

# STRATEGIC FORMATION OF COLLABORATIVE NETWORKS

PHILIP SOLIMINE<sup>1</sup> AND LUKE BOOSEY<sup>2</sup>

[Click here for most recent version]

**ABSTRACT.** We examine behavior in an experimental collaboration game that incorporates endogenous network formation. The environment is modeled as a generalization of the voluntary contributions mechanism. 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. Providing players more detailed information about the sharing behavior of others drastically increases efficiency, and positively affects a number of other key outcomes. To understand the driving causes of these changes in behavior we develop and estimate a structural model for actions and small network panels and identify how social preferences affect behavior. We find that the treatment reduces altruism but stimulates reciprocity, helping players coordinate to reach mutually beneficial outcomes. In a set of counterfactual simulations, we show that increasing trust in the community would encourage higher average contributions at the cost of mildly increased free-riding. Increasing overall reciprocity greatly increases collaborative behavior when there is limited information but can backfire in the treatment, suggesting that negative reciprocity and punishment can reduce efficiency. The largest returns would come from an intervention that drives players away from negative and toward positive reciprocity.

---

*Date:* April 12, 2024.

The authors would like to thank Christopher Brown, Arun Chandrasekhar, Krishna Dasaratha, Dean Eckles, Matthew Gentry, Ben Golub, Mark Isaac, Matthew Jackson, Annalise Maillet, Angelo Mele, Luca Paolo Merlino, Anke Meyer-Baese, Svetlana Pevnitskaya, Tim Rooney, 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 seminar participants at UBC, ESA Job Market Candidates Seminar, Conference(s) 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 FSU Computational Exposition. 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 in working form under the title “Resource sharing on endogenous networks”.

<sup>1</sup> Vancouver School of Economics, University of British Columbia

*Correspondence:* philip.solimine@ubc.ca

<sup>2</sup> Department of Economics, Florida State University.

## 1. INTRODUCTION

We study a voluntary collaboration game in which players choose how much, and with whom, to share. In this setting, shared resources generate externalities for other players, but the player sharing the resources does not derive any immediate or direct benefit from doing so. Viewing the decisions about who to share with as a process of endogenous network formation, we estimate a structural model on preferences and use it to infer the value of different social preferences for trust and reciprocity.

Consider, for example, 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 effort is a scarce resource and spreading effective contributions across multiple projects reduces the impact on each project individually. By creating efficiency gains, this type of collaboration can play a valuable role in economic systems in which it is embedded. Apart from scientific collaboration, a number of relevant real-world examples fit into the basic framework of collaboration with congestion. On the internet, web pages make linking decisions; they are paid advertising revenue for traffic, and they are indexed by a search engine which favors well-linked pages, such as Google’s PageRank algorithm (Page et al., 1999). Linking to another site diverts a fraction traffic and amplifies search appearances for the linked site. Another example is networks between firms who form links by collaborating on R&D projects (see, e.g. Dasaratha, 2023), and determine how much intellectual property they will share with other firms through this process.

A pervasive feature of these collaborative environments, in both digital and analog settings, is that a player who invests their effort 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.

While the examples above are distinct from each other in several ways, they each highlight some key common characteristics of voluntary sharing in collaborative environments. First, the decision to share naturally entails some individual cost to the sharer. Second, shared resources are, in many cases, congestible. They congest in the sense that sharing the same resources across multiple recipients reduces 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 reciprocally; firms and coauthors may be

more inclined to dedicate their time to a collaborative project if they notice that their fellow researchers are similarly dedicated.

In such settings, trust and reciprocity play a key roles in helping players coordinate to reach efficient outcomes (Fehr et al., 1997; Fehr and Gächter, 1998). People with social preferences tend to enjoy helping others who have helped them. By controlling the specific structure and content of the information a player has about their neighbors' actions, we aim to reason about how changes in these social preferences can be identified by the changing patterns of collaborative behavior.

In order to examine the underlying influences on voluntary resource sharing, we designed an experiment to study a simple sharing environment that incorporates the three key features outlined above. 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.<sup>1</sup> 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.

Our experiment allows us to examine the impact of the information structure on reciprocal behavior. In the control, players are given information only about 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 altruism or generalized reciprocity. In the information treatment, players are informed about the incoming benefit from each of the other three group members. In the treatment, they can specifically identify which other group members shared with them, and how much. This enables direct reciprocity to influence both contribution and network formation decisions. In each session of the experiment, groups first played 15 rounds of this collaboration game in the control condition with limited information. Then, they were informed that they would play another 15 rounds with the same group (albeit with shuffled IDs). In control sessions, the second part simply repeated the control condition for 15 more rounds. In the treatment sessions, these rounds included the additional information.

---

<sup>1</sup>This 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 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.<sup>2</sup> 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 main contributions is methodological in nature – we implement a new approach (see, e.g., Overgoor et al., 2019; Gupta and Porter, 2022) to estimate the network formation game as a combination of individual evolutionary discrete choice strategies.

Our other main contributions are more conceptual in nature. The voluntary sharing environment we study is novel, but it 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; Kinatered 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.

Our reduced-form results alone point to striking effects of the information treatment. We find that when subjects can observe who shares with them, as well as how much they shared, contributions and the average number of links increase substantially. Fewer nodes are isolated, facilitating the formation of efficient structure, and contribution profiles are more balanced (decentralized) across group members. As one might expect, the availability of more detailed information leads to contribution and linking patterns that exhibit dramatically higher levels of reciprocal behavior than in the baseline condition.

