligned} & p_u u^*(s + 2) + p_d u^*(s) + (1 - p_u - p_d) u^*(s + 1) + p_d u^*(s - 1) + (1 - p_d) u^*(s) \\ & \leq p_d u^*(s) + (1 - p_d) u^*(s + 1) + p_u u^*(s + 1) + p_d u^*(s - 1) + (1 - p_u - p_d) u^*(s). \end{aligned}$$

This simplifies to

$$u^*(s + 2) + u^*(s) \leq 2u^*(s + 1),$$

which follows from Lemma B.1. \square

We can now prove Lemma 5.1 and Theorem 5.2.

LEMMA 5.1. Consider the games $G_n = (T, \vec{f}, h, m, n)$ (where T , \vec{f} , h , and m are fixed, but n may vary). There exists a k such that for all n , s_k is an optimal policy for $\mathcal{P}_{G_n, \vec{s}(\bar{k}), t}$. The threshold k is the maximum value of κ such that

$$\alpha_t \leq E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u, p_d)} \gamma_t], \quad (4)$$

where $J(\kappa, p_u, p_d)$ is a random variable whose value is the first round in which an agent starting with κ dollars, using strategy s_κ , and with probabilities p_u and p_d of earning a dollar and of being chosen given that he volunteers, respectively, runs out of money.

PROOF. Fix n . Suppose that an agent is choosing between a threshold of κ and a threshold of $\kappa + 1$. These policies only differ when the agent has κ dollars: he will volunteer with the latter but not with the former. If he volunteers when he has κ dollars and is chosen, he will pay a cost of α_t and he will have $\kappa + 1$ dollars. As in the proof of Lemma B.1, we can define a bijection on histories such that, in corresponding histories of equal probability, an agent who started with κ dollars and is using s_κ will always have one less dollar than an agent who started with $\kappa + 1$ dollars and is using $s_{\kappa+1}$, until the first round r in which the agent using $s_{\kappa+1}$ has zero dollars. This means that in round $r - 1$ the agent using $s_{\kappa+1}$ had a request satisfied but the agent using s_κ

was unable to because he had no money. Thus, if the agent volunteers when he has κ dollars and pays a cost of α_t in the current round, the expected value of being able to spend that dollar in the future is $E[(1 - (1 - \delta_t)/n)^{J(\kappa+1, p_u, p_d)}] \gamma_t$. Since this expectation is strictly increasing in κ (an agent with more money takes longer to spend it all), the maximum κ such that Equation (4) holds is an optimal threshold policy.

Taking the maximum value of κ that satisfies Equation (4) ensures that, for the n we fixed, we chose the maximum optimal threshold. We now need to show that this maximum optimal threshold is independent of n , which we do by showing that the expected utility of every threshold policy s_k is independent of n . The expected utility of a policy depends on the initial amount of money, but since an agent's current amount of money is a random walk whose transition probabilities are determined by p_u and p_d , there is a well-defined limit probability

$$x_i^* = \lim_{r \rightarrow \infty} \Pr(\text{agent has } i \text{ dollars in round } r)$$

determined by the ratio p_u/p_d (this is because the limit distribution satisfies the detailed balance condition: $x_i^* p_u = x_{i+1}^* p_d$). This distribution has the property that if the agent starts with i dollars with probability x_i^* , then in every round the probability he has i dollars is x_i^* . Thus, in each round his expected utility is $\gamma p_d(1 - x_0^*) - \alpha p_u(1 - x_k^*)$. We can factor out n to write $p_u = p'_u/n$ and $p_d = p'_d/n$ where p'_u and p'_d are independent of n . Note that $p_u/p_d = p'_u/p'_d$, so the x_i^* 's are independent of n . Thus, we can rewrite the agent's expected utility for each round as c/n , where $c = \gamma p'_d(1 - x_0^*) - \alpha p'_u(1 - x_k^*)$ is independent of n . Therefore, the expected utility of s_k is

$$\sum_{r=0}^{\infty} \left(1 - \frac{1 - \delta_t}{n}\right)^r \frac{c}{n} = \frac{c}{1 - \delta_t},$$

which is independent of n . \square

THEOREM 5.2. *For all games $G = (T, \vec{f}, h, m, n)$, all vectors \vec{k} of thresholds, and all $\varepsilon > 0$, there exist n_ε^* and $\delta_{\varepsilon, n}^*$ such that for all $n > n_\varepsilon^*$, types $t \in T$, and $\delta_t > \delta_{\varepsilon, n}^*$, an optimal threshold policy for $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ is an ε -best reply to the strategy profile $\vec{S}(\vec{k})_{-i}$ for every agent i of type t .*

