

RESOURCE SHARING ON ENDOGENOUS NETWORKS

PHILIP SOLIMINE¹ AND LUKE BOOSEY²

[Click here for most recent version]

ABSTRACT. We examine behavior in a voluntary resource sharing game that incorporates endogenous network formation; an incentive problem that is increasingly common in contemporary digital economies. By varying the information structure in a controlled laboratory experiment, we examine the underlying mechanisms of reciprocity that generate emergent patterns in linking and contribution decisions. Specifically, we vary whether players are given information about which other players in their group are sharing with them. Reduced-form estimates find significant effects of this information treatment on a number of key outcomes such as efficiency and balanced decentralization. To further understand the driving causes of these observed changes in behavior, we develop and estimate a discrete-choice framework, using computationally efficient panel methods to identify the structure of social preferences in this setting. We find that subjects react to this new information by focusing reciprocity, which helps players coordinate to reach mutually beneficial outcomes. We also find that direct and indirect reciprocity are not perfect substitutes, and this interaction further helps players support efficient outcomes.

1. INTRODUCTION

The digital age has played host to the rapid development of massive sharing economies. From social media platforms to open-source software and even down to the microstructure of the World Wide Web, our entire infrastructure often relies heavily on voluntary interaction for its prosperity. Digital platforms permeate nearly every aspect of modern life, from online shopping to social

Date: November 13, 2021.

The authors would like to thank Christopher Brown, Krishna Dasaratha, Matthew Gentry, Ben Golub, Mark Isaac, Angelo Mele, Luca Paolo Merlino, Anke Meyer-Baese, Svetlana Pevnitskaya, Jonathan Stewart, and Cynthia Fan Yang for their advice and feedback, valuable discussions at various stages of production, and/or comments on earlier drafts. Thanks also to participants at the Economic Science Association Job Market Candidates Seminar, Conference of Network Science in Economics, Conference(s) of the Southern Economic Association, NetSci satellite on Statistical Physics of Financial and Economic Networks, XS/FS and Quantitative Methods working groups at FSU, and the Florida State University Computational Xposition. Finally, we are grateful to the Charles & Persis Rockwood Fellowship, the L. Charles Hilton Center for Economic Prosperity and Individual Opportunity, and the FSU College of Social Sciences & Public Policy for their generous support. This paper has previously been presented under the title “Strategic formation of cooperative networks”.

¹ *Job market paper. Corresponding author:* psolimine@fsu.edu. Department of Economics and Department of Scientific Computing, Florida State University, Tallahassee, FL.

² Department of Economics, Florida State University, Tallahassee, FL.

media, travel, peer-to-peer payments, career advancement, entertainment, education, and beyond. The ability to form networks instantaneously has led to a dramatic shift in the way we interact. But too little is known about how these platforms affect our behavior and social organization. There is thus a clear need to develop our understanding of the motives that power this social interaction and their resulting effects on the structure of social systems. A pervasive feature of these environments, in both digital and analog settings, is that an individual who invests in generating externalities can often choose where to direct them. This pattern of interaction generates a network whose structure can be informative of the process that governs its formation and evolution. Understanding how reputation and information affect the evolutionary behavior of interconnected social systems can help inform platform designers and policymakers on ways to design digital platforms that promote efficiency and generate positive change.

In this paper, we study a voluntary resource sharing game in which players choose how much, and with whom, to share. Shared resources generate externalities for other players, but the player sharing the resources does not derive any immediate or direct benefit from doing so. The decisions about who to share with can be viewed as a process of endogenous network formation. A number of relevant real-world examples fit into the basic framework of voluntary resource sharing. For instance, consider web page linking decisions; a series of web pages are paid advertising revenue for traffic, and they are indexed by a search engine which favors well-linked pages.¹ Each web page (the source) chooses a set of other pages to link to, and chooses how prominently to display this set of hyperlinks. When a user clicks these links, they are redirected away from the source, diverting a fraction of their traffic to another site. In this case, linking to another website has an explicit cost but no explicit or direct monetary benefit for the source page. In addition to the diverted traffic, the linked page receives a boost in search results due to the incoming link. For each additional link chosen by the source, however, more traffic is lost to the source, and each linked page receives a smaller portion of the diverted traffic.

For a second example, consider a group of friends that engage in password sharing for a collection of streaming services. Each has the option to share their password with their friends, providing them with access to the service – however, the platforms in question have strict limits how many users can access the platform concurrently from a single account. This environment has recently been studied in detail by Gerke et al. (2019). In this case, the more users that share an account,

¹See, e.g. Page et al. (1999)

the higher the probability of disruption due to concurrent usage. In this example, the platform sets the level of congestion by limiting the number of concurrent users; sharing is essentially free if a sufficient number of screens are allowed and there is no risk of disruption, but becomes costly as simultaneous usage is limited. Finally, consider the time allocation problem faced by a researcher involved in multiple projects with different sets of coauthors. The researcher has a limited amount of time and concentration power to dedicate towards coauthored projects and her own research activity. Allocating attention to coauthored projects benefits the coauthors, but effective contributions towards multiple projects reduces the impact on each.

While the examples above are distinct from each other in several ways, they each highlight some key common characteristics of voluntary sharing environments. First, the decision to share naturally entails some cost to the sharer. Second, shared resources are, in many cases, congestible in the sense that sharing the same resources across multiple recipients may reduce the effective benefit derived by each recipient. Third, sharing is conventionally associated with altruistic or reciprocal motivations – for example, webpage linking decisions may be forged with anticipation that links will be reciprocated; subscribers may share their passwords with a few friends in exchange for reciprocal sharing of their passwords to other streaming services; coauthors may be more inclined to dedicate their time to a project if their coauthors are similarly dedicated.

In such settings, reciprocity plays a key role in helping players coordinate to reach efficient outcomes. People tend to enjoy helping others who have helped them, and they can do this through either direct or indirect channels. By controlling the specific structure and content of the information a player has about their neighbors' actions, we can change how they reciprocate. Here, we are specifically interested in two types of reciprocity and how these types interact. The first is generalized or indirect reciprocity. Generalized reciprocity can be thought of as a “pay it forward effect”; it refers to a setting in which players who receive a high value of incoming benefits from others are then encouraged to share with the community, in a way that is not specifically aimed or targeted only at their benefactors. This contrasts with direct reciprocity, in which reciprocal behavior is targeted specifically toward others who a player knows has shared with them in the past. Given a voluntary contributions mechanism on an exogenous structure, it would be difficult to distinguish these two forms of reciprocity, and even more difficult to determine how these social preferences interact with each other. By varying the structure and content of information that each player receives about the behavior of their neighbors in a controlled laboratory setting, we

can enable or disable each of these forms of reciprocity. By observing people’s behavior in this setting, and looking at the patterns of interaction that they choose, we can identify how these social preferences actually interact.

In order to examine the underlying influences on voluntary resource sharing, we designed an experiment to study a simple sharing environment that incorporates these three key features. In the experiment, subjects were randomly assigned to groups of four players. Groups interacted with each other for 15 rounds (i.e., without random rematching between rounds). In each round, players privately and simultaneously made two choices; (i) how much of their endowment (20 tokens) to share (referred to as their contribution), and (ii) with whom to share, forming a network. Players retained any endowment not shared. Player i ’s contribution (her shared resources) generated benefits for her selected recipients and for player i herself; however, the shared resources were congestible, so that each received only some fraction of a token for each contributed token shared by player i . More specifically, every contributed token was multiplied by a factor of 1.6, with the resulting proceeds shared equally between player i and her selected recipients.² At the end of the round, player i ’s payoff was calculated as the sum of her retained endowment, the return on her own shared resources, and the incoming benefits derived from others who selected player i as a recipient of their own shared resources.

We examine two treatment conditions in the experiment, designed to identify the impact of the information structure on reciprocal behavior. The first is a control treatment in which players learn the total inflow of benefits from others after each round, but cannot identify the source of those benefits. In this way, direct reciprocity is precluded, although sharing behavior may nevertheless be driven in part by generalized (or indirect) reciprocity. The second condition is an information treatment, in which players are informed about the incoming benefit from each of the other three group members. As such, they can identify who shared with them, enabling direct reciprocity to influence both the contribution and network formation decisions. In each session of the experiment, groups first played 15 rounds of the control treatment. After completing this first part, they were informed that they would play another 15 rounds with the same group (albeit with shuffled IDs). In three sessions, the second part simply repeated the control treatment condition. In the other five sessions, we implemented the information treatment in the second part.

²This specification of the benefits from shared resources is, by design, comparable to the widely adopted implementation of the *marginal per capita return* (MPCR) in VCM or linear public goods experiments.

By virtue of the endogenous linking decisions, our study naturally contributes to the growing literature on network formation games, and especially to more recent work examining simultaneous action and link formation decisions. Network formation problems are notoriously difficult to analyze, but have far-reaching consequences and applications in economics and related fields alike.³ Fundamentally, the observation of a panel of small networks over time, rather than that of a single large cross-section, allows us to adopt a more sophisticated and innovative structural approach. These methods build on the recent techniques developed by Badev (2021) and Mele (2017) for large cross-sectional observations of social and economic networks. More specifically, panel data allows us to estimate the discrete choice process directly, avoiding the computational issues associated with large-network cross-sectional estimation. As such, one of our contributions is methodological in nature – we implement a new approach (see, e.g., Overgoor et al., 2019; Gupta and Porter, 2020) to estimate the network formation game as a combination of individual evolutionary discrete choice strategies.

The voluntary sharing environment exhibits some similarities with other well-studied strategic decision settings, including dictator giving, the provision of public goods or club goods, and public goods games played on networks. As such, our study is closely related to the literature examining network public goods games (Bramoullé et al., 2007; Elliott and Golub, 2019; Boosey, 2017), and especially to work that extends the game to incorporate endogenous linking decisions (Galeotti and Goyal, 2010; Kinaterder and Merlino, 2017, 2021). Importantly, while existing literature has focused on settings in which links describe the decision to access another player’s provision of the good, our environment captures the reverse situation – the links formed by each player describe the decision to *grant access* to the benefits generated by her own contribution.