Building off these stark reduced-form results, we develop and estimate a structural model that incorporates both generalized reciprocity and direct reciprocity as core behavioral considerations for the participants. Our estimates suggest that in the baseline condition (without specific information), sharing behavior is driven partly by generalized reciprocity. By contrast, when subjects are provided with more specific information, sharing behavior is heavily

---

<sup>2</sup>See Chandrasekhar (2016) for a comprehensive introduction to the study of network formation.

influenced by direct reciprocity, and subjects substitute sharply away from generalized reciprocity concerns in favor of this more focused form of reciprocity.

The estimates of our structural parameters are robust to the inclusion of several heterogeneous individual characteristics, which we construct based on subjects' responses to a post-experiment survey. We distill the responses down to three key attributes using a principal components analysis, representing trust, overall reciprocity, and positive reciprocity. Incorporating these heterogeneous characteristics into our estimations, we observe clear and intuitive effects of heterogeneity on behavior. In the baseline condition, higher scores on the trust attribute are consistent with trustworthy behavior (stronger effects of generalized reciprocity), but predict less trusting decisions as subjects tend to be more cautious about sharing altruistically, perhaps out of mistrust in the notion that others will reciprocate. In the treatment condition, knowing that others observe the source of incoming benefits allows subjects with higher trust attributes to be more trusting.

The two principal components of reciprocity—overall reciprocity and positive reciprocity—also explain much of the variation in behavioral patterns. Overall reciprocity, which captures both a taste for positive and negative reciprocity, has a nuanced effect in the two conditions. In the baseline, subjects with a higher overall reciprocity attribute are more altruistic and exhibit a stronger preference for generalized reciprocity. In the treatment condition, they shift more weight towards direct reciprocity, as expected, but also exhibit less altruism and less concern for generalized reciprocity. One interpretation for these differences is that in the treatment condition, the additional information about sharing decisions of others facilitates negative reciprocity (punishment of those who do not share) more than it enhances positive reciprocity.

Subjects with higher scores on the positive reciprocity attribute are also more altruistic in the baseline condition. However, in contrast with our results for the overall reciprocity attribute, in the treatment condition, they share generously with the entire group in hopes of stimulating the reciprocity of others, rather than resorting to punishment out of negative reciprocity.

Finally, we use our structural framework to conduct three counterfactual simulations, each examining the effects of a uniform increase in one of the three principal attributes—trust, overall reciprocity, and positive reciprocity. Consistent with our estimates for the model with individual heterogeneity, an increase in the trust attribute improves key outcomes in the information treatment, but has a mild, negative impact on the baseline condition. Overall reciprocity improves outcomes in the baseline condition, but can backfire substantially due to negative reciprocity in the treatment condition. However, a uniform increase in subjects' positive reciprocity attribute has a substantial positive effect on all key outcomes in both the treatment and baseline conditions.

Our study naturally contributes to three important and active bodies of literature in economics. We briefly review these connections—to network public goods games, to structural estimation of network formation data, and to social preferences for trust and reciprocity—in the remainder of this section.

**1.1. Network provision of public goods.** 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; Parise and Ozdaglar, 2023).

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 provision to environments with endogenous linking include Galeotti and Goyal (2010), which furthers the specialization result of Bramoullé et al.

(2007). These papers emphasize the prevalence of core-periphery architectures as equilibrium networks, but in a setting where players choose incoming, rather than outgoing links. Kinaterder 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).

In another relevant study by Rand et al. (2011), the authors conducted an experiment to gauge the effects of endogenous networks on cooperation in a repeated prisoner's dilemma. By varying the opportunity for network updates, they showed that subjects are able to take advantage of their ability to change social ties in order to refine their neighborhoods and increase efficiency. Our results will show that, while the endogeneity of the network itself does allow for this fine-tuning of the social neighborhood, the dynamics of the network alone are not sufficient to support long-term efficient outcomes. Instead, a platform that aims to nudge players toward efficient social structure should take advantage of its ability to shape and distribute information to its users.

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 in the previous literature on public goods and sharing on endogenous networks (e.g. Galeotti and Goyal, 2010; Kinaterder and Merlino, 2017; Brown, 2024), in that individuals choose the *outgoing* direction of their externalities, rather than the incoming direction of others' externalities. The cost of linking in this study's environment is explicitly tied to the effort or contribution level and can be flexibly specified to represent pure or impure (congestive) externalities. 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.

Hiller (2022) presents a similar model of network formation with a continuous effort choice and where effort levels are strategic complements. In their 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 our setting, actions are perfect substitutes and the only complementarities among players' actions are those created by their preferences for reciprocity. Finally, Dasaratha (2023) also derives theoretical results for a similar model of collaboration with a specific application to collaborative innovation in industrial networks.

**1.2. Structural estimation from network data.** 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). Thus far, the literature has focused predominantly on preference estimation from large-network cross-sectional data. This is not without reason—panel data consisting of repeated observations from small networks over time is rare and difficult to find in practice. By collecting experimental data, however, we can 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. More specifically, 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, examined only very recently by Overgoor et al. (2019) and Gupta and Porter (2022).

Cross-sectional network formation estimators rely on assumptions about 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, but differs by leveraging panel structure to avoid the assumption that the network has already converged to its steady-state distribution after sufficient iterations of an individual revision process.

**1.3. Social preferences for trust and reciprocity.** There is a large literature in behavioral and experimental economics that points toward the importance of various behavioral traits and heterogeneous characteristics of trust and reciprocity in sharing behavior. A



large series of studies including Fehr et al. (1997); Fehr and Gächter (1998, 2000); Camerer (2003); Cox (2004) have leveraged experimental evidence to highlight the use of trust and reciprocity as devices for contract enforcement and for driving cooperation in markets and sharing games. Throughout these works is emphasized that reciprocity manifests not only positively, but can also be used to characterize punishment behavior. While trust and reciprocity in these settings often drive efficiency gains, we show through counterfactuals that which characteristics matter the most is highly context dependent; both trust and reciprocity can backfire as tools to promote efficiency, depending on the information structure.