PROOF. By Lemma B.2, $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ has an optimal threshold policy. However, this might not be a best reply for agent i in the actual game if the other agents are playing $\vec{S}(\vec{k})$. $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ assumes that the probabilities of earning or spending a dollar in a given round are always exactly p_u and p_d respectively. Theorem 4.1 guarantees only that, in the game, the corresponding probabilities are close to p_u and p_d with high probability after some amount of time that can depend on n . A strategy S for player i in G defines a policy π_S for $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$ in the obvious way; similarly, a policy for the MDP determines a strategy for player i in the game. The expected utility of π_S is close to $U_i(S, \vec{S}(\vec{k})_{-i})$, but is, in general, not equal to it, because, as we noted, p_u and p_d may differ from the corresponding probabilities in the game. They differ for three reasons: (1) they are close, but not identical; (2) they are only close with high probability, and (3) they are only close after some amount of time. As we now show, given ε , the difference in the expected utility due to each reason can be bounded by $\varepsilon/6$, so the expected utility of any strategy is within $\varepsilon/2$ of the value the corresponding policy in $\mathcal{P}_{G, \vec{S}(\vec{k}), t}$. Thus, an optimal strategy for the MDP is an ε -best reply.

As we have seen, the probabilities p_u and p_d are determined by the number of agents of each type that volunteer (i.e., the expressions $v_{t'}$ for each type t'). The distance between $d^{\vec{x}}$ and d^* bounds how much the actual number of agents of type t' that wish

to volunteer in round r can differ from $v_{t'}/\beta_{t'}$. Even if exactly $v_{t'}/\beta_{t'}$ agents wish to volunteer for each type t' , there might not be exactly $v_{t'}$ agents who actually volunteer because of the stochastic decision by nature about who can volunteer and because i cannot satisfy his own requests. However, for sufficiently large n , the effect on p_u and p_d from these two factors is arbitrarily close to zero. Applying Theorem 4.1, there exist n_1 and r_1 such that if there are at least n_1 agents, for all round $r > r_1$, $d^{\vec{x}}$ and d^* are sufficiently close that the difference between the utility of policy $\pi_{S'}$ in the MDP and $U_i((S', \vec{S}_{-i})$ in rounds $r > r_1$ where d^* is sufficiently close is at most $\varepsilon/6$.

Note that the maximum possible difference in utility between a round of the MDP and a round of the game is $\gamma + \alpha$ (if agent i spends a dollar rather than earning one). Applying Theorem 4.1 again, for $e = \varepsilon/6(\gamma + \alpha)$, there exist n_2 and r_2 such that the probability of the distribution not being within e of d^* is less than e . Thus, the difference between the expected utility of policy $\pi_{S'}$ in the MDP and $U_i((S', \vec{S}_{-i})$ in rounds $r > r_2$ where d^* is not sufficiently close is at most $e(\gamma + \alpha) = \varepsilon/6$.

Let $n_\varepsilon^* = \max(n_1, n_2)$ and $r^* = \max(r_1, r_2)$. The values of n_ε^* and r^* do not depend on δ , so we can take $\delta_{\varepsilon, n}^*$ to be sufficiently close to 1 that the total utility from the first r^* rounds is at most $\varepsilon/6$, completing the proof of the theorem. \square

Recall that BR_G maps a vector \vec{k} describing the threshold strategy for each type to a vector \vec{k}' of best replies.