In Section 2, we walk through the relevant literature. We first discuss how our work is tied to the existing literature on network public goods games (with exogenous or endogenous networks). Then, we relate our empirical approach to the recent literature on discrete choice modeling and the estimation of network formation models. This literature has focused predominantly on large-network cross-sectional estimation. The introduction of experimental data, however, allows us to carefully control group size and monitor progress over time; generating a panel of networks from which estimation is particularly tractable and relies on fewer assumptions.

³See Chandrasekhar (2016) for a comprehensive introduction to the study of network formation.

Section 3 contains a brief overview of the theoretical framework that describes this decision setting. Notably, the game is characterized by a free-rider problem – it admits a unique stage-game Nash equilibrium of no investment and an empty network with no sharing. While this result is straightforward to obtain, it does not accurately describe the behavior we observe. This is reminiscent of the robust findings in classical linear public goods experiments that contributions depart significantly from the free-riding prediction. These findings have spurred a rich literature examining alternative models of conditional cooperation, usually by incorporating some form of social preferences, inequality aversion, or reciprocity. In a similar fashion, we introduce behavioral preferences based on forms of reciprocity that are allowed by the information structure, and discuss stability concepts for networks under these behavioral preferences. While stable networks are a fixed point of the dynamics, convergence results do not necessarily obtain when players make simultaneous moves. However, we can empirically investigate convergence. Part of the benefit to the panel approach to modeling network formation is that preferences no longer need fit into the strict form required to aggregate to a potential game, since we have the benefit of actually witnessing individual strategic evolution in a repeated choice setting.

In Section 6, we describe panel estimators for network formation modeling. We are among the first to adopt this approach, perhaps because of the difficulty that is associated with collecting data on repeated choice in replicated groups. The Maximum Likelihood Estimator (MLE) can be tractable if the group size is limited, but computation time still grows exponentially with the size of the network. As the size of each network cross-section grows larger, the Maximum Pseudo-Likelihood Estimator (MPLE) can provide an effective alternative. This approach works by replacing the likelihood function with the product of conditional distributions. Under certain conditions, the MPLE can produce consistent estimates for large networks, and its required computation time grows only quadratically with the network size; offering an exponential speedup relative to the MLE. We also discuss quantification of uncertainty in this setting. Clustered standard errors can be produced based on partial likelihood methods, but these may understate the variability of estimates in data of limited size. Particularly for pseudolikelihood methods, standard errors based on the inverse Hessian matrix are known to underestimate true variability, sometimes drastically. We discuss (semi-)parametric and fully nonparametric bootstrap methods designed to produce cluster robust standard errors that work for the MPLE and more accurately reproduce the true variability of the estimators in the case of limited data.

Section 4 describes the design of the laboratory experiment. In the experiment, we fixed the parameters of the environment and varied the information that was shown to the players. In the baseline, players can only observe the total amount that was shared with them at the end of each round. In treatment sessions, we give players more detailed information about the behavior of the rest of their group – specifically, we inform them of how much was shared with them by each other member of their group.

In Section 5 we discuss the results of this series of experiments. We first examine the effects of this treatment on key outcomes such as efficiency and complementarity. We find that players use information effectively, and even this simple mechanism for reputation building and trust manages to support highly efficient and reciprocal outcomes and avoid free-riding. These reduced-form results are contained in Section 5.1. In Section 6.3, we implement structural estimators to characterize the nature of the tradeoff between generalized and direct reciprocity. Notably, we find that generalized reciprocity appears to persist even after the intervention; players who received more benefits in the previous round tend to share more, even if it is not specifically directed at those who shared with them. Because we are able to strictly control group size in the experiment, we are able to estimate both the MLE and MPLE on experimental data. We find that the MPLE estimates from these panels of small networks are indeed biased. However, even in these small network panels, the MPLE produces estimates that are qualitatively similar to the MLE, pointing to its usefulness in cases where the MLE might not be tractable. Using estimates from the MLE, we characterize which stable network structures are frequently observed.

We then introduce individual heterogeneity into the model, using characteristics generated by questionnaires in the experiment designed to elicit measures of individuals' trust and their preference for reciprocity. These characteristics seem to play a significant role in decision making. Under the more conservative nonparametric bootstrap, however, it appears that more data would be required to precisely reconstruct the impact of observed heterogeneity. Notably however, average estimates of behavioral effects do not appear substantially biased by the absence of individual heterogeneity. Finally, in Section 7, we summarize the results and conclude.

2. RELEVANT LITERATURE

There has been prior work examining both the extension of public goods to static (exogenous) networks, and the provision of public goods on endogenous networks. In particular, Bramoullé

et al. (2007) launched research into this environment, by showing that given a network shape, specialized Nash equilibria (in which a small subset of players provide effort, and those connected to them free-ride) tend to be not only stable but also the most efficient in terms of overall welfare. Further, they find that, in an environment with agent heterogeneity, these specialized equilibria are often unique. This is especially relevant since, in our endogenous network environment, an agent's marginal cost of effort is dependent on their position in the network, and in particular how many others they choose to share with.

Elliott and Golub (2019) characterize outcomes in public goods games on exogenous networks by the spectrum of a matrix called the benefits matrix, in which each entry gives the marginal rate of substitution between decreasing own contribution and increased benefits from a neighbor in a fixed network. Their results tie the existence of Pareto-efficient outcomes to the spectral radius of the benefits matrix, and characterize the Lindahl outcomes as those with effort proportional to each individual's eigenvector centrality in the graph described by the benefits matrix. Many of their results rely on the connectedness of the benefit graph, which is not guaranteed in random graphs. In the case of endogenous formation this assumption may not be satisfied, rendering spectral methods difficult to implement outside of particular special cases involving links that are fully independent (e.g. Dasaratha, 2020).

Finally, in the exogenous/fixed network case, Boosey (2017) uses data from a laboratory experiment to examine the mechanisms for cooperation in a repeated network public goods game. Experimental results showed a significant portion of subjects playing strategies of *conditional cooperation*, in which subjects play strategies which react strongly to the behavior of their neighbors in previous rounds. We will incorporate this phenomenon into our structural model, by placing strategies on an evolutionary spectrum from reactionary to predictive. When playing a purely reactive strategy under bounded rationality, simultaneous play may not converge to a stage-game equilibrium (Alós-Ferrer and Netzer, 2010; Hommes and Ochea, 2012).

Prior extensions of public goods to environments with endogenous linking include Galeotti and Goyal (2010), which furthers the specialization result of Bramoullé et al. (2007). In a setting where individuals choose *incoming* links (or, in other words, pay a fixed cost to receive benefits from the contributions of another individual), they find that equilibria in this environment correspond to a *core-periphery* architecture, in which a small number of individuals (the *core*) provide costly effort, and a larger number of players (the *periphery*) link exclusively to the core, and free-ride off their

effort. Kinatered and Merlino (2017) expand on this result, showing how heterogeneity in costs and benefits to effort in such an environment generates equilibrium networks which take the form of a tiered multipartite graph. In the heterogeneous case, individuals with comparative advantage in production costs bear most of the brunt of contribution, while peripheral players, corresponding to those with higher valuations of the good, pay the linking costs associated with linking to the core (or inner levels of the multipartite network).

Our model departs from the existing literature on public goods in endogenous networks in a number of ways. Primarily, we model a situation in which individuals choose others with whom they would like to share the externalities generated by their resources. This is the reverse of the situations studied by Galeotti and Goyal (2010) and Kinatered and Merlino (2017), in that individuals choose the *outgoing* direction of their externalities, rather than the incoming direction of others' externalities. This describes a different scenario entirely. Most notably, the cost of linking in this environment is explicitly tied to the effort or contribution level, and can be flexibly specified to represent pure or impure (congestive) externalities, or even costs which decrease with the number of links formed. Also in this voluntary sharing environment, there is a unique stage-game Nash equilibrium of no contributions and no linking. This is, however, in stark contrast to what we actually observe in the laboratory implementation, and provides a rich environment to identify and analyze the structure of social preferences.

In addition to the novelty of the sharing environment, we also contribute to the growing literature on estimation in network formation games; a problem which is notoriously difficult to analyze, but with important applications (Chandrasekhar, 2016). By observing a panel of small networks over time, rather than that of a single large cross-section, we are able to estimate the discrete choice process directly, reducing the set of computational issues typically encountered in large-network cross-sectional estimation. Estimation of a network formation game as a combination of individual evolutionary discrete choice strategies is a new approach, which has begun to be examined only very recently by Overgoor et al. (2019) and Gupta and Porter (2020).

Hiller (2020) presents a similar model of network formation with a continuous effort choice. In his paper, while there are congestion costs associated with linking highly, these congestion costs are framed as fixed costs associated with additional links rather than as an increased marginal cost of effort. In addition, he assumes that effort levels are strategic complements, which is true in our model under reciprocity. Similar to Hiller's framework, we allow for the possibility of global substitution

effects if generalized reciprocity becomes negative under the treatment condition. However, our evidence points instead to persistent generalized reciprocity, representing global complementarity. Our model thus allows for the specification of a broader scope of social preferences.

Dasaratha (2019) presents a similar model, with specific application to innovation in industrial networks. Specifically, they model the behavior of firms choosing to invest in research and development, choosing a level of openness for their own research, and learning from other firms about their innovations. This is characterized by a multiplicative interaction rate, in which reciprocal complementarity is generated by an interaction between firms' level of sharing and openness. They study the asymptotic behavior of this theoretical model as the network grows large, and find similar results to ours; efficient outcomes are characterized by the emergence of a giant component. Similarly in our model, a network is only capable of reaching efficient outcomes if all players participate by sharing with at least one other player. Our model differs in the endogeneity of returns to contribution – we allow for the cost of sharing to vary based on a player's chosen number of neighbors, and can then describe which networks are capable of supporting efficient outcomes based on the scaling and curvature of per-capita returns to investment.

Cross-sectional network formation estimators work by making assumptions on the meeting process and dynamics that guarantee convergence to a stochastically stable stationary distribution, also called a Quantal Response Equilibrium (QRE). While the QRE (McKelvey and Palfrey, 1995, 1998) is a fixed point stationary distribution of the logit-response (logit best-reply) dynamics, in the case of simultaneous revision opportunities this fixed point is potentially unstable. This means that play may instead exhibit a Hopf bifurcation and converge to a limit cycle or stable orbit, rather than to the fixed point QRE distribution (Alós-Ferrer and Netzer, 2010). Estimating the individual strategies, however, rather than imposing stability and estimating the QRE, allows us to comment on the convergence of the calibrated system, as well as to draw from the steady-state distribution under arbitrary specifications of the revision opportunity (that is, since we estimate the utility parameters directly). This means that we can simulate draws of the steady-state distribution for large networks under sequential-move individual revisions, using the methods of Badev (2021), the first paper to examine identification in discrete-choice games taking place on endogenous networks – in which agents choose both a set of links and an action or investment level. The method used is closely related to the one used by Mele (2017) to estimate structural parameters in such a setting,

assuming that the network has converged to its steady-state distribution after sufficient iterations of an individual revision process.