This result agrees with some more recent work examining the role of these characteristics in supporting positive market outcomes. For example, Choi and Storr (2022) finds evidence suggesting that providing reputation systems in experimental markets interacts with preferences primarily by giving participants more information about whom not to trust. Subsequent work by Solimine and Isaac (2023) supported this result and further emphasized the role of the information in determining the effectiveness of trust in promoting positive market outcomes. Our counterfactual findings involving trust agree with these findings; through the way that trust interacts with preferences for reciprocity and altruism, promoting trust in the community dramatically improves outcomes when subjects are provided with detailed information about others' behavior. When information is more limited, however, introducing higher levels of trust increases trustworthy behavior by some but may backfire by allowing others to take advantage of this change.

These results also agree with other experimental work that has specifically focused on the role of trust in experimental sharing environments. In particular, Glaeser et al. (2000) found that survey questions about trust (similar to the questions we used to measure trust) are effective at predicting trustworthy behavior but less so at predicting trusting behavior. The distinction between trusting and trustworthy behavior is mirrored in Anderson et al. (2004), who found that certain measures of trust are negatively associated with sharing in public goods experiments.

Studies that focus specifically on reciprocity have found similar results. Fehr and Gächter (2000) explains that “[reciprocity] means that in response to friendly actions, people are frequently much nicer and much more cooperative than predicted by the self-interest model; conversely, in response to hostile actions they are frequently much more nasty and even brutal.” Indeed, reciprocity is not always positive and evidence from the intersection of neuroscience and economics (Fehr et al., 2005) agrees that reciprocity can really be broken down into two pieces—rewards from mutual cooperation and a taste for punishment of unfair behavior. Through a factor analysis of survey responses, we break reciprocity into these two components; one that measures “overall reciprocity” and positively weights both types, and one component that distinguishes between positive and negative reciprocity. We

find that while increasing trust can mildly increase free-riding in the absence of information, and promoting overall reciprocity can backfire in the presence of information by stimulating punishment. However, an intervention that is specifically targeted at moving players away from punishment and instead focuses on rewarding mutual cooperation can be effective at promoting collaboration in both information environments.

**1.4. Outline of the paper.** In Section 2, we lay out a simple theoretical framework for the collaborative sharing environment. Section 3 describes the design and procedures for the laboratory experiment testing the effects of different information structure on collaboration patterns. We present and discuss the reduced-form results of the experiment in Section 4. Then, in Section 5 we develop and estimate a structural empirical framework to analyze social preferences using our experimental data, and discuss the results of the estimation and the goodness of fit. This section also presents the three counterfactual simulations designed to measure the value of social preferences for trust and reciprocity. Section 6 concludes.

## 2. CONCEPTUAL FRAMEWORK

We represent player decisions in a single period by directed networks with implicit self-links, coupled with a vector of actions. 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  $\omega_i$  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  $i^{\text{th}}$  column of  $A$  with the  $i^{\text{th}}$  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  $\pi_i$  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 stage game, the optimal contribution for any linking pattern is  $c_i = 0$  for all  $i$ .  $\square$

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  $\beta_i$  is a general function specifying the social preferences. Since the externality term  $\kappa$  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  $\theta$ , 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, it is more restrictive than the classical notion of “pairwise stability” (Jackson and Wolinsky, 1996). The conditions in Definition 1 are straightforward from a game theoretic perspective. Condition 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 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 2.** 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 3.** 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  $|N_i| = n$  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

$$(6) \quad \sum_{i=1}^n \pi_i(A, c) = \sum_{i=1}^n \left( \frac{k}{|N_i|+1} c_i + \sum_{j=1}^n A_{ji} \frac{k}{|N_j|+1} c_j \right) = \sum_{i=1}^n k c_i = k \sum_{i=1}^n c_i,$$

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.  $\square$

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.

### 3. 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.

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.


#### 4. REDUCED-FORM RESULTS

In this section, we highlight several aggregate, reduced-form results showing the treatment effects of information provision on key outcomes of interest. The reduced-form estimates are presented in Table 1. Moreover, Figure 1 illustrates the evolution of the mean level (for each key outcome) across experimental periods, with bars representing standard errors clustered at the group level. All estimations highlight a strong impact of the information treatment on key outcomes. For computational convenience, in each estimation, we normalize contributions to lie in the interval between 0 and 1.


**4.1. Contributions and links.** 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 (1) and (2) 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, exhibit 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 mildly relative to groups in the baseline sessions.

**Result 1.** *Contributions and average degree (number of links) are both significantly higher in the treatment condition than in the baseline condition.*


**4.2. Efficient Structure.** The estimation in column (3) 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 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




(A) Contributions (in Experimental Currency Units)




(B) Links




(C) Isolated Nodes



(D) Individual Costs



(E) Reciprocity



(F) Centralization



FIGURE 1. Dynamics of key outcomes



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 (3). The likelihood of an isolated node, however, continues to increase over time after the intervention; and 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.

**Result 2.** *The information treatment has a significant positive impact on the formation of efficient network structure, by reducing the prevalence of isolated nodes.*

**4.3. Individual Costs.** 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 are illustrated 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.

**Result 3.** *Individual costs of sharing are significantly higher in the treatment condition than in the baseline, although they exhibit a similar downward trend over time in both conditions.*

**4.4. Reciprocity.** 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). Our main hypothesis is that individuals in the treatment condition 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 exhibits 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 differentially. In the treatment condition, reciprocity displays a large positive level effect

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

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

Standard errors in parentheses, clustered for 46 groups

and switches from following a downward to an upward trend. These trends are apparent in Figure 1e and supported by the coefficient estimates in column (5) of Table 1. The upward trend in reciprocity after the treatment intervention indicates that the downward trending individual costs in treatment sessions can be explained primarily by players progressively pruning unreciprocated links. It also suggests that, even using information only from the past, players can succeed in coordinating on stable concurrently reciprocated links.