LEMMA 5.3. *Consider the family of games $G_m = (T, \vec{f}, h, m, n)$ and the strategies $\vec{S}(\vec{k})$, for $mhn < \sum_t f_t k_t hn$. For this family of game, $\lambda_{m, \vec{k}}$ is non-decreasing in m and non-increasing in \vec{k} ; $p_u^{m, \vec{k}}$ is non-decreasing in m and non-increasing in \vec{k} ; and the function BR_G is non-decreasing in \vec{k} and non-increasing in m .*

PROOF. We first show that that $\lambda_{m, \vec{k}}$ is monotone in m and \vec{k} . We then show that $p_u^{m, \vec{k}}$ is a monotone function of $\lambda_{m, \vec{k}}$ and that BR_G is a monotone function of $p_u^{m, \vec{k}}$, completing the proof.

We now show that $\lambda_{m, \vec{k}}$ is non-decreasing in m . Fix a vector of thresholds \vec{k} and let

$$g_{\vec{k}}(\lambda) = \sum_{t, i} i \frac{f_t \lambda^i q_{\vec{k}}(t, i)}{\sum_{j=0}^{k_t} \lambda^j q_{\vec{k}}(t, j)}, \quad (9)$$

where $q_{\vec{k}}$ is the value of q from Equation (5) (we add the subscript \vec{k} to stress the dependence on \vec{k}). The definition of $\lambda_{m, \vec{k}}$ in Equations (2) and (3) in Lemma 4.2 ensures that, for all m , $m = g_{\vec{k}}(\lambda_{m, \vec{k}})$. A relatively straightforward computation shows that $g_{\vec{k}}'(\lambda) > 0$ for all λ . Thus, if $m' > m$, $g_{\vec{k}}(\lambda) = m$, and $g_{\vec{k}}(\lambda') = m'$, we must have $\lambda' > \lambda$. It follows that $\lambda_{m, \vec{k}}$ is increasing in m . (Note that $\lambda_{m, \vec{k}}$ is undefined for $m \geq \sum_t f_t k_t$, which is why monotonicity holds only for values of m such that $mhn < \sum_t f_t k_t$.)

We next show that $\lambda_{m, \vec{k}}$ is non-increasing in \vec{k} . Since we have a finite set of types, it suffices to consider the case where a single type t^* increases its threshold by 1. Let \vec{k} denote the initial vector of thresholds, and let \vec{k}' denote the vector of thresholds after agents of type t^* increase their threshold by 1; that is, $k_t = k'_t$ for $t \neq t^*$, and $k'_{t^*} = k_{t^*} + 1$.

The first step in showing that $\lambda_{m, \vec{k}}$ is non-increasing in \vec{k} is to show that $g_{\vec{k}'}(\lambda_{m, \vec{k}}) > g_{\vec{k}}(\lambda_{m, \vec{k}}) = m$. We do this by breaking the sum in the definition of g in Equation (9) into two pieces; those terms where $t \neq t^*$, and those where $t = t^*$.

It follows immediately from Equation (5) that there exists a constant c such that, for all i and $t \neq t^*$, we have $q_{\vec{k}'}(t, i) = cq_{\vec{k}}(t, i)$. It follows from Equation (2) that for all i and $t \neq t^*$, since $k_t = k'_t$, we have

$$i \frac{f_t \lambda_{m, \vec{k}}^i q_{\vec{k}'}(t, i)}{\sum_{j=0}^{k'_t} \lambda_{m, \vec{k}}^j q_{\vec{k}'}(t, j)} = i \frac{f_t \lambda_{m, \vec{k}}^i c q_{\vec{k}}(t, i)}{\sum_{j=0}^{k'_t} \lambda_{m, \vec{k}}^j c q_{\vec{k}}(t, j)} = i \frac{f_t \lambda_{m, \vec{k}}^i q_{\vec{k}}(t, i)}{\sum_{j=0}^{k_t} \lambda_{m, \vec{k}}^j q_{\vec{k}}(t, j)}; \quad (10)$$

that is, the corresponding terms in the sum for $g_{\vec{k}'}(\lambda_{m, \vec{k}})$ and $g_{\vec{k}}(\lambda_{m, \vec{k}})$ are the same if $t \neq t^*$.