3. THEORY

We can represent player decisions in a single period by weighted, directed networks with implicit self-links. Consider a set of players (agents) $N = \{1 \dots n\}$. An individual agent in this set is denoted as i . Players have a resource constraint of ω_i units, and simultaneously choose a contribution level c_i from the interval $[0, \omega_i]$. Thus $(\omega_i - c_i)$ are kept at no cost, and used to produce one unit of value for player i .

In addition to choosing a contribution level, players choose a subset N_i of the other players, with whom they would like to split the production cost and benefits. The cost of contribution and the magnitude of externalities are flexibly specified by a *Marginal Per Capita Return* (MPCR) function $m_i(N_i)$, which maps the player's chosen neighborhood to a cost-externality structure. In this way, the marginal cost of contributing is $(1 - m_i(N_i))$. For example, in the case of homogeneity in costs and congestive externalities, which will be a focal point in this study, we could set $m_i(N_i) = m(N_i) = \frac{1.6}{|N_i|}$ for $|N_i| > 0$. Without loss of generality, we fix $m_i(\emptyset) = 1$. Because we are primarily interested in examining behavior in situations where sharing is not apparently rational under standard, self-interested preferences, we operate under the following assumption:

Assumption 1. $m_i(N_i) = 1$ for the empty neighborhood $N_i = \emptyset$, and $m_i(N_i) < 1$ for all nonempty neighborhoods $N_i \neq \emptyset$.

This assumption guarantees that players cannot generate efficiency gains without sharing with another player. The outcome of the game in any round can be summarized by an $n \times n$ *adjacency matrix* A , with the elements $A_{ij} = 1$ if player j shares with player i , and $A_{ij} = 0$ otherwise. Since the benefits of investment are divided between the contributor and their neighborhood, we will fix $A_{ii} = 1$. We will denote A_i as the i^{th} column of A with the i^{th} element removed. Thus A_i is an $(n - 1)$ -vector from $\{0, 1\}^{(n-1)}$ that contains all the information of the neighborhood N_i of a node i . In addition, we collect the contributions c_i into a single vector c . By slight abuse of notation, we will overload the marginal per-capita return function $m_i(\cdot)$ to accept arguments in the form A_i , with $m_i(A_i) \triangleq m_i(\sum_{j=1}^{n-1} (A_i)_j) = m_i(|N_i|)$.

Then we can write players' concrete (observable) monetary payoffs in the following form:

$$(1) \quad \pi_i(A, c) = \pi_i(A_i, c_i; A_{-i}, c_{-i}) = \left(\sum_{j=1}^n A_{ij} c_j m_j(A_j) \right) - c_i$$

Importantly, the summation includes $j = i$ because i shares in the benefits of their own investment, although in our cases of interest, they cannot earn positive net return on their own contribution. Incentives can be decomposed into a combination of effort cost and externalities from others' contributions. That is, we can write π_i as:

$$(2) \quad \pi_i(A, c) = (m_i(A_i) - 1)c_i + \kappa_i(A_{-i}, c_{-i})$$

where $\kappa_i(A_{-i}, c_{-i}) = \sum_{j=1, j \neq i}^n A_{ij} c_j m_j(A_j)$ models pure externalities, which do not affect the decision problem of agent i on the margin, since they cannot be controlled. It is straightforward to show, by backward induction (and using Assumption 1), that the following proposition holds:

Proposition 1. *The unique Nash equilibrium of the sharing game (under Assumption 1) is $c_i = 0$ and $N_i = \emptyset$.*

Proof. Since $\kappa(A_{-i}, c_{-i})$ is independent of both c_i and A_i , this term is dropped from the marginal decision problem when taking first order conditions. We have:

$$(3) \quad \frac{\partial \pi_i}{\partial c_i} = m_i(A_i) - 1 < 0$$

by Assumption 1. Thus in the last-period stage game, the optimal contribution for any linking pattern is $c_i = 0$ for all i . From there, play unravels and no investment should be made in sharing. \square

The free-riding hypothesis is a prominently studied feature of the voluntary contributions mechanism. Robust evidence from both the lab and field, however, shows that while this *is* the unique equilibrium, realistic play very rarely agrees with this theoretical result of full free-riding (Isaac and Walker, 1988; Fisher et al., 1995).

Explanations for the failure of the free-riding hypothesis typically fall into two categories – bounded rationality and behavioral or social preferences. Bounded rationality models assert that players make errors in their computation of the optimal strategy. Social preferences models, on the other hand, aim to explain systematic deviations by asserting that the strategies are, in fact, rational

and optimal when considering other non-monetary concerns of the players – such as preference for reciprocity or fairness concerns – which bias actions toward prosocial behavior. We describe the mean utility as allowing flexibly specified social preferences:

$$(4) \quad u_i(A, c|\theta) = \theta_1 \pi_i(A, c) + \beta_i(A, c|\theta)$$

where β_i is a general function specifying the social preferences. Since the externality term κ is entirely independent of the action choice of player i in period t , we can aid computation by truncating utility to form *individual potentials*:

$$(5) \quad \phi_i(A, c|\theta) = \theta_1 (\pi_i(A, c) - \kappa_i(A_{-i}, c_{-i})) + \beta_i(A, c|\theta)$$

An effective alternative solution concept should incorporate behavioral preferences, taking θ , which is a vector of parameters, as given. We present the following stability concept referred to as behavioral Nash stability:

Definition 1. A state of the network and contribution profiles is *behaviorally Nash stable* if the following conditions hold for all players i :

- (1) $\phi_i(A_i, c_i|A_{-i}, c_{-i}, \theta) \geq \phi_i(A'_i, c_i|A_{-i}, c_{-i}, \theta) \quad \forall A'_i \in \{0, 1\}^{n-1}$
- (2) One of the following:
 - $\frac{\partial}{\partial c_i} \beta_i(A, c|\theta) = \theta_1(1 - m_i(A_i))$
 - $c_i = \omega_i$ and $\frac{\partial}{\partial c_i} \phi_i(A, c|\theta) > 0$
 - $c_i = 0$ and $\frac{\partial}{\partial c_i} \phi_i(A, c|\theta) < 0$

This equilibrium concept is similar to those used in prior literature on games on endogenous networks (Golub and Sadler, 2021), and in particular the notion of “k-stability” introduced by Badev (2021). Since players can simultaneously and unilaterally update any subset of their links, there is no need to consider “pairwise stability” (Jackson and Wolinsky, 1996) as such. The conditions in Definition 1 are straightforward from a game theoretic perspective. Conditions 1 guarantees that no player could unilaterally gain by changing their link strategy. Condition 2 stipulates that the contribution profile must be a Nash equilibrium of the game induced by the behavioral preferences and the network A .

It may be useful to define a weaker notion of stability to later empirically investigate convergence. The following describes a network that is *topologically stable*.

Definition 2. A network configuration is *topologically stable* under the contribution profile c if it satisfies condition (1) from Definition 1.

That is, the network is topologically stable if each player’s linking decisions are best responses, conditional on the contribution profile. Behavioral Nash stability is thus a refinement of topological stability, since all behaviorally Nash stable networks are necessarily topologically stable. However, topological stability does not require that the contribution profile is a Nash equilibrium of the game induced by the network.

It is also straightforward to characterize what is meant by an efficient outcome. When discussing efficiency, we refer only to the direct monetary outcomes of gameplay.

Definition 3. An outcome (A, c) is *efficient* if $\sum_{i=1}^n \pi_i(A, c) \geq \sum_{i=1}^n \pi_i(A', c')$ for all $(A', c') \in \mathcal{A} \times \mathcal{C}$

In other words, an outcome is efficient if it generates the maximum possible monetary gains for the players. Finally, consider the following classification of congestion effects.

Definition 4. The decision setting is called

- *Purely congestive* if $m_i(|N_i|) = \frac{k}{|N_i|+1}$ for $|N_i| > 0$ and $k \in \mathbb{R}$ with $1 < k < 2$
- *Subcongestive* if $m_i(|N_i|) = \frac{k(|N_i|)}{|N_i|+1}$ for $|N_i| > 0$ and $k : \mathbb{Z}_n \setminus \{0\} \rightarrow \mathbb{R}$ is a decreasing function with $1 < k(1) < 2$
- *Supercongestive* if $m_i(|N_i|) = \frac{k(|N_i|)}{|N_i|+1}$ for $|N_i| > 0$ and $k : \mathbb{Z}_n \setminus \{0\} \rightarrow \mathbb{R}$ is an increasing function with $1 < k(m) < m$ for all $m \in \mathbb{Z}_n \setminus \{0\}$

This characterization is inspired by the literature on group size effects in voluntary contributions environments (e.g. Isaac and Walker, 1988). Pure congestion describes a situation in which players generate fixed efficiency gains by sharing, and gains from their investment are divided between the sharing player and their neighbors of choice. Subcongestivity and supercongestivity refer to settings in which the marginal per-capita return decreases faster or slower with the addition of a link than the purely congestive return. Supercongestivity also contains an interesting special case, which might be termed *anticongestive*, in which the good exhibits a specific type of “network effect” – that is, sharing becomes more effective as the group size increases.

This special case lends itself to a convenient characterization of efficient structure, based on how a player’s MPCR scales with their out-degree. That is, the efficient outcome depends on the curvature of the MPCR function. It transitions between these phases when the game is purely

congestive, at the boundary between subcongestivity and supercongestivity. Along this boundary, any network in which every player is participating can support efficient outcomes.