**Result 4.** *Reciprocity is significantly higher in the treatment condition than in the baseline condition. Moreover, it increases over time in the treatment, while it decreases over time in the baseline.*

**4.5. Centralization.** 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 as in the baseline. However, our estimate of this effect is not quite significant at the 10% level, with a p-value of 0.112.

**Result 5.** *In the treatment condition, we observe more decentralized networks—exhibiting a more balanced profile of contributors—than in the baseline condition.*

Taken together, these reduced-form results suggest that subjects are highly effective at using the specific information provided in the treatment condition to coordinate on more

efficient collaboration networks. To further investigate the driving forces behind these treatment effects, we adopt a structural approach. In the next section, we introduce our structural framework and implement panel structural estimators of network formation in order to better understand the underlying behavioral traits that generate our reduced-form results.

## 5. STRUCTURAL ESTIMATION

**5.1. Structural Model.** 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:

$$(7) \quad \mathbb{E}[U_i(A_t, c_t | \theta) | \Omega_{it}] = \theta_1 (\pi_i(A_t, c_t) - \kappa_i(A_{-i}, c_{-i})) + \mathbb{E}[\beta_i(A, c | \theta) | \Omega_{it}] + \epsilon_{it}$$

where  $\Omega_{it}$  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  $\pi_i$ , a behavioral term  $\beta_i$  describing social preferences, and a structural error term  $\epsilon_{it}$  which is i.i.d. across link structure alternatives and time. The structural shock  $\epsilon_{it}$  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:

$$(8) \quad \mathbb{E}_i[U_i(A_i, c_i | \theta) | \Omega_{it}] \geq \mathbb{E}_i[U_i(B_i, d_i | \theta) | \Omega_{it}] \quad \forall (B_i, d_i) \in \{0, 1\}^{n-1} \times [0, \omega]$$

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.<sup>3</sup> 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  $\Omega_t$ , 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

---

<sup>3</sup>This 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.

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.

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 expected utility, including the preference shock which is subject to the following assumption:

**Assumption 2.** The structural preference shock  $\epsilon_{it}$  is i.i.d. Extreme Value type 1.

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

$$(9) \quad f_i(A_i, c_i | \Omega_{it}, \theta) = \frac{\exp(\phi_i(A_i, c_i | \Omega_{it}, \theta))}{\int_0^{\omega_i} \sum_{G_i \in \{0,1\}^{n-1}} \exp(\phi_i(G_i, \delta_i | \Omega_{it}, \theta)) d\delta}$$

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

$$(10) \quad f(A, c | \Omega_t, \theta) = \prod_{i=1}^n f_i(A_i, c_i | \Omega_{it}, \theta)$$

**5.2. Methodology.** 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:

$$(11) \quad \ell(\theta) = \sum_{t=2}^T \sum_{k=1}^N \sum_{i=1}^n \ln(f_i(A_{it}^k, c_{it}^k | \Omega_{it}^k, \theta)) = \sum_{t=2}^T \sum_{k=1}^N \sum_{i=1}^n \phi_i(A_{it}^k, c_{it}^k | \Omega_{it}^k, \theta) - \ln(Z_{it}(\theta, \Omega_{it}^k))$$

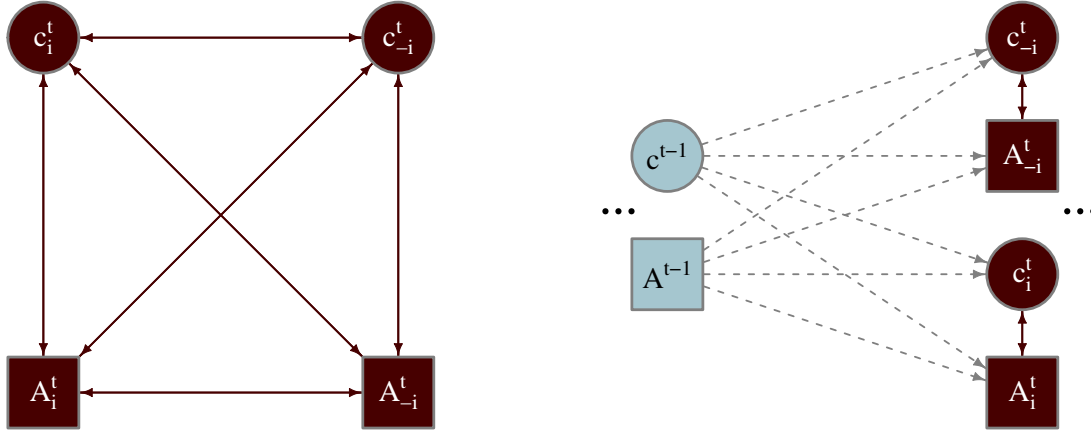
where  $Z_{it}(\theta, \Omega_{it}^k)$  is the partition function from the denominator of the density (9),

$$(12) \quad Z_{it}(\theta, \Omega_{it}^k) = \int_0^{\omega_i} \sum_{G_i \in \{0,1\}^{n-1}} \exp(\phi_i(G_i, \delta_i | \Omega_{it}^k, \theta)) d\delta.$$

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

$$(13) \quad \hat{\theta} = \arg \max_{\theta \in \Theta} \ell(\theta).$$

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, and bidirectional links are highlighted in dark red, while dashed grey links are not bidirectional. A light blue node is observed by the



(A) Cross-sectional

(B) Panel

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

decision maker, and thus can be conditioned on in estimation, while all interconnected dark 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  $1.27 \times 10^{30}$  to just 1,024 operations per cross-section.