Now consider the corresponding terms for type t^* . First observe that for all $i < k'_{t^*}$,

$$\frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}'}(t^*, i)}{\sum_{j=0}^{k'_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}'}(t^*, j)} < \frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}}(t^*, i)}{\sum_{j=0}^{k_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}}(t^*, j)}; \quad (11)$$

the two terms have essentially the same numerator (the use of $q_{\vec{k}'}$ instead of $q_{\vec{k}}$ cancels out as in Equation (10)), but the first has a larger denominator because $k'_{t^*} = k_{t^*} + 1$, so there is one more term in the sum. Since $f_{t^*} = \sum_{i=0}^{k_{t^*}} d_{q_{\vec{k}}}^*(t^*, i) = \sum_{i=0}^{k'_{t^*}} d_{q_{\vec{k}'}}^*(t^*, i)$, by Equations (2) and (3),

$$\sum_{i=0}^{k_{t^*}} \frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}}(t^*, i)}{\sum_{j=0}^{k_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}}(t^*, j)} = \sum_{i=0}^{k'_{t^*}} \frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}'}(t^*, i)}{\sum_{j=0}^{k'_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}'}(t^*, j)}. \quad (12)$$

It follows that

$$\sum_{i=0}^{k'_{t^*}} i \frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}'}(t^*, i)}{\sum_{j=0}^{k'_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}'}(t^*, j)} > \sum_{i=0}^{k_{t^*}} i \frac{f_{t^*} \lambda_{m, \vec{k}}^i q_{\vec{k}}(t^*, i)}{\sum_{j=0}^{k_{t^*}} \lambda_{m, \vec{k}}^j q_{\vec{k}}(t^*, j)}. \quad (13)$$

To see this, note that the two expressions above have the form $\sum_{i=0}^{k_{t^*}+1} ic_i$ and $\sum_{i=0}^{k_{t^*}} id_i$, respectively. By Equation (12), $\sum_{i=0}^{k_{t^*}+1} c_i = \sum_{i=0}^{k_{t^*}} d_i = f_{t^*}$; by Equation (11), $c_i < d_i$ for $i = 0, \dots, k_{t^*}$. Thus, in going from the right side to the left side, weight is being transferred from lower terms to $k_{t^*} + 1$.

Combining Equations (10) and (13) gives us $g_{\vec{k}'}(\lambda_{m, \vec{k}}) > g_{\vec{k}}(\lambda_{m, \vec{k}}) = m$, as desired. Since $g_{\vec{k}'}(\lambda_{m, \vec{k}'}) = m$, by definition, it follows that $g_{\vec{k}'}(\lambda_{m, \vec{k}}) > g_{\vec{k}'}(\lambda_{m, \vec{k}'})$. Since, as shown above, $g_{\vec{k}'}$ is an increasing function, it follows that $\lambda_{m, \vec{k}} > \lambda_{m, \vec{k}'}$. Thus, $\lambda_{m, \vec{k}}$ is decreasing in \vec{k} .

We now show that the monotonicity of $\lambda_{m, \vec{k}}$ implies the monotonicity of $p_u^{m, \vec{k}}$. To do this, we show that, for all types t , $p_u^{m, \vec{k}} = p_d \lambda_{m, \vec{k}} \omega_t$. Since ω_t and p_d are independent of m and \vec{k} , it then follows that the monotonicity of $\lambda_{m, \vec{k}}$ implies the monotonicity of $p_u^{m, \vec{k}}$. (Recall that $\omega_t = \beta_t \chi_t / \rho_t$ was defined in Section 3.)

Fix a type t' . Then, dropping superscripts and subscripts on p_u , d , and λ for brevity, we have the following sequence of equalities (where the explanation for some of these

lines is given following the equations):

$$p_u = \left(\sum_t \rho_T (f_t - d(t, 0)) \right) \left(\frac{\chi_{t'} \beta_{t'}}{n \sum_t \chi_t \beta_t (f_t - d(t, k_t))} \right) \quad (14)$$

$$= \left(\frac{\sum_t \sum_{i=1}^{k_t} \rho_t d(t, i)}{\sum_t \sum_{i=0}^{k_t-1} \chi_t \beta_t d(t, i)} \right) \left(\frac{\chi_{t'} \beta_{t'}}{n} \right) \quad (15)$$

$$= \left(\frac{\sum_t \sum_{i=0}^{k_t-1} \rho_t \lambda \omega_t d(t, i)}{\sum_t \sum_{i=0}^{k_t-1} \chi_t \beta_t d(t, i)} \right) \left(\frac{\chi_{t'} \beta_{t'}}{n} \right) \quad (16)$$