Proposition 2. *A network is efficient in a:*

- *Purely congestive game if and only if $|N_i| \geq 1$ and $c_i = \omega_i$ for all players i .*
- *Subcongestive game if and only if $|N_i| = 1$ and $c_i = \omega_i$ for all players i .*
- *Supercongestive game if and only if the network is dense and $c_i = \omega_i$ for all players i .*

Proof. In a purely congestive setting, efficiency gains are characterized by the multiplicative constant k . More specifically, since each player’s contribution is multiplied by k and divided equally among themselves and their neighborhood, we have $\sum_{i=1}^n \pi_i(A, c) = \sum_{i=1}^n \left(\frac{k}{|N_i|+1} c_i + \sum_{j=1}^n A_{ji} \frac{k}{|N_j|+1} c_j \right) = \sum_{i=1}^n k c_i = k \sum_{i=1}^n c_i$, as long as each player forms at least one link. This efficiency is linear and strictly increasing in c_i for all i , because the definition of pure congestion prescribes $k > 1$. Therefore the sum of all monetary payoffs is strictly increasing in contributions, meaning that an efficient configuration must be characterized by full contribution.

When the setting is sub- or super-congestive, $\sum_{i=1}^n \pi_i(A, c) = \sum_{i=1}^n k(|N_i|)c_i$. Since both $k(\cdot)$ and c_i are restricted to be positive, the function is linear and increasing in both terms. Thus, efficiency will prescribe that the player choose the network that generates the largest value of $k(\cdot)$. If $k(\cdot)$ is monotone, this means choosing either the smallest (if the game is subcongestive) or largest (if the game is supercongestive) possible number of links. Finally, Assumption 1 guarantees that $k(0) = 1$ (and thus that $m(0) = m(\emptyset) = 1$). Since $k(\cdot) > 1$ for nonzero arguments in either subcongestive or supercongestive games, any fully efficient structure in these settings can never include an isolated node. □

For the experiment, we turn our focus to the purely congestive case, using the number of free-riders in a group to describe an efficient structure.

4. EXPERIMENTAL DESIGN

In this section, we describe the design and procedures of the laboratory experiment in greater detail. The experiment was conducted using undergraduate students in the XS/FS Experimental Social Sciences Laboratory at the Florida State University. We collected data from a total of 184 subjects across eight sessions. Subjects were recruited using ORSEE (Greiner, 2015), and played a

computerized version of the game programmed using zTree (Fischbacher, 2007). Instructions used in the experiment, including screenshots of the decision screens, are contained in Appendix A.

Subjects played a repeated version of a purely congestive resource sharing game, which consisted of a fixed group size of $n = 4$, and a homogenous cost/externality structure defined by an MPCR function of $m_i(N_i) = m(N_i) = \frac{1.6}{|N_i|}$. Subjects were informed at the beginning of the session that there would be two parts to the experiment, but were not informed about any details of the second part until the first part was completed. Each subject in a group was assigned a unique ID (1–4). In the first part, which consisted of 15 rounds, no information was shown to subjects regarding the individual decisions made by others in their group. Instead, subjects were simply shown their own payoff between rounds. At the end of round 15, subjects were redirected to a waiting screen, at which point the instructions for the second part were read.

The treatment variation was implemented in the second part. In three baseline sessions, consisting of a total of 72 subjects in 18 groups, subjects were told that the second part of the experiment would be exactly the same as the first part, except that subject IDs would be randomly reassigned. In five treatment sessions, consisting of 112 subjects across 28 groups, subjects were also told that they would play the game for another 15 rounds. However, in addition to reassigning ID's, subjects were also told that they would be shown how much benefit they received in the previous round from each other subject in their group. Although providing subjects with information about the past behavior of others in their group does not change the unique Nash equilibrium, this information treatment facilitates direct reciprocity where the baseline sessions do not. After the two main parts of the experiment were finished, subjects completed a series of questionnaires designed to elicit behavioral characteristics. Questions from this section are shown in Appendix B. Sessions lasted no longer than an hour. At the end of the session, subjects were paid privately by check, earning an average of \$16.69, including a \$10 show-up fee.

We use evidence from our laboratory experiment to explore information design for this environment and to compare the MLE and MPLE panel estimators. While providing subjects information about the past behavior of others in their group does not change the unique Nash equilibrium, we expect that it will have a nontrivial impact on efficiency and performance in the market. We estimate the effects of the treatment on a number of outcome measures, and use structural methods to investigate the root causes of these changes in behavior.


5. RESULTS

In this section, we present and discuss the results of the experiments and the performance of the various estimators. In Section 5.1 we highlight several aggregate, reduced-form results showing the treatment effects of information provision on key outcomes of interest. Then in Section 6.3, we implement panel structural estimators of network formation in order to dig deeper into the underlying behavioral traits that generate these observed treatment effects. We decompose reciprocity into two terms, direct reciprocity and generalized or altruistic reciprocity (the “pay it forward” effect). In Section 6.4, we discuss the estimation results and their consequences for our understanding of behavioral effects that drive the gains in observed coordination and efficiency presented in Section 5.1. In Section 6.5, we examine network configurations that emerge in the data which are stable under the parameter estimates. We investigate the role of individual heterogeneity in Section 6.6, another benefit of the repeated-choice modelling paradigm. Finally in Section 6.7, we discuss robustness and the performance of our estimators in these data.


5.1. Reduced-form Results. First, we examine the effects of the treatment on several quantities of interest, to determine how this information affects the evolution of gameplay. The results of these reduced-form estimates are shown in Table 1. Figure 1 shows the evolution of the means of these quantities through experimental periods, and confidence intervals constructed using standard errors clustered at the group level. All estimations highlight a strong impact of the information treatment on key outcomes. In estimations, for computational convenience, we normalize contributions to lie in the interval between 0 and 1.

A natural question is how contributions and average degree (number of outgoing links) are impacted by the information treatment. The dynamics of these variables can be found in Figures 1a and 1b along with corresponding estimation results in columns (2) and (3) in Table 1. We find that the treatment substantially increases both contributions and linking. Contributions still show a tendency to decrease over time, and the rate of this decay is not significantly impacted by the treatment. Links, on the other hand, display an additional differential dynamic; although there is a strong initial boost from the treatment intervention, decay in the number of links actually appears to accelerate relative to groups in the baseline sessions.


The estimation in column (1) of Table 1 shows the effects of the treatment on subjects’ ability to coordinate on efficient structure. By efficient structure, we refer to a network topology which




(A) Contributions




(B) Links




(C) Isolated Nodes



(D) Individual Costs



(E) Reciprocity



(F) Centralization




FIGURE 1. Dynamics of key outcomes

satisfies the conditions required for efficiency by Proposition 2, without necessarily satisfying the requirement of full contributions. In other words, this is a binary indicator taking the value 1 if the observed network structure is theoretically capable of supporting efficient outcomes, and takes the value 0 otherwise. Because the version of the game we implemented was purely congestive, an efficient structure is one in which there are no “isolated” players who do not form any outgoing links. Figure 1c shows how the number of isolated players changes over time. As shown in the figure and estimates, the treatment has an immediate impact in supporting efficient structure. This is evidenced by a large downward jump in the number of isolated players at the time of intervention – and correspondingly, a significant positive effect of Treatment on Efficient Structure in column (1). The likelihood of an isolated node, however, continues to increase over time after the intervention; similar to the average degree (number of links), its growth appears to accelerate. This indicates that degree reductions are not coming exclusively from highly connected players, which in turn impacts the structural ability of the social system to reach efficient outcomes.

Contribution decisions and link decisions in isolation may not capture the true dynamic of behavior. This is because both variables together determine the cost of sharing – the marginal return of a player’s contribution in this purely congestive game decreases as they share with more other players. For each observed action, we compute the direct cost to the sharing individual as $(1 - m_i(A_i))c_i$. The dynamics of these individual costs can be seen in Figure 1d. Due to the decreasing trend in both contributions and linking, costs follow a downward trend. Treatment sessions see a large fixed jump in individual costs, and like contributions, the differential effect is not significantly different than zero, indicating that costs continue to decrease at the same rate as they did before the intervention, albeit from a significantly higher starting point.

A commonly cited behavioral driver of sharing behavior in similar environments, and a natural candidate for behavioral preferences, is reciprocity (e.g. Fehr et al., 1997; Fehr and Gächter, 1998). We postulate that individuals in the treatment session can use their new information to coordinate with other players who are sharing with them. This type of reciprocity would make contribution decisions locally complementary. Thus, a natural measure for such reciprocity is the product of incoming and outgoing benefit flows from each player. We compute this measure for each network cross-section as $\sum_{i=1}^n \sum_{j=1, j \neq i}^n A_{ij}A_{ji}m_i(A_i)m_j(A_j)c_i c_j$. This measure shows a negative trend in the baseline sessions as players fail to successfully coordinate on reciprocal outcomes in the absence of information. The information intervention has a positive effect both immediately and

TABLE 1. **Treatment effects on key outcomes**

	(1) Efficient structure	(2) Contributions	(3) Links	(4) Costs	(5) Reciprocity	(6) Centralization
Period	-0.0105*** (0.00295)	-0.0319*** (0.00455)	-0.0253** (0.0118)	-0.0106*** (0.00212)	-0.000667*** (0.000160)	0.00640*** (0.00100)
Treatment	0.375*** (0.0593)	0.783*** (0.0900)	1.668*** (0.252)	0.353*** (0.0446)	0.0252*** (0.00482)	-0.124*** (0.0173)
T × Period	-0.0147*** (0.00436)	0.0108 (0.0101)	-0.0644** (0.0274)	7.58×10^{-6} (0.00577)	0.00202** (0.000805)	-0.00243 (0.00150)
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	1380	1380	1380	1380	1380	1362

Standard errors in parentheses, clustered for 46 groups

** $p < 0.10$, *** $p < 0.05$, **** $p < 0.01$

differentially; after the intervention, reciprocity switches from a downward to an upward trend. This trend can be observed in Figure 1e and in the estimation in column (5) of Table 1. The upward trend in reciprocity after the treatment intervention, indicates that the downward trend in individual costs in treatment periods can be explained by players progressively pruning unreciprocated links. It also indicates that even using information only from the past, players can succeed in coordinating on stable concurrently reciprocated links.

Finally in Figure 1f and column (6) of Table 1, we estimate the impact of the treatment on (de)centralization. To measure centralization, we use the Herfindahl-Hirschman index (HHI), which is computed as $\sum_{i=1}^n \left(\frac{1.6 * c_i}{\sum_{i=1}^n 1.6 * c_i} \right)^2$ (restricting focus to situations in which there is at least one player who is sharing). We observe that centralization tends to increase over time, as a small set of players tend to emerge who end up generating most of the efficiency gains for the group. The treatment has an immediate impact in encouraging decentralized networks, characterized by a more balanced profile of actively sharing players. Point estimates also suggest a differential impact of the treatment, in that its tendency is to not increase as quickly in treatment sessions. However, our estimate of this effect is not quite significant at the 10% level, with a p-value of 0.112.