**5.3. Inference.** We also estimate uncertainty in these environments by computing asymptotic standard errors using a partial likelihood approach, to account for possible misspecification of the joint density.<sup>4</sup> The partial likelihood approach assumes only that the conditional distribution is properly specified by  $f$ , and focuses on a partial likelihood function  $\ell$  which is defined as

$$(14) \quad \ell_k^p(\theta) = \sum_{t=1}^T \sum_{i=1}^n \log f_t(A_{it}^k, c_{it}^k \mid \Omega_{it}^k, \theta)$$

<sup>4</sup>For example, misspecification may arise due to serial correlation in the error term within a group. See chapter 13 of Wooldridge (2010) for details of this approach.

Point estimates from the MLE and PMLE are identical, but the asymptotic variance of the partial likelihood estimator adjusts for clustering by group.

That is, in the limit as  $N \rightarrow \infty$ , the distribution of  $\hat{\theta}_{\text{PMLE}}$  is determined by the central limit theorem

$$(15) \quad \sqrt{N} \left( \hat{\theta}_{\text{PMLE}} - \theta_0 \right) \sim \mathcal{N} \left( \mathbf{0}, A_0^{-1} B_0 A_0^{-1} \right)$$

where  $A_0 = -\mathbb{E} [\nabla_{\theta}^2 \ell_k^p(\theta_0)]$  and  $B_0 = \mathbb{E} \left[ \nabla_{\theta} \ell_k^p(\theta_0) (\nabla_{\theta} \ell_k^p(\theta_0))^{\top} \right]$ . The key difference between the PMLE and MLE standard error formulas lies in the summation over  $t$  within the function  $\ell_k^p$ , which must be accounted for in computation of the expected score and Hessian. In this exponential family, the scores can be expressed as the sufficient statistics minus their expected values under  $\theta$ . Typically in a large network this would need to be estimated using Monte-Carlo methods, but conveniently, because the limited group size allows us to directly compute the normalizing constant for each  $f_{it}$ , we can efficiently compute the scores directly during the calculation of the normalizing constant. Crucially, the summation across time periods must take place before the outer product in an estimate of the Fischer information matrix.

Following Wooldridge (2010), we estimate  $A_0$  and  $B_0$  as

$$(16) \quad \hat{A}_0 = -\frac{1}{N} \sum_{t=1}^T \sum_{k=1}^N \sum_{i=1}^n \nabla_{\theta} \ell_{it}^k(\theta) (\nabla_{\theta} \ell_{it}^k(\theta))^{\top}$$

$$(17) \quad \hat{B}_0 = \hat{A}_0 + \frac{1}{N} \sum_{k=1}^N \sum_{t=1}^T \sum_{r \neq t} \sum_{i=1}^n \nabla_{\theta} \ell_{it}^k(\theta) (\nabla_{\theta} \ell_{ir}^k(\theta))^{\top}$$

where the additional cross products in  $\hat{B}_0$  compensate for potential serial correlation. Finally, the asymptotic variance matrix is estimated as  $\frac{1}{N} \hat{A}_0^{-1} \hat{B}_0 \hat{A}_0^{-1}$ .

Because of the reasonably limited size of our dataset, we also construct p-values through a fully nonparametric bootstrap based on resampling of full group sequences. The nonparametric clustered bootstrap approach works as its name suggests – sequences 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, it is crucial that 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 leverages the argument that 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. As such, the nonparametric bootstrap should generally provide a more conservative estimate of the true variability of the estimators.

Further, we do not expect the MLE to reach its asymptotic distribution within a limited sample size. Therefore, we make no assumptions on the normality of the estimator evaluated on our sample, which would be implicit in using the standard error as our metric of uncertainty and suggesting that it would imply t-statistics or p-values. Instead, we present nonparametric estimates of the p-values which are computed directly from the bootstrap results as the empirical probability that the estimator takes the opposite sign to what was estimated on the full sample.

**5.4. Estimation Results.** We turn next to the structural estimation of the discrete choice framework presented above. We first estimate the following simple specification of behavior, in order to understand how the treatment changes incentives by directing attention toward direct and generalized reciprocity, (as in Fehr et al. (1997); Fehr and Gächter (1998)). We also incorporate some persistence parameters to account for dynamic effects or inertia in decisions:

$$\begin{aligned}
\beta_{it}(A, c|\theta) = & \underbrace{\theta_2 T \kappa_i(A_{-i}, c_{-i})}_{\text{Change in altruism}} + \underbrace{\frac{\theta_5 T}{n-1} \sum_{j=1, j \neq i}^n A_{ij,t} A_{ji,t-1} m_i(A_{i,t}) m_j(A_{j,t-1}) c_{i,t} c_{j,t-1}}_{\text{Reciprocity}} \\
& + \underbrace{\frac{\theta_3 + \theta_4 T}{(n-1)^2} \sum_{j=1, j \neq i}^n A_{ij,t} m_i(A_{i,t}) c_{i,t} \sum_{k=1, k \neq i}^n A_{ki,t-1} m_i(A_{k,t-1}) c_{k,t-1}}_{\text{Generalized reciprocity}} \\
& + \underbrace{\theta_6 (c_{i,t} - c_{i,t-1})^2}_{\text{Contribution inertia}} + \underbrace{\theta_7 \sum_{j=1, j \neq i}^n |A_{ij,t} - A_{ij,t-1}|}_{\text{Link inertia}}
\end{aligned}$$

We estimate this model using the data from our laboratory experiment, and present the results of these estimations in Table 2.

We focus on the estimates of five primary parameters.<sup>5</sup> The first parameter,  $\theta_1$ , captures how players weight the cost of sharing. Not surprisingly, all estimates of this parameter are positive. This means that players follow their incentives and that, holding all else constant, they are more likely to choose actions that are less costly.