$$= \lambda \left(\frac{\sum_t \sum_{i=0}^{k_t-1} \chi_t \beta_t d(t, i)}{\sum_t \sum_{i=0}^{k_t-1} \chi_t \beta_t d(t, i)} \right) (\omega_{t'} p_d) \quad (17)$$

$$= \lambda \omega_{t'} p_d$$

Equation (14) is just the definition of p_u from Equation (6). Equation (15) follows from the observation that, by Equation (2), $f_t = \sum_i d(t, i)$. Equation (16) follows from the observation that, again by Equation (2), $d(t, i) = \omega_t \lambda d(t, i-1)$. Equation (17) follows from the definitions of ω_t and p_d (see Equation (7)). Thus, as required, $p_u^{m, \vec{k}} = p_d \lambda_{m, \vec{k}} \omega_t$.

Finally, we show that the monotonicity of $p_u^{m, \vec{k}}$ implies the monotonicity of BR_G . Let $\vec{k}'' = BR_G(\vec{k})$. By Lemma 5.1, k_t'' is the maximum value of κ such that

$$\alpha_t \leq E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u^{m, \vec{k}}, p_d)}] \gamma_t.$$

We (implicitly) defined the random variable $J(\kappa, p_u, p_d)$ as a function on histories. Instead, we can define $J(\kappa, p_u, p_d)$ as a function on random bitstrings (which intuitively determine a history). With this redefinition, it is clear that, if $p_u < p_u'$, for all bitstrings b , we have $J(\kappa, p_u, p_d)(b) < J(\kappa, p_u', p_d)(b)$. It easily follows that

$$E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u', p_d)}] < E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u, p_d)}]$$

for all κ . Thus, the monotonicity of BR_G follows from the monotonicity of $p_u^{m, \vec{k}}$. \square

LEMMA 5.4. *For all games $G = (T, \vec{f}, h, m, n)$, there exists a $\delta^* < 1$ such that if $\delta_t > \delta^*$ for all t , there is a vector \vec{k} of thresholds such that $BR_G(\vec{k}) > \vec{k}$.*

PROOF. Take \vec{k} to be such that $k_t = \lceil m \rceil + 1$ for each type t . Then by Theorem 5.2, there exists a \vec{k}' such that $BR_G(\vec{k}) = \vec{k}'$. By Lemma 5.1, k_t' is the maximum value of κ such that

$$\alpha_t \leq E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u^{\vec{k}}, p_d)}] \gamma_t. \quad (4)$$

As δ_t approaches 1, $E[(1 - (1 - \delta_t)/n)^{J(\kappa, p_u^{\vec{k}}, p_d)}]$ approaches 1, and so the right hand side of Equation (4) approaches γ_t . For any standard agent, $\alpha_t < \gamma_t$. Thus, there exists a δ_t such that

$$\alpha_t \leq E[(1 - (1 - \delta_t)/n)^{J(k_t, p_u^{\vec{k}}, p_d)}] \gamma_t.$$

For this choice of δ_t , we must have $k_t' \geq k_t + 1 > k_t$. Take $\delta^* = \max_t \delta_t$. \square

ACKNOWLEDGMENTS

We would like to thank Randy Farmer, Peter Harremoës, Shane Henderson, Jon Kleinberg, David Parkes, Dan Reeves, Emin Gün Sirer, Michael Wellman, and anonymous referees and editors for helpful suggestions,

discussions, and criticisms. Eric Friedman, Ian Kash, and Joseph Halpern are supported in part by NSF under grant ITR-0325453. Joseph Halpern is also supported in part by NSF under grants CTC-0208535, IIS-0812045, and IIS-0534064; by ONR under grant N00014-01-10-511; by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725; and by AFOSR under grants F49620-02-1-0101, FA9550-08-1-0438, FA9550-05-1-0055, and FA9550-09-1-0266. Eric Friedman is also supported in part by NSF under grant CDI-0835706.