6. STRUCTURAL ESTIMATION

We use network formation econometrics to structurally model the environment and investigate the root causes of the behavioral changes observed between information environments. When

incorporating both prosociality and unobserved decision factors, players make decisions based on the following evolutionary random utility model:

$$(6) \quad \mathbb{E} [U_{it}(A_t, c_t | \theta) | \Omega_{it}] = \theta_1 (\pi_i(A_t, c_t) - \kappa_i(A_{-i}, c_{-i})) + \mathbb{E} [\beta_i(A, c | \theta) | \Omega_{it}] + \epsilon_{it}$$

where Ω_{it} is used to denote the information available to player i at time t . Here, utility is linearly decomposed into three separate forces – a weighting of preference for monetary earnings π_i , a behavioral term β_i describing social preferences, and a structural error term ϵ_{it} which is i.i.d. across link structure alternatives and time. The structural shock ϵ_{it} represents alternative decision factors, observed by the decision maker but unobserved by the econometrician.

Each player then chooses their strategy to maximize this expected utility, so that:

$$(7) \quad \mathbb{E}_i [U_i(A_i, c_i | A_{-i}, c_{-i}, \theta) | \Omega_{it}] \geq \mathbb{E}_i [U_i(B_i, d_i | \theta) | \Omega_{it}] \quad \forall (B_i, d_i) \in \{0, 1\}^{n-1} \times [0, \omega]$$

Using a model based on beliefs allows us to estimate individual strategic evolution, while avoiding the bias that arises as a result of the simultaneous nature of decision making in this setting. This is enabled by our panel of small networks – the rare but ideal setting in which to study network formation.⁴ Individuals form expectations over the incoming value that will be sent to them from each other player, conditional on the local information that they are given. We denote the information available to the player in time period t as Ω_t , representing a set of observations of previous behavior upon which players can form reasonable expectations about their neighbors' future actions.

We can also relax the usual assumption that players move individually, which is common in the cross-sectional and large network estimators (Mele, 2017; Badev, 2021). This assumption is typically used to ensure the convergence of logit-response dynamics to a steady state distribution (Foster and Young, 1990; Alós-Ferrer and Netzer, 2010) from which parameters are then estimated. Observing a panel of linking decisions by a subset of nodes, set in small networks, allows us to directly (and tractably) estimate utility parameters from the evolution of gameplay. While the assumption of players using logit best-response is somewhat strict, it is a substantial relaxation of the assumptions used in cross-sectional estimators. Specifically, we no longer need to make assumptions regarding the meeting process for individual revisions.

⁴This is as opposed to the potential games approach which uses a single large cross-section and estimates parameters from the stationary distribution, given some assumptions on the symmetry of preferences, and the assumption of nonzero meeting probabilities for all dyads.

At time t , some (known) subset of nodes are selected to update their linking strategy. We place no restrictions on who or how many are chosen for revision. The nodes that are selected have the opportunity to change their link selection and contribution. They select the combination link set and contribution which maximize their utility, including the preference shock which is subject to the following assumption:

Assumption 2. The structural preference shock ϵ_{it} is i.i.d. Extreme Value type 1.

This allows us to estimate agents' best responses from variations of the multinomial logit:

$$(8) \quad f_i(A_i, c_i | \Omega_{it}, \theta) = \frac{\exp(\phi_i(A_i, c_i | \Omega_{it}, \theta))}{\int_0^{\omega_i} \sum_{G_i \in \{0,1\}^{n-1}} \exp(\phi_i(G_i, \delta_i | \Omega_{it}, \theta)) d\delta}$$

Using this logit best-response process, the probability of observing network A and contribution profile c in period t is:

$$(9) \quad f(A, c | \Omega_t, \theta) = \prod_{i=1}^n f_i(A_i, c_i | \Omega_{it}, \theta)$$

We observe a panel of N groups over T time periods. We denote the network from group k in period t as A_t^k . The resulting log-likelihood function is:

$$(10) \quad \ell(\theta) = \sum_{t=2}^T \sum_{k=1}^N \sum_{i=1}^n \ln \left(f_i(A_{it}^k, c_{it}^k | \Omega_{it}^k, \theta) \right) = \sum_{t=2}^T \sum_{k=1}^N \sum_{i=1}^n \phi_i(A_{it}^k, c_{it}^k | \Omega_{it}^k, \theta) - \ln \left(Z_{it}(\theta, \Omega_{it}^k) \right)$$


where $Z_{it}(\theta, \Omega_{it}^k)$ is the partition function from the denominator of the density (8),

$$(11) \quad Z_{it}(\theta, \Omega_{it}^k) = \int_0^{\omega_i} \sum_{G_i \in \{0,1\}^{n-1}} \exp(\phi_i(G_i, \delta_i | \Omega_{it}^k, \theta)) d\delta.$$

The MLE is then defined as the vector of parameters that maximizes this objective,

$$(12) \quad \hat{\theta} = \arg \max_{\theta \in \Theta} \ell(\theta).$$

The graphs in Figure 2 highlight the computational benefits associated with panel estimators of network formation compared to a pooled cross-sectional estimator. In these graphs, nodes represent random variables; a circle represents a single variable (the contribution of Player i), and a square represents a dense subgraph in which all of the variables contained in the labeled set are completely connected. A link indicates statistical dependence of the receiving variable on the linking variable. A link to or from a square node indicates dependence between all elements in the union of the sets,



(A) Cross-sectional

(B) Panel

FIGURE 2. **Dependence graphs for a two-player game.**

and bidirectional links are highlighted in red. A green node is observed by the decision maker, and thus can be conditioned on in estimation, while all interconnected red nodes must be treated as a single outcome variable. Since each player has $n - 1$ potential links in the network, each set A_i^t contains $n - 1$ interdependent random variables. Panel estimation improves computation time by conditioning on each player’s available information to sever concurrent interdependence between links and contributions across players. This works by essentially performing n simultaneous independent multinomial logit estimations across n variables (each potential link plus a contribution level), as opposed to a single multinomial logit over n^2 outcomes, saving a total of $2^{(n^2)} - 2^n$ operations when the support of c_i is binary. For perspective, in a network of size $n = 10$, panel estimation reduces the number of required floating point operations from 1.27×10^{30} to just 1,024 operations per cross-section.

6.1. Estimation by the pseudolikelihood approach. While not quite as extreme as the computational burden imposed by the full stationary distribution models (Mele, 2017; Badev, 2021), there is some substantial computational complexity arising with estimation from this form of the density, particularly for panels of larger networks. Specifically, computation of the full log-likelihood is $O(n2^{n-1}qNT)$ where n is the number of nodes in each network, q is the number of quadrature nodes (or the magnitude of the discretization of c), N is the number of distinct groups and T is

the total number of time periods. This complexity comes from the need to compute the partition function (11) separately for each information set Ω_{it}^k .

One approach to reduce this computation time is to adopt a pseudolikelihood approach (Stewart and Schweinberger, 2020; Boucher and Mourifié, 2017; Leifeld et al., 2018; Strauss and Ikeda, 1990). This involves introducing small intentional misspecification to the marginal densities, to approximate the log likelihood at a reduced computational burden. Estimates from the pseudolikelihood approach can be shown to be consistent, and to have good asymptotic properties in settings with weaker dependence (Stewart and Schweinberger, 2020). We consider the conditional distribution of each link:

$$(13) \quad f_i(A_{ij}|A_{-ij}, c_i, \theta, \Omega_{it}) = \frac{\exp(\phi_i(A, c_i|\theta, \Omega_{it}))}{\exp(\phi_i(A - ij|A_i^{-ij}, c_i, \theta, \Omega_{it})) + \exp(\phi_i(A + ij|A_i^{-ij}, c_i, \theta, \Omega_{it}))}$$

and the conditional distribution of each contribution level as:

$$(14) \quad f(c_i|A_i, \theta, \Omega_{it}) = \frac{\exp(\phi_i(A, c_i|\theta, \Omega_{it}))}{\int_0^{\omega_i} \exp(\phi_i(c_i = \delta|A_i, \theta, \Omega_{it}))d\delta}.$$

The pseudolikelihood approach works by approximating the joint distribution by the product of conditionals. The log pseudolikelihood function to be maximized is:

$$(15) \quad \ell(\theta) \approx \ell_p(\theta) = \sum_{t=2}^T \sum_{k=1}^N \sum_{i=1}^n \left(\ln(f(c_i|A_i, \theta, \Omega_{it})) + \sum_{j=1, j \neq i}^n \ln(f_i(A_{ij}|A_i^{-ij}, c_i, \theta, \Omega_{it})) \right)$$

This can be computed in $O((n^2 + q)NT)$, offering a substantial computational advantage over the full likelihood in the case of large data. The Maximum PseudoLikelihood Estimator is then defined as

$$(16) \quad \hat{\theta}_p = \arg \max_{\theta \in \Theta} \ell_p(\theta).$$

Although it offers an exponential speedup, the MPLE is known to be biased in certain cases. In particular, it is known to perform poorly when faced with complex dependence (Van Duijn et al., 2009).

6.2. Uncertainty. Prior work involving the MPLE has noted that it tends to understate uncertainty. Biased standard errors generated by the Hessian matrix are known to produce overly tight confidence intervals for the estimator (Van Duijn et al., 2009; Desmarais and Cranmer, 2012). To

compensate for this, we measure uncertainty in the estimators instead using Monte-Carlo simulations and bootstrapping with sequences of networks.

The first proposed estimator is obtained by a semi-parametric bootstrap. First, sequences of networks from the baseline and treatment conditions are drawn using an iterated Metropolis-Hastings sampler. Once iterated to stationarity, this sampler draws from the transition distribution associated with game evolution from the current network cross-section. After a suitable number of iterations have passed, a cross-section is taken from the sampler, and becomes the next network observation in the time series. This process is repeated to replicate the number of time periods in an observation, and then this sequence of networks and contribution profiles is repeated to generate a large sample representing the estimated distribution of network sequences. Estimates are then bootstrapped from this large sample of network sequences.