The other four main coefficients concern the behavioral component to payoffs. The positive sign of  $\theta_2$ , which is the coefficient on the interaction between contribution cost and the treatment, indicates that having access to this new information makes players more careful about who they share with. We can interpret this coefficient as a reduction in altruism due to the treatment, since it indicates a tendency of players to focus more on their individual costs of sharing.

<sup>5</sup>Both of the inertia terms are significantly negative, indicating a tendency of players to bias their actions toward those that they took in the previous round.

TABLE 2. **Structural parameter estimates**

	Estimate	MLE	
		SE	Bootstrap p-val
$(\theta_1)$ Contribution costs	9.355	0.107	0.000
$(\theta_2)$ Treatment $\times$ Contribution costs	2.629	0.224	0.214
$(\theta_3)$ Generalized reciprocity	7.011	0.090	0.005
$(\theta_4)$ Treatment $\times$ Generalized reciprocity	-9.177	0.279	0.038
$(\theta_5)$ Treatment $\times$ Reciprocity	26.155	0.382	0.000
$(\theta_6)$ Contribution inertia	-5.082	0.038	0.000
$(\theta_7)$ Link inertia	-0.542	0.009	0.000
Number of groups	46	46	46
Number of treatment groups	28	28	28
Number of time periods	30	30	30
Bootstrap samples			1000
Asymptotic standard errors clustered by group			

The coefficients  $\theta_3$  and  $\theta_4$  are related to generalized reciprocity. The interpretation of generalized reciprocity is that it measures the tendency of players to share more when they receive more overall benefits as a result of their group’s sharing in the previous round, without regard to who exactly shared with them. The estimate of  $\theta_3$  is positive indicating that in the baseline sessions, generalized reciprocity drives some of the observed sharing behavior. The negative coefficient on the treatment effect ( $\theta_4$ ), however, drives behavior in the opposite direction. Because the sum of effects is significantly negative, this suggests a strict tradeoff between generalized and direct reciprocity, and highlights how the information provided in the treatment helps to focus reciprocity toward active collaborators.

The final coefficient of primary interest ( $\theta_5$ ) is on the interaction between the treatment effect and our measure of direct reciprocity. The large magnitude and significance of this estimate indicate that direct reciprocity is indeed a strong driver of the more successful collaboration observed in treatment groups.

**5.5. Individual heterogeneity.** After the main parts of the experiment, subjects were asked a series of survey questions to elicit measures of their individual characteristics. All of these questions were aimed at eliciting a subject’s heterogeneous preferences toward trust and reciprocity. The full set of survey questions are reproduced in Appendix B. In order to distill answers from these questions into a reasonable number of attributes to describe






FIGURE 3. **Scree plot of the eigenvalues determining reciprocity characteristics**

reciprocity and trust, we performed a principal components analysis and selected the top two principal components as representative characteristics of reciprocity, along with the first principal component of the trust questionnaire responses to represent trust. The decision to select two components for reciprocity was determined by referring to a Scree plot of the singular values shown in Figure 3. From this, we can see that the first two components of the reciprocity questionnaire describe more than 70% of the variation in survey responses. The first principal component, referred to as overall reciprocity, captures both positive and negative reciprocity, while the second principal component, referred to as positive reciprocity, favors positive reciprocity and eschews negative reciprocity.

We estimate the same model as above with additional interaction terms between the main parameters of interest and the subject’s individual characteristics (the trust attribute, overall reciprocity attribute, and positive reciprocity attribute). The results are reported in Table 3. We find strong evidence that individually heterogeneous preferences drive differences in behavior. Importantly, the point estimates for  $\theta_1$ – $\theta_7$  are robust to the introduction of these heterogeneous individual characteristics.

Because there are only three trust questions, the first principal component summarizes most of the information from the trust questionnaire. It places positive weight on the question that involves trust and negative weights on two questions that suggest mistrust. Perhaps surprisingly, this measure of trust is associated with a positive interaction on contribution costs in the baseline, which indicates that individuals who score highly on trust are less altruistic and more careful about where they direct effort in the baseline. This agrees with the

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

	Estimate	MLE	
		SE	Bootstrap p-val
Contribution costs	10.471	0.110	0.000
T × Contribution costs	2.480	0.229	0.241
Generalized reciprocity	9.385	0.084	0.003
T × Generalized reciprocity	-13.033	0.305	0.002
T × Reciprocity	25.039	0.415	0.000
Contribution inertia	-5.626	0.037	0.000
Link inertia	-0.516	0.008	0.000
<b>Individual Heterogeneity</b>			
Trust × Contribution costs	1.951	0.055	0.011
Overall reciprocity × Contribution costs	-2.966	0.066	0.051
Positive reciprocity × Contribution costs	-6.916	0.117	0.002
Trust × T × Contribution costs	-1.951	0.123	0.091
Overall reciprocity × T × Contribution costs	4.413	0.143	0.052
Positive reciprocity × T × Contribution costs	4.266	0.131	0.068
Trust × Generalized reciprocity	2.128	0.048	0.099
Overall reciprocity × Generalized reciprocity	4.286	0.064	0.065
Positive reciprocity × Generalized reciprocity	-8.645	0.130	0.021
Trust × T × Generalized reciprocity	0.347	0.277	0.428
Overall reciprocity × T × Generalized reciprocity	-6.300	0.330	0.052
Positive reciprocity × T × Generalized reciprocity	9.356	0.314	0.121
Trust × T × Reciprocity	3.151	0.392	0.167
Overall reciprocity × T × Reciprocity	3.289	0.373	0.149
Positive reciprocity × T × Reciprocity	-2.345	0.382	0.302
Number of groups	46		
Number of treatment groups	28		
Number of time periods	30		
Number of bootstrap samples			1000
Asymptotic standard errors and bootstrap p-values clustered for 46 groups.			