REFERENCES

- E. Adar and B. A. Huberman. 2000. Free Riding on Gnutella. *First Monday* 5, 10 (2000).
- Sachin Adlakha and Ramesh Johari. 2010. Mean Field Equilibrium in Dynamic Games with Complementarities. In *IEEE Conference on Decision and Control (CDC)*.
- Sachin Adlakha, Ramesh Johari, Gabriel Y. Weintraub, and Andrea Goldsmith. 2010. Mean Field Analysis for Large Population Stochastic Games. In *IEEE Conference on Decision and Control (CDC)*.
- K. G. Anagnostakis and M. Greenwald. 2004. Exchange-Based Incentive Mechanisms for Peer-to-Peer File Sharing. In *International Conference on Distributed Computing Systems (ICDCS)*. 524–533.
- Christina Aperjis, Michael J. Freedman, and Ramesh Johari. 2008. Peer-assisted content distribution with prices. In *2008 ACM Conference on Emerging Network Experiment and Technology (CoNEXT)*. 17.
- Christina Aperjis and Ramesh Johari. 2006. A peer-to-peer system as an exchange economy. In *GameNets '06: Proceeding from the 2006 Workshop on Game Theory for Communications and Networks*. 10. DOI: <http://dx.doi.org/10.1145/1190195.1190206>
- Alvin AuYoung, Brent Chun, Chaki Ng, David Parkes, Amin Vahdat, and Alex Snoeren. 2007. *Practical Market-Based Resource Allocation*. Technical Report. UC San Diego. CS2007-0901.
- Mira Belenkiy, Melissa Chase, C. Christopher Erway, John Jannotti, Alptekin Küpcü, Anna Lysyanskaya, and Eric Rachlin. 2007. Making p2p accountable without losing privacy. In *2007 ACM Workshop on Privacy in the Electronic Society (WPES)*. 31–40.
- Aleksander Berentsen. 2002. On the Distribution of Money Holdings in a Random-Matching Model*. *International Economic Review* 43, 3 (2002), 945–954. DOI: <http://dx.doi.org/10.1111/1468-2354.t01-1-00042>
- Aleksander Berentsen and Guillaume Rocheteau. 2002. On the efficiency of monetary exchange: how divisibility of money matters. *Journal of Monetary Economics* 49, 8 (2002), 1621 – 1649. DOI: [http://dx.doi.org/10.1016/S0304-3932\(02\)00178-2](http://dx.doi.org/10.1016/S0304-3932(02)00178-2)
- O.J. Blanchard and S. Fischer. 1989. *Lectures on Macroeconomics*. The MIT Press.
- J. Brunelle, P. Hurst, J. Huth, L. Kang, C. Ng, D. Parkes, M. Seltzer, J. Shank, and S. Youssef. 2006. Egg: An Extensible and Economics-Inspired Open Grid Computing Platform. In *Third Workshop on Grid Economics and Business Models (GECOM)*. 140–150.
- B. Chun, P. Buonadonna, A. AuYoung, C. Ng, D. Parkes, J. Schneidman, A. Snoeren, and A. Vahdat. 2005. Mirage: A Microeconomic Resource Allocation System for SensorNet Testbeds. In *Second IEEE Workshop on Embedded Networked Sensors*. 19–28.
- Bram Cohen. 2003. Incentives Build Robustness in BitTorrent. In *First Workshop on the Economics of Peer-to-Peer Systems (P2PECON)*.
- T. Cover and J. Thomas. 1991. *Elements of Information Theory*. John Wiley & Sons, Inc., New York.
- Pranav Dandekar, Ashish Goel, Ramesh Govindan, and Ian Post. 2011. Liquidity in credit networks: a little trust goes a long way. In *Proceedings of the 12th ACM conference on Electronic commerce (EC)*. ACM, New York, NY, USA, 147–156. DOI: <http://dx.doi.org/10.1145/1993574.1993597>
- G. Ellison. 1994. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies* 61 (1994), 567–588.
- eMule Project. 2009. eMule web site. (2009). <http://www.emule-project.net/>.
- M. Feldman, K. Lai, I. Stoica, and J. Chuang. 2004. Robust incentive techniques for peer-to-peer networks. In *ACM Conference on Electronic Commerce (EC)*. 102–111.
- Eric. J. Friedman, Joseph. Y. Halpern, and Ian A. Kash. 2006. Efficiency and Nash Equilibria in a Scrip System for P2P Networks. In *Seventh ACM Conference on Electronic Commerce (EC)*. 140–149.
- E. J. Friedman and P. Resnick. 2001. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy* 10, 2 (2001), 173–199.
- Drew Fudenberg and Jean Tirole. 1991. *Game Theory*. MIT Press.
- Edward J. Green and Ruilin Zhou. 1998. A Rudimentary Random-Matching Model with Divisible Money and Prices. *Journal of Economic Theory* 81, 2 (1998), 252 – 271. DOI: <http://dx.doi.org/10.1006/jeth.1997.2356>
- Edward J. Green and Ruilin Zhou. 2002. Dynamic Monetary Equilibrium in a Random Matching Economy. *Econometrica* 70, 3 (2002), 929–969. DOI: <http://dx.doi.org/10.1111/1468-0262.00315>