Metropolis-Hastings samplers for network games can be slow to mix (Mele, 2017). This can mean that long sequences of large networks could be prohibitively expensive to compute, since each new observation is associated with a new conditional distribution that must reach stability. For this reason, an attractive alternative is the nonparametric bootstrap. This is possible since we have replicated data and would not need to rely on subsampling networks, a difficulty arising in nonparametric bootstrap for large network cross-sections. This approach also avoids distributional assumptions, generating the most conservative estimates of variability.

A nonparametric clustered bootstrap approach works as the name suggests – data are drawn from the sample (with replacement) to form a new sample, matching the total number of observations and proportion of observations of treatment groups. In order to ensure proper clustering, each sequence of networks from one group is treated as a single observation. The estimation routine is then run again on this new dataset, and the estimates are stored. Standard errors are obtained as the standard deviation of estimates made by repeating this process many times. This approach works because drawing from the observed sequences with replacement is the best approximation we can make to drawing from the true distribution of sequences without making any distributional assumptions. Since these distributional assumptions would otherwise provide more information and a closer approximation to the true distribution, the nonparametric bootstrap provides the most conservative estimates of the true variability of the estimators.

6.3. Structural estimation results. We turn next to the structural estimation of the discrete choice framework presented in Section 6. We estimate the model using both the MLE and MPLE

approaches and compare results. In order to generate confidence intervals for estimates, we construct standard errors based on bootstrap samples of 1000 sequences of networks generated by the model using the estimated parameters, as well as 1000 sequences drawn by a nonparametric bootstrap.

For panels of small networks, the maximum likelihood estimator is straightforward to define and estimate. However, its computation time grows exponentially with the size of the network cross-sections. The maximum pseudolikelihood estimator provides an attractive alternative. It works by approximating the true likelihood function by a product of its conditionals, essentially conducting an individual logistic regression for each link and contribution in the network profile. This approximation is known to produce consistent estimates as the size of the network grows large. Fortunately, this is precisely when the MPLE shines – its required computation time grows only quadratically with the size of the networks, offering an exponential speedup relative to the MLE.

We estimate the following simple specification of behavior, in order to understand how the treatment focuses reciprocity from being generalized (commonly known as a “pay it forward” effect (Gray et al., 2014; Baker and Bulkley, 2014; Van Doorn and Taborsky, 2012; Pfeiffer et al., 2005)) to directed (focused only toward others who shared with the player, as in Fehr et al. (1997); Fehr and Gächter (1998)):

$$\beta_i(A, c|\theta) = \underbrace{\frac{\theta_2 + \theta_3 T}{(n-1)^2} \sum_{j=1, j \neq i}^n A_{ij} m_i(A_i) c_i \sum_{l=1, l \neq i}^n A_{li} m_l(A_l) c_l}_{\text{Generalized reciprocity}} + \underbrace{\frac{\theta_4 T}{n-1} \sum_{j=1, j \neq i}^n A_{ij} A_{ji} m_i(A_i) m_j(A_j) c_i c_j}_{\text{Direct reciprocity}}$$

Since our network data come from a laboratory experiment, we can easily control the group size and ensure that we generate data on which both estimators are tractable. We find that the MPLE successfully reproduces estimates which are qualitatively similar to those from the MLE, although it shows some substantial bias for terms describing complex dependence. Results of these estimations are provided in Table 2.

We also estimate uncertainty in these environments; standard asymptotic standard errors generated by Fisher information matrices, even when cluster robust, are known to serially underestimate variability in the MPLE (Desmarais and Cranmer, 2012; Van Duijn et al., 2009). Even for networks small enough to accommodate the MLE, estimates may also fail to reach asymptotics if data are limited. Therefore, we propose two algorithms to generate confidence intervals for the estimators – the semi-parametric Monte-Carlo bootstrap and the fully nonparametric bootstrap. Using these

TABLE 2. **Structural parameter estimates**

	Point estimates		Standard errors					
	MLE	MPLE	(Asymptotic)		(MC Bootstrap)		(NP Bootstrap)	
			MLE	MPLE	MLE	MPLE	MLE	MPLE
(θ_1) Contribution costs	5.2368	4.2836	0.0198	0.0181	0.1752	0.1705	0.3925	0.3633
(θ_2) Generalized reciprocity	20.1884	30.8815	0.1464	0.1543	1.4988	1.3102	2.4442	2.2928
(θ_3) Treatment \times Generalized reciprocity	-6.9893	-13.5047	0.2237	0.2406	2.3551	2.1204	4.2045	4.1546
(θ_4) Treatment \times Direct reciprocity	24.2407	24.4307	0.1136	0.1541	1.0105	1.0686	1.6560	2.3312
Number of groups	46	46			46	46	46	46
Number of treatment groups	28	28			28	28	28	28
Number of time periods	30	30			30	30	30	30
Sample size					1000	1000	46	46
Bootstrap samples taken					1000	1000	1000	1000
AIC	48530.3360							
AICc	48530.3602							
BIC	48551.7846							

Standard errors clustered for 46 groups

methods, we show that the true variances of both estimators are understated in our experimental data.

6.4. Discussion of results. Structural estimation gives us the ability to distinguish between different types of reciprocal preferences by the specific coordination patterns that they generate. Notably, we find strong statistical evidence that generalized reciprocity does not decline completely with the treatment; effort remains globally complementary in this sense. That is, players who receive more tend to then give more, even if that sharing is not specifically directed at those who shared with them. This result holds when using the confidence intervals generated by asymptotic standard errors, and is robust to the most conservative estimates of variability generated by a nonparametric bootstrap. We also verify this using Monte-Carlo simulation, the results of which are displayed in Figure 3.

We estimate four primary parameters in order to analyze how subjects react to new information. The first parameter, θ_1 , captures how players weight the cost of sharing.⁵ θ_2 measures the “pay it

⁵ θ_1 is analogous to the price coefficient in standard IO models of discrete choice.

forward” effect, which we refer to as altruistic or generalized reciprocity. This term is essentially a limited-information form of reciprocity, in which players who receive benefits from the sharing of others pay it forward by sharing, albeit not necessarily with those who shared with them. This generalized reciprocity works by determining how much, in total, was shared with a player, and averaging it across the number of players – in essence they reciprocate to their best possible guess that a player could make of how much each other group member shared with them given limited information. θ_3 gives a treatment effect on generalized reciprocity, to test whether the introduction of an ability to directly reciprocate increases or decreases this altruistic reciprocity. Finally, θ_4 represents preferences for direct reciprocity. Since this is only possible in the information treatment (beyond its approximation through generalized reciprocity), it is interacted with a treatment dummy and assumed to be zero in the baseline.

Estimation shows a positive and highly significant impact of both generalized and direct reciprocity. Further, it is clear that subjects react to direct incentives; they are less likely to share in ways that incur higher individual costs. The negative treatment effect θ_3 characterizes players’ substitution away from generalized reciprocity when direct reciprocity is enabled by the information treatment. Interestingly, θ_2 and θ_3 are jointly significant. This indicates that in our specification, the two forms of reciprocity are not perfect substitutes and that there is a persistence of altruistic reciprocity through the treatment intervention.

The persistence of generalized reciprocity suggests that even in the treatment sessions with more information, preferences may not fit the strict functional form required to be aggregable by an exact potential function. This has implications for the convergence of these dynamics to stable configurations. Notably, existing cross-sectional estimators rely on the existence of a potential function to guarantee convergence to the stochastically stable quantal response equilibrium defining the stationary distribution from which they are estimated. Because we have a panel of small networks, we can sidestep this requirement to estimate preferences that may not fit the strict form required to admit an aggregate potential.

6.5. Convergence and empirical stability. It is well-known that theoretical convergence guarantees do not obtain from simultaneous move logit-response dynamics, even in the case of potential games (Alós-Ferrer and Netzer, 2010). Simultaneous revision opportunities can lead to oscillations and instability of the fixed point. This does not, however, negate the existence of such fixed points.




FIGURE 3. **Persistence of generalized reciprocity in 1000 MC bootstrap samples**

Using our framework, we can empirically check whether an observed cross-section has reached a Nash- or topologically-Nash stable configuration.

Our data consist of a total of 1380 network cross sections. Of these, only 40 are topologically stable. The modal stable network remains the empty network (or at least, a network configuration with no contributions); of these 40, 22 of the observed stable networks are nonempty. Of these, 20 are observed in treatment sessions. These structures are summarized in Figure 4. The majority of stable networks, however, are nontrivial. Of the nonempty stable networks that we observe, the most common to emerge is the single reciprocal dyad, which appears 11 times across 4 groups, 10 of which occur during treatment sessions. This is also the only example of a nontrivial fully behaviorally Nash stable configuration that we observe; the corresponding Nash equilibrium contribution level requires both linking players to contribute fully. This fully stable outcome was remarkably persistent, appearing in the 23rd and subsequently in the 25th-30th time periods for the group that achieved it.

Apart from the single bilateral link, we observe three other stable configurations. The second most common topology to emerge is dual pairwise links, which we observe in six cross sections across two treatment groups. This type of network structure we would expect to be the most prevalent – due to the structure of the congestion incentives, a single link is the least costly at the individual




FIGURE 4. **Observed topologically stable motifs**

level and therefore stable states with individual links should have the widest basin of attraction. However, we observe two other interesting examples of topologically stable network configurations. In two time periods for one treatment group, a dense cluster of three nodes emerged. Notably, this is later replaced by two stable bilateral pairs, as the originally isolated player becomes involved.

The other anomaly occurred in a baseline session, in which a stable 4-cycle (“ring”) network emerged. In this stable configuration, each player was sharing with one other and no links were reciprocated. This network configuration is topologically stable under generalized reciprocity, since there are no benefits that could be obtained by a player switching outgoing benefit link to someone who is sharing with them in the absence of direct reciprocity. When nontrivial stable configurations arose in baseline (low information) sessions, they tended to emerge in earlier time periods (round 4 in both cases) and never persisted for more than a single period.

6.6. Individual heterogeneity. Neglecting individual heterogeneity is known to produce bias in some multinomial logit models of discrete choice. We also introduce individual heterogeneity into the model using observed characteristics collected from a questionnaire after the experiment. While individual heterogeneity appears to play a significant role in certain aspects of gameplay, accounting for it does not appear to reduce bias in the MPLE.