results of Glaeser et al. (2000), who suggest that such trust questionnaires predict trustworthy behavior but do not necessarily predict trusting behavior. Further in line with these results is a strong positive interaction of the trust characteristic with generalized reciprocity in the baseline. This suggests that these individuals are trustworthy in that they respond more to sharing by others by increasing their own effort. However, they are less likely to share blindly and trust that others will reciprocate. In the treatment, estimates of the effect of trust are less precisely estimated but suggest a reversal of this phenomenon; they trust that others will reciprocate when they know that others will be aware of their sharing behavior. This is captured by the negative estimate of the interaction between trust, the treatment indicator, and contribution costs, together with the positive estimate on the interaction between trust, the treatment indicator, and direct reciprocity. This sheds more light on information as a

mechanism driving the mixed results regarding trust and sharing behavior in public goods games in previous work (Anderson et al., 2004).

The characteristic that we describe as overall reciprocity consists of positive weights on the answers to all of the questions in the reciprocity questionnaire. This includes both questions about positive reciprocity (e.g. “If someone does me a favor, I am prepared to return it”), as well as negative reciprocity (“If someone puts me in a difficult position, I will do the same to them”). Estimates of the interaction between this characteristic and the behavioral utility terms suggests that these individuals are more altruistic in the baseline and behave more in line with generalized reciprocity. At the onset of the treatment, they also shift more weight toward direct reciprocity. However, this shift toward direct reciprocity is potentially offset by a decrease in altruism (measured by additional weight placed on the costs of contributing) coupled with a strong decrease in generalized reciprocity. This suggests that individuals who have a high overall reciprocity attribute use new information to discriminate between collaborators as a mechanism for punishment.

On the other hand, the second component of reciprocity places positive weight on questions involving positive reciprocity and negative weight on questions involving negative reciprocity or punishment. Individuals who align with this characteristic place much lower weight on the actual cost of contributing, suggesting some altruism. While there is some tradeoff in the treatment, the sign of the aggregate interaction term remains negative in the treatment suggesting that these players are still behaving more altruistically than average. Perhaps surprisingly, there is a strong negative coefficient on the interaction between positive reciprocity and generalized reciprocity in the baseline. These together suggest that their increased sharing is not conditional on having received more benefits from their group, possibly representing a tendency to share in anticipation that others will behave reciprocally. This interpretation is reinforced by a large positive effect of the treatment on generalized reciprocity for this group, offset by a small decrease in direct reciprocity. In other words, these individuals reciprocate by sharing with the entire group, and trusting in the reciprocity of others, rather than by using new information as a tool for punishment.

**5.6. Simulations and goodness of fit.** To address concerns about how well the structural model fits the data, and to enable our counterfactual investigations involving behavioral interventions, we start by conducting some simulations. The simulations are conducted semi-parametrically, by bootstrap sampling initial conditions with replacement and using them as the starting point for a sequence of network and contribution profiles, drawn with Markov-Chain Monte Carlo (MCMC) sampling. Although a full MCMC chain would be slow to mix, we can exploit conditional independence relationships, due to the panel structure, to build a partial Gibbs sampler that draws from each player’s next strategies individually—treating a

tuple  $(A_i, c_i)$  individually using a combination of MCMC and rejection sampling. Whenever possible, we avoid distributional assumptions and opt for semi-parametric procedures.

The partial Gibbs MCMC sampler works as follows; first we draw a starting panel from an initial distribution  $f_0(A, c)$ . For this we resample from the starting panels observed in the data either in the treatment or control. Next, we iteratively draw each player's next strategy using a standard Metropolis-Hastings sampler. This leverages the fact that players' strategies are independent conditional on the previously observed network. Using a proposal distribution  $(A_i^m, c_i^m) \sim g_i(A_i, c_i)$ , the acceptance ratio for sample  $m$  is determined by  $\alpha = \frac{\exp(\phi_i(A_i^m, c_i^m | \Omega_{it}, \theta))}{\exp(\phi_i(A_i^{m-1}, c_i^{m-1} | \Omega_{it}, \theta))} \frac{g_i(A_i^{m-1}, c_i^{m-1})}{g_i(A_i^m, c_i^m)}$ , and sample  $m$  is accepted with probability equal to  $\min\{1, \alpha\}$ . Due to the standard theory of Metropolis-Hastings sampling, we then have  $\lim_{m \rightarrow \infty} (A_i^m, c_i^m) \sim f_i(A_i^m, c_i^m | \Omega_{it}^k, \theta)$ . In practice, we choose a large number  $M$  and choose the action  $(A_i^M, c_i^M)$  to be included in the next network panel.<sup>6</sup> This is repeated for each of the  $n$  players, and the resulting panel  $(A^M, c^M)$  is then taken to be the next step in the sequence, and this process is then repeated for each of the  $T$  time periods to generate a single full sequence. This partial Gibbs sampler (that is, drawing each player's strategy variables separately) greatly improves the number of accepted samples, and thus the speed of mixture, by dramatically decreasing the magnitude of the sample space and thus increasing acceptance probabilities.