- A. J. Grove, J. Y. Halpern, and D. Koller. 1994. Random Worlds and Maximum Entropy. *J. Artif. Intell. Res. (JAIR)* 2 (1994), 33–88.
- R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. 2004. Propagation of trust and distrust.. In *Conference on the World Wide Web (WWW)*. 403–412.
- M. Gupta, P. Judge, and M. H. Ammar. 2003. A reputation system for peer-to-peer networks.. In *Network and Operating System Support for Digital Audio and Video (NOSSDAV)*. 144–152.
- T. Hens, K. R. Schenk-Hoppe, and B. Vogt. 2007. The Great Capitol Hill Baby Sitting Co-op: Anecdote or Evidence for the Optimum Quantity of Money? *J. of Money, Credit and Banking* 9, 6 (2007), 1305–1333.
- M. Humbert, H. Manshaei, and J-P Hubaux. 2011. One-to-n scrip systems for cooperative privacy-enhancing technologies. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. 682–692. DOI : <http://dx.doi.org/10.1109/Allerton.2011.6120234>
- J. Ioannidis, S. Ioannidis, A. D. Keromytis, and V. Prevelakis. 2002. Fileteller: Paying and Getting Paid for File Storage.. In *Financial Cryptography*. 282–299.
- Ithaca Hours Inc. 2005. Ithaca Hours. (2005). <http://www.ithacahours.org/>.
- E. T. Jaynes. 1978. Where do we stand on maximum entropy? In *The Maximum Entropy Formalism*, R. D. Levine and M. Tribus (Eds.). MIT Press, Cambridge, Mass., 15–118.
- Kris Johnson, David Simchi-Levi, and Peng Sun. 2014. Analyzing Scrip Systems. *Operations Research* 62, 3 (2014), 524–534. DOI : <http://dx.doi.org/10.1287/opre.2014.1260>
- S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. 2003. The Eigentrust algorithm for reputation management in P2P networks.. In *Conference on the World Wide Web (WWW)*. 640–651.
- M. Kandori. 1992. Social norms and community enforcement. *Review of Economic Studies* 59 (1992), 63–80.
- Ian. A. Kash, Eric. J. Friedman, and Joseph. Y. Halpern. 2007. Optimizing Scrip Systems: Efficiency, Crashes, Hoarders and Altruists. In *Eighth ACM Conference on Electronic Commerce (EC)*. 305–315.
- Ian A. Kash, Eric J. Friedman, and Joseph Y. Halpern. 2009. Manipulating Scrip Systems: Sybils and Collusion. In *First Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*.
- Ian A. Kash, Eric J. Friedman, and Joseph Y. Halpern. 2012. Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion. *Distributed Computing* 25, 5 (2012), 335–357. DOI : <http://dx.doi.org/10.1007/s00446-012-0170-z>
- Nobuhiro Kiyotaki and Randall Wright. 1989. On Money as a Medium of Exchange. *Journal of Political Economy* 97, 4 (1989), pp. 927–954. <http://www.jstor.org/stable/1832197>
- L. Lovasz and P. Winkler. 1995. Mixing of Random Walks and Other Diffusions on a Graph. In *Surveys in Combinatorics*, P. Rowlinson (Ed.). Cambridge University Press, 119–154. citeseer.csail.mit.edu/lovasz95mixing.html
- Geoffrey Mainland, Laura Kang, Sébastien Lahaie, David C. Parkes, and Matt Welsh. 2004. Using virtual markets to program global behavior in sensor networks. In *11th ACM SIGOPS European Workshop*. 1.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, U.K.
- P. Milgrom and J. Roberts. 1990. Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica* 58, 6 (1990), 1255–1277.
- M. S. Miller and K. E. Drexler. 1988. Markets and Computation: Agoric Open Systems. In *The Ecology of Computation*, B. Huberman (Ed.). Elsevier, 133–175.
- C. Ng, P. Buonadonna, B. Chun, A. Snoeren, and A. Vahdat. 2005. Addressing Strategic Behavior in a Deployed Microeconomic Resource Allocator. In *Third Workshop on Economics of Peer-to-Peer Systems (P2PECON)*. 99–104.
- Ryan S. Peterson and Emin Gün Sirer. 2009. Antfarm: Efficient Content Distribution with Managed Swarms. In *Networked Systems Design and Implementation (NSDI)*.
- M. Piatek, T. Isdal, A. Krishnamurthy, and T. E. Anderson. 2008. One Hop Reputations for Peer to Peer File Sharing Workloads. In *5th USENIX Symposium on Networked Systems Design & Implementation (NSDI)*. 1–14.
- M. L. Puterman. 1994. *Markov Decision Processes*. Wiley.
- Rameez Rahman, David Hales, Tamas Vinko, Johan A. Pouwelse, and Henk J. Sips. 2010. No more crash or crunch: Sustainable credit dynamics in a P2P community. In *P2010 International Conference on High Performance Computing & Simulation (HPCS)*. 332–340.
- D. M. Reeves, B. M. Soule, and T. Kasturi. 2007. Yootopia! *ACM SIGecom Exchanges* 6, 2 (2007), 1–26.
- S. I. Resnick. 1992. *Adventures in Stochastic Processes*. Birkhauser.
- Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259.