After the experiment, subjects were asked survey questions to get an idea of their individual characteristics. Five of these questions were aimed at eliciting a subject’s preference for reciprocity. These survey questions are reproduced in Appendix B. In order to distill answers from these questions into attributes to describe reciprocity, we performed a principal components analysis and selected the top two principal components as representative characteristics. The decision to select two components was determined by referring to a Scree plot of the singular values shown in Figure




FIGURE 5. **Scree plot of the eigenvalues determining reciprocity characteristics**

5. From this, we can see that the first two components describe more than 70% of the variation in survey responses. The results are reported in Table 3. We find some evidence of individually heterogeneous preferences, although nonparametric bootstrap results suggest that more data are needed for precise estimates. Most importantly, the point estimates for θ are robust to the introduction of our heterogeneous characteristics.

6.7. Performance of the estimators. We develop a pair of estimators designed to measure preferences by players' reactions to information. By incorporating panel data, replicated sequences of small networks, these estimators can overcome many of the computational issues that plague exponential family cross-sectional estimators of network formation (Mele, 2017; Badev, 2021; Chandrasekhar, 2016). Further, temporal and panel estimators rely on fewer assumptions regarding meeting processes, and do not require that the process has iterated for long enough to reach the fully stochastically stable stationary distribution.

The Maximum PseudoLikelihood Estimator (MPLE) produces estimates qualitatively similar to those provided by the Maximum Likelihood Estimator (MLE) at a drastically improved computational cost. Estimates produced by the MPLE may be biased. But the MPLE shines for panels of larger networks; asymptotic consistency obtains for large networks, and its computation presents an exponential speedup when compared to the MLE in this case. Both options, however, provide estimates at a substantially reduced computational cost relative to popular cross-sectional methods.

TABLE 3. **Structural parameter estimates with heterogeneity**

	Point estimates		Standard errors			
	MLE	MPLE	(Asymptotic)		(NP Bootstrap)	
			MLE	MPLE	MLE	MPLE
Contribution costs	5.4801	4.5122	0.0213	0.0177	0.3963	0.3489
Generalized reciprocity	20.8539	31.9597	0.1745	0.1781	2.8141	2.5623
T × Generalized reciprocity	-7.2340	-12.7533	0.2608	0.2469	4.6541	4.4426
T × Direct reciprocity	24.1115	23.7034	0.1208	0.1577	1.9101	2.4041
Individual Heterogeneity						
Trust × Contribution costs	0.1122	0.2104	0.0097	0.0087	0.2069	0.2213
Reciprocity1 × Contribution costs	-0.3664	-0.3324	0.0129	0.0127	0.3800	0.3624
Reciprocity2 × Contribution costs	-0.2661	0.0524	0.0168	0.0167	0.4642	0.5358
Trust × T × Direct reciprocity	-0.5424	-1.8040	0.0818	0.1254	1.3002	1.5565
Reciprocity1 × T × Direct reciprocity	1.3030	2.3572	0.0750	0.0895	1.4117	1.5481
Reciprocity2 × T × Direct reciprocity	-1.8272	-2.1358	0.0896	0.1106	2.0228	2.3875
Trust × Generalized reciprocity	0.8059	1.0751	0.4024	0.0739	1.2834	1.4696
Reciprocity1 × Generalized reciprocity	1.9023	1.1086	0.1057	0.1283	2.2625	2.7251
Reciprocity2 × Generalized reciprocity	1.1025	0.5944	0.1287	0.1855	3.3285	4.0953
Trust × T × Generalized reciprocity	-0.3302	2.8440	0.1693	0.1772	3.6599	3.6307
Reciprocity1 × T × Generalized reciprocity	-4.5654	-5.9976	0.1745	0.1737	4.3829	4.2076
Reciprocity2 × T × Generalized reciprocity	3.0949	5.6090	0.2188	0.2105	5.4058	5.5810
Number of groups	46	46			46	46
Number of treatment groups	28	28			28	28
Number of time periods	30	30			30	30
Sample size					46	46
Bootstrap samples taken					1000	1000
AIC	48460.8808					
AICc	48461.2128					
BIC	48547.4754					

Standard errors clustered for 46 groups

This reduction in computational complexity comes from a relaxation of the assumption that the game has iterated long enough to reach its stationary distribution which requires an entire network and contribution profile as a single outcome in the multinomial logit estimation. Rather, if we observe transitions and temporal behavior, we need only focus on outcomes at the individual level as transitions from previous states. This leads to an information-rich environment that can help us distinguish between different forms of behavior that may not accommodate aggregation by a potential function. These include the global complementarities introduced by generalized reciprocity as well as non-reciprocal individual heterogeneity.

7. CONCLUSION

In this paper, we have developed a model of resource sharing on an endogenous network. Unlike previous models of voluntary resource sharing, we allow individuals to specifically choose the beneficiaries of their externalities. This intuitive extension provides a convenient generalization of the classical voluntary contributions mechanism, and allows for flexible specification of group size effects in a way designed to capture realistic incentive problems.

We develop a theoretical framework for the voluntary sharing game and estimate preferences using panels of dynamic network data. Using data collected in the lab, we examine the impact of different information structures on sharing and link formation decisions, with a particular focus on the relative importance of generalized and direct reciprocity. We find that individuals in these environments effectively use information to coordinate on more efficient outcomes. While voluntary sharing is, theoretically, plagued by free-riding incentives, behavioral biases such as a preference for reciprocity can generate complementarities that allow for nontrivial stable outcomes that involve considerable sharing among groups.

When players are shared with, they tend to react by sharing with others. Even in the absence of detailed information, this generalized reciprocity can help to support some sharing behavior. When provided with more detailed information, players exhibit a clear preference for targeting the positive externalities they generate toward others who have shared with them, reciprocating directly. Using structural estimation methods, we characterize the tradeoff between direct and generalized reciprocity. We find that subjects tend to rely less heavily on generalized reciprocity when given the ability to use direct reciprocity. This substitution, however, does not completely

eliminate the role of generalized reciprocity; behavioral complementarities persist both locally and globally when directed reciprocation is enabled.

While we have demonstrated the effectiveness of our modeling paradigm and established proof-of-concept, the real power of this methodology is in its capacity to explain far more sophisticated patterns of learning and behavior. In this respect, we hope that these findings will lay the groundwork for future study in a number of possible directions, to better understand the motivations and patterns of learning that guide the formation and evolution of prosocial behavior in the sharing economy.

REFERENCES

- Alós-Ferrer, Carlos and Nick Netzer (2010), “The logit-response dynamics.” *Games and Economic Behavior*, 68, 413–427.
- Badev, Anton (2021), “Nash equilibria on (un) stable networks.” *Econometrica*, 89, 1179–1206.
- Baker, Wayne E and Nathaniel Bulkley (2014), “Paying it forward vs. rewarding reputation: Mechanisms of generalized reciprocity.” *Organization science*, 25, 1493–1510.
- Boosey, Luke A (2017), “Conditional cooperation in network public goods experiments.” *Journal of Behavioral and Experimental Economics*, 69, 108–116.
- Boucher, Vincent and Ismael Mourifié (2017), “My friend far, far away: a random field approach to exponential random graph models.” *The econometrics journal*, 20, S14–S46.
- Bramoullé, Yann, Rachel Kranton, et al. (2007), “Public goods in networks.” *Journal of Economic Theory*, 135, 478–494.
- Chandrasekhar, Arun (2016), “Econometrics of network formation.” *The Oxford Handbook of the Economics of Networks*, 303–357.
- Dasaratha, Krishna (2019), “Innovation and strategic network formation.” *arXiv preprint arXiv:1911.06872*.
- Dasaratha, Krishna (2020), “Distributions of centrality on networks.” *Games and Economic Behavior*, 122, 1–27.
- Desmarais, Bruce A and Skyler J Cranmer (2012), “Statistical inference for valued-edge networks: The generalized exponential random graph model.” *PloS One*, 7, e30136.
- Elliott, Matthew and Benjamin Golub (2019), “A network approach to public goods.” *Journal of Political Economy*, 127, 730–776.

- Fehr, Ernst and Simon Gächter (1998), “Reciprocity and economics: The economic implications of homo reciprocans.” *European economic review*, 42, 845–859.
- Fehr, Ernst, Simon Gächter, and Georg Kirchsteiger (1997), “Reciprocity as a contract enforcement device: Experimental evidence.” *Econometrica: journal of the Econometric Society*, 833–860.
- Fischbacher, Urs (2007), “z-tree: Zurich toolbox for ready-made economic experiments.” *Experimental economics*, 10, 171–178.
- Fisher, Joseph, R Mark Isaac, Jeffrey W Schatzberg, and James M Walker (1995), “Heterogenous demand for public goods: Behavior in the voluntary contributions mechanism.” *Public Choice*, 85, 249–266.
- Foster, Dean and H Peyton Young (1990), “Stochastic evolutionary game dynamics.” *Theoretical Population Biology*, 38, 219–232.
- Galeotti, Andrea and Sanjeev Goyal (2010), “The law of the few.” *American Economic Review*, 100, 1468–92.
- Gerke, Stefanie, Gregory Gutin, Sung-Ha Hwang, and Philip Neary (2019), “Netflix games: local public goods with capacity constraints.” *arXiv preprint arXiv:1905.01693*.
- Golub, Benjamin and Evan Sadler (2021), “Games on endogenous networks.” *arXiv preprint arXiv:2102.01587*.
- Gray, Kurt, Adrian F Ward, and Michael I Norton (2014), “Paying it forward: Generalized reciprocity and the limits of generosity.” *Journal of experimental psychology: General*, 143, 247.
- Greiner, Ben (2015), “Subject pool recruitment procedures: organizing experiments with orsee.” *Journal of the Economic Science Association*, 1, 114–125.
- Gupta, Harsh and Mason A Porter (2020), “Mixed logit models and network formation.” *arXiv preprint arXiv:2006.16516*.
- Hiller, Timo (2020), “A simple model of network formation with congestion effects.” Technical report, PUC Rio.
- Hommes, Cars H and Marius I Ochea (2012), “Multiple equilibria and limit cycles in evolutionary games with logit dynamics.” *Games and Economic Behavior*, 74, 434–441.
- Isaac, R Mark and James M Walker (1988), “Group size effects in public goods provision: The voluntary contributions mechanism.” *The Quarterly Journal of Economics*, 103, 179–199.
- Jackson, Matthew O and Asher Wolinsky (1996), “A strategic model of social and economic networks.” *Journal of economic theory*, 71, 44–74.

- Kinateder, Markus and Luca Paolo Merlino (2017), “Public goods in endogenous networks.” *American Economic Journal: Microeconomics*, 9, 187–212.
- Kinateder, Markus and Luca Paolo Merlino (2021), “The evolution of networks and local public good provision: A potential approach.” *Games*, 12, 55.
- Leifeld, Philip, Skyler J Cranmer, and Bruce A Desmarais (2018), “Temporal exponential random graph models with btergm: Estimation and bootstrap confidence intervals.” *Journal of Statistical Software*, 83.
- McKelvey, Richard D and Thomas R Palfrey (1995), “Quantal response equilibria for normal form games.” *Games and Economic Behavior*, 10, 6–38.
- McKelvey, Richard D and Thomas R Palfrey (1998), “Quantal response equilibria for extensive form games.” *Experimental Economics*, 1, 9–41.
- Mele, Angelo (2017), “A structural model of dense network formation.” *Econometrica*, 85, 825–850.
- Overgoor, Jan, Austin Benson, and Johan Ugander (2019), “Choosing to grow a graph: modeling network formation as discrete choice.” In *The World Wide Web Conference*, 1409–1420, ACM.
- Page, Lawrence, Sergey Brin, Rajeev Motwani, and Terry Winograd (1999), “The pagerank citation ranking: Bringing order to the web.” Technical report, Stanford InfoLab.
- Pfeiffer, Thomas, Claudia Rutte, Timothy Killingback, Michael Taborsky, and Sebastian Bonhoeffer (2005), “Evolution of cooperation by generalized reciprocity.” *Proceedings of the Royal Society B: Biological Sciences*, 272, 1115–1120.
- Stewart, Jonathan R and Michael Schweinberger (2020), “Pseudo-likelihood-based m -estimation of random graphs with dependent edges and parameter vectors of increasing dimension.” *arXiv preprint arXiv:2012.07167*.
- Strauss, David and Michael Ikeda (1990), “Pseudolikelihood estimation for social networks.” *Journal of the American statistical association*, 85, 204–212.
- Van Doorn, Gerrit Sander and Michael Taborsky (2012), “The evolution of generalized reciprocity on social interaction networks.” *Evolution: International Journal of Organic Evolution*, 66, 651–664.
- Van Duijn, Marijtje AJ, Krista J Gile, and Mark S Handcock (2009), “A framework for the comparison of maximum pseudo-likelihood and maximum likelihood estimation of exponential family random graph models.” *Social networks*, 31, 52–62.

APPENDIX A. EXPERIMENTAL INSTRUCTIONS

Experimental Instructions

Introduction

Thank you for participating in today's experiment. I will read through the script so that everyone receives the same information. Please remain quiet and do not communicate with other participants during the experiment. Raise your hand if you have any questions. Someone will come to you to answer the question privately.

For your participation in today's experiment, you will receive the show-up fee of \$7. In addition, during the experiment, you will have the opportunity to earn more money. Your additional earnings will depend on the decisions you make and on the decisions made by other participants. At the end of the experiment, you will be paid anonymously by check. No other participant will be informed about your payment.

The experiment consists of multiple parts. The instructions for each part will only be distributed and read after previous parts have been completed.

Part 1

All amounts in this part will be expressed in **tokens**. At the end of the experiment, your earnings will be converted to dollars according to the exchange rate, **80 tokens = \$1**.

This part of the experiment consists of 15 decision rounds. Before the rounds begin, you will be randomly divided into groups consisting of 4 participants. Each player in your group will be assigned a letter ID (W, X, Y, or Z). Group members and letter IDs are fixed for the duration of this part.

In each round, you and the other members of your group will each be given 20 tokens. You can divide these 20 tokens between two accounts: **Account A** and **Account B**.

Account A. Any tokens you put into Account A will be kept as earnings for the round. Thus, if you put 10 tokens in Account A, your earnings from Account A will be 10 tokens.

Account B. Any tokens you put into Account B generate earnings for you and for any other group members **of your choosing**. Specifically, in addition to dividing your 20 tokens between Account A and Account B, you must designate, in each round, which of the other group members will receive earnings from your Account B.

The tokens you put in Account B will be multiplied by a factor of **1.6** and then divided equally between you and the group members who you designate as recipients.

For example, suppose you designate two other group members as recipients, and put 10 tokens into Account B. Those 10 tokens will be multiplied by 1.6, so that there are 16 tokens to divide between you and your designated recipients. In this case, you and the two group members you designated would each receive 5.33 tokens in earnings from your Account B.

For another example, suppose you designated one other group member as a recipient and put 10 tokens into Account B. As before, those 10 tokens will be multiplied by 1.6, so that there are 16 tokens to divide. In this case, you and the group member you designated as a recipient would each receive 8 tokens in earnings from your Account B.

Note that if you do not designate any recipients, then all 20 tokens must be placed into Account A. Thus, it is not possible for you to put tokens into Account B and receive the multiplied amount entirely to yourself.

How do you make your decisions?

A screenshot of the decision screen is shown below. In the top left panel, you can enter the number of tokens you want to put into Account B. You can enter any integer number of tokens from 0 to 20.

In the right panel, you can designate recipients by clicking on the boxes for the other members of your group. Your own box is automatically selected and highlighted in gray. To select another player as a recipient, click on their box and it will be highlighted in blue. If you want to unselect a player, click on the highlighted box again and it will be unhighlighted.

Round 1 of 15

You are Player W
You have 20 tokens to allocate

Please select your designated recipients (if any) below:

How much would you like to allocate to Account B?
1

OK

History of your decisions and earnings by round:

Round	Account B Allocation	Designated Recipients	Round earnings:

Player W
Player X
Player Y
Player Z

When you are ready to submit both your designated recipients and your allocation decision, click the “OK” button.

Your Earnings

In each round, you can collect earnings in three ways.

- (1) **From your own Account A** - any tokens you put in Account A will be paid to you as earnings at the end of the round
- (2) **From your own Account B** - any earnings generated for you out of your own Account B will be paid at the end of the round
- (3) **From the Account B of each player who selected you as one of their designated recipients** - if another player chose you as a designated recipient, you will receive some earnings from their Account B. How much you earn from them depends on how much they put in their Account B and how many other recipients they designated.

Example Scenario. Suppose you put 10 tokens into Account A and 10 tokens into Account B, while you selected one other player in your group to be a designated recipient. In addition, suppose that two of the other players (say, X and Y) selected you as one of their designated recipients. Suppose player X put 20 tokens into his Account B and selected 2 designated recipients (one being you), while player Y put 10 tokens into her Account B and selected only 1 designated recipient (you). How would your earnings be calculated in this scenario?

- First, you would receive **10 tokens** in earnings from your Account A.
- Second, both you and your one designated recipient would receive **8 tokens** in earnings from your Account B, since the 10 tokens you allocated will be multiplied by 1.6 and divided between the two of you.
- Third, you would receive some earnings from player X and player Y, each of whom selected you as a designated recipient. From Player X, your earnings will be equal to $(20 \times 1.6)/3 = \mathbf{10.66 \text{ tokens}}$, since the multiplied amount in his Account B is divided between him and his two designated recipients. From Player Y, your earnings will be equal to $(10 \times 1.6)/2 = \mathbf{8 \text{ tokens}}$, since the multiplied amount in her Account B is divided between her and her one designated recipient (which is you).

Adding these together, your earnings for the round in this scenario would be

$$(10 + 8 + 10.66 + 8) = \mathbf{36.66 \text{ tokens.}}$$

At the end of the experiment, your total earnings from Part 1 will be the sum of your earnings from all 15 rounds.

Feedback

At the end of each round, you will see the following information:

- The number of tokens you allocated to Account B
- A list of the players you selected as designated recipients
- Your earnings from Account A
- The earnings received by you and any designated recipients you selected from your allocation to Account B
- Your total earnings for the round, which includes any earnings you received as a designated recipient for other players' Account B allocations.

In the bottom panel on the left of the decision screen, the history table shows all of your decisions (allocation and designated recipients) and payoffs from previous rounds.

Part 2

This part of the experiment is almost identical to the previous part. As in Part 1, all amounts are expressed in tokens and your earnings will be converted at the end of the experiment according to the exchange rate, **80 tokens = \$1**.

Part 2 also consists of 15 decision rounds. Your group of 4 participants will be the same as in Part 1. However, before the first round in this part, you will be randomly assigned a new letter ID (W, X, Y, or Z). The new letter IDs will remain fixed for the duration of this part.

The only differences between this part and the previous part relate to the feedback that you will receive. In this part, at the end of each round, you will see all of the same information as in Part 1. In addition, you will see **a summary of the earnings you received from the Account B of each player who selected you as a designated recipient**. Furthermore, on the decision screen, there will be an additional history panel that shows you the Account B earnings you received from each player in all previous rounds of Part 2 (this panel will only show up after Round 1 of this part).

At the end of the experiment, your total earnings from Part 2 will be the sum of your earnings from all 15 rounds.

Part 3 (no need to distribute)

This part of the experiment is an individual survey. All of the instructions are provided on the screen. There are four (4) screens to complete.

APPENDIX B. BEHAVIORAL CHARACTERISTIC QUESTIONNAIRES

Reciprocity Questionnaire

On a scale from 1 (this statement does not apply to me at all) to 7 (this statement applies to me perfectly) please indicate how well you believe the following statements apply to you personally.

1. If someone does me a favor, I am prepared to return it.
2. If I suffer a serious wrong, I will take revenge as soon as possible, no matter what the cost.
3. If somebody puts me in a difficult position, I will do the same to them.
4. I go out of my way to help somebody who has been kind to me before.
5. If somebody insults or offends me, I will offend or insult them back.
6. I am ready to undergo personal costs to help somebody who has helped me before.

Trust Questionnaire

Please indicate to what extent you agree or disagree with the following statements.

1: Strongly disagree, 2: Disagree somewhat, 3: Agree somewhat, 4: Strongly agree

1. In general, one can trust people.
2. These days you cannot rely on anybody else.
3. When dealing with strangers, it is better to be careful before you trust them.