- Michael Sirivianos, Jong Han Park, Xiaowei Yang, and Stanislaw Jarecki. 2007. Dandelion: Cooperative Content Distribution with Robust Incentives. In *2007 USENIX Annual Technical Conference*. 157–170.
- M. Stonebraker, P. M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Stelin, and A. Yu. 1996. Mariposa: a wide-area distributed database system. *The VLDB Journal* 5, 1 (1996), 48–63.
- I. E. Sutherland. 1968. A Futures Market in Computer Time. *Commun. ACM* 11, 6 (1968), 449–451.
- A. Tarski. 1955. A Lattice-Theoretical Fixpoint Theorem and Its Applications. *Pacific J. Math.* 5 (1955), 285–309.
- R. H. Timberlake. 1987. Private Production of Scrip-Money in the Isolated Community. *J. of Money, Credit and Banking* 19, 4 (1987), 437–447.
- D. M. Topkis. 1979. Equilibrium Points in Nonzero-Sum n-Person Submodular Games. *Siam Journal of Control and Optimization* 17 (1979), 773–787.
- Mihaela van der Schaar, Jie Xu, and William Zame. 2013. Efficient online exchange via fiat money. *Economic Theory* 54, 2 (2013), 211–248. DOI: <http://dx.doi.org/10.1007/s00199-013-0744-4>
- V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. 2003. KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing. In *First Workshop on Economics of Peer-to-Peer Systems (P2PECON)*.
- Washington Metropolitan Area Transit Commission. 1970. Scrip System of the D.C. Transit System. (1970).
- Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. 2011. Industry dynamics: Foundations for models with an infinite number of firms. *Journal of Economic Theory* 146, 5 (2011), 1965 – 1994. DOI: <http://dx.doi.org/10.1016/j.jet.2011.05.007>
- Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. 2008. Markov Perfect Industry Dynamics With Many Firms. *Econometrica* 76, 6 (2008), 1375–1411. DOI: <http://dx.doi.org/10.3982/ECTA6158>
- Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. 2010. Computational Methods for Oblivious Equilibrium. *Operations Research* 58, 4-Part-2 (2010), 1247–1265. DOI: <http://dx.doi.org/10.1287/opre.1090.0790>
- L. Xiong and L. Liu. 2002. Building Trust in Decentralized Peer-to-Peer Electronic Communities. In *International Conference on Electronic Commerce Research (ICECR)*.