In practice, because we want the proposal distribution to be close to the true distribution, we draw from a proposal distribution  $g_i$  that is constructed as follows; the empirical distribution of contributions in the baseline and treatment sessions along with a random binary vector of links with equal probability on each feasible configuration. That is,

$$(18) \quad g_i(A_i^m, c_i^m) = \begin{cases} \frac{1}{2^{n-1}-1} \frac{1}{nNT} \sum_{t=1}^T \sum_{k=1}^N \sum_{i=1}^n \mathbb{I}[c_{it}^k = c_i^m] & \text{if } c_i^m \neq 0 \text{ and } A_i^m \neq \mathbf{0} \\ \frac{1}{nNT} \sum_{t=1}^T \sum_{k=1}^N \sum_{i=1}^n \mathbb{I}[c_{it}^k = 0] & \text{if } c_i^m = 0 \text{ and } A_i^m = \mathbf{0} \\ 0 & \text{otherwise} \end{cases}$$

After drawing a sample of 1000 simulated groups, 500 in the treatment and 500 in the control, we evaluate the fit of the model by replicating our reduced form estimations from Table 1. The results of this replication for the MLE estimates are shown in Tables 4 and 5. These evaluations of the goodness-of-fit of the structural model are entirely promising – point estimates match up nearly perfectly in sign and closely in magnitude to the reduced form static and dynamic treatment effects. A further look at the simulations in Figures 4, 5, 6 reinforces this finding. In these panels, the cross-group averages of 10,000 simulated groups for each condition are shown as pale lines overlaying the dashed lines (which show the patterns from the original data.) These figures highlight the ability of this structural model to reproduce observed patterns of behavior on all of our key metrics. Crucially, because the

<sup>6</sup>For this paper, we used a sample length of 10,000 in between accepted actions.

TABLE 4. **Reduced form replication under MLE estimates**

	(1) Efficient structure	(2) Contributions	(3) Links	(4) Costs	(5) Reciprocity	(6) Centralization
Period	-0.00434 (0.000526)	-0.00896 (0.000746)	-0.00297 (0.00211)	-0.00193 (0.000196)	-0.000319 (0.0000287)	0.00102 (0.000191)
Treatment	0.459 (0.0363)	1.113 (0.0554)	1.969 (0.126)	0.342 (0.0155)	0.0391 (0.00567)	-0.110 (0.00923)
T × Period	-0.00803 (0.00164)	-0.0113 (0.00266)	-0.0467 (0.00583)	-0.00470 (0.000741)	0.000195 (0.000253)	0.00175 (0.000438)
Group FE	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	30000	30000	30000	30000	30000	30000

Standard errors in parentheses, clustered for 1000 groups

TABLE 5. **Reduced form replication with estimates with heterogeneity**

	(1) Efficient structure	(2) Contributions	(3) Links	(4) Costs	(5) Reciprocity	(6) Centralization
Period	-0.00217 (0.000581)	-0.00422 (0.000893)	0.00364 (0.00213)	-0.000567 (0.000250)	-0.000156 (0.0000389)	0.000735 (0.000181)
Treatment	0.449 (0.0340)	0.876 (0.0536)	1.702 (0.123)	0.274 (0.0156)	0.0379 (0.00576)	-0.0911 (0.00836)
T × Period	-0.00810 (0.00163)	-0.000625 (0.00271)	-0.0412 (0.00587)	-0.00152 (0.000779)	0.000845 (0.000246)	0.00104 (0.000406)
Group FE	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	30000	30000	30000	30000	30000	30000

Standard errors in parentheses, clustered for 1000 groups

vast majority of these metrics (with the sole exception of individual costs) are not directly fit by the structural model, this provides some confidence that the model is accurately capturing key features of subject behavior.

**5.7. Counterfactual behavioral interventions.** In our estimated model, it is clear that behavioral heterogeneity features (trust, overall reciprocity, positive reciprocity) drive changes in collaboration patterns across the experiments. In a set of counterfactual simulations, we examine the effects of uniform upward shifts in each one of these characteristics across the population. These can be described as “benign behavioral interventions”, because they mimic the effects of priming players toward a certain social preference but without directly altering payoffs or the information structure. We use these counterfactuals to reason about which behavioral traits would be the most valuable to promote collaboration.

From a platform design perspective, we can think of these as policy counterfactuals to investigate marketing strategies that would emphasize certain behavioral traits. More generally, due to the complex interactions between these characteristics and the different forms of altruism and reciprocity, it is useful to know which behavioral traits are the most valuable in promoting collaboration. The results, shown in Figures 4, 5, and 6, highlight the effects of these counterfactual adjustments. In each figure, we show the average of 10,000 simulations in the counterfactual regime, shown as the darker lines, juxtaposed against 10,000 simulations from the standard model as well as the values actually observed in the data (shown as a dashed line).

Figure 4 shows the effect of a one-standard deviation increase in the trust component, applied uniformly across the population. As the figure shows, this intervention improves outcomes along all measures, including contributions, links, reciprocity, and centralization, in the treatment condition. In the baseline however, the effect is moderately reversed. This is explained by the reduction in altruism and overall willingness to internalize sharing costs as well as an increase in free-riding to take advantage of those who are more trusting.

The next set of simulations, shown in Figure 5, shows the effects of a uniform increase in overall reciprocity. From this figure, we can see that our measure of overall reciprocity is essentially capturing the opposite effect as trust; in the baseline overall reciprocity creates substantial increases in linking and contributions. In the treatment, however, when subjects are given access to information that can be used as a tool for punishment, observed networks are actually less efficient than before the change.

Finally, Figure 6 shows the effects of an increase in positive reciprocity, and thus a shift away from a punishment mindset and toward a focus on gains from mutual effort and collaboration. As we can see, shifting players toward positive reciprocity has some of the largest effects in the baseline, with large increases in linking behavior and increases in contributions that are similar to those from a shift in overall reciprocity. Unlike the other examples, however, the increases from positive reciprocity are consistent across both the baseline and the treatment, generating substantial efficiency gains in both settings. From this perspective, it seems that positive reciprocity is the most consistently valuable of these social preferences if the goal is to encourage efficient collaboration.

## 6. CONCLUSION

In this paper, we have developed a model of collaboration as 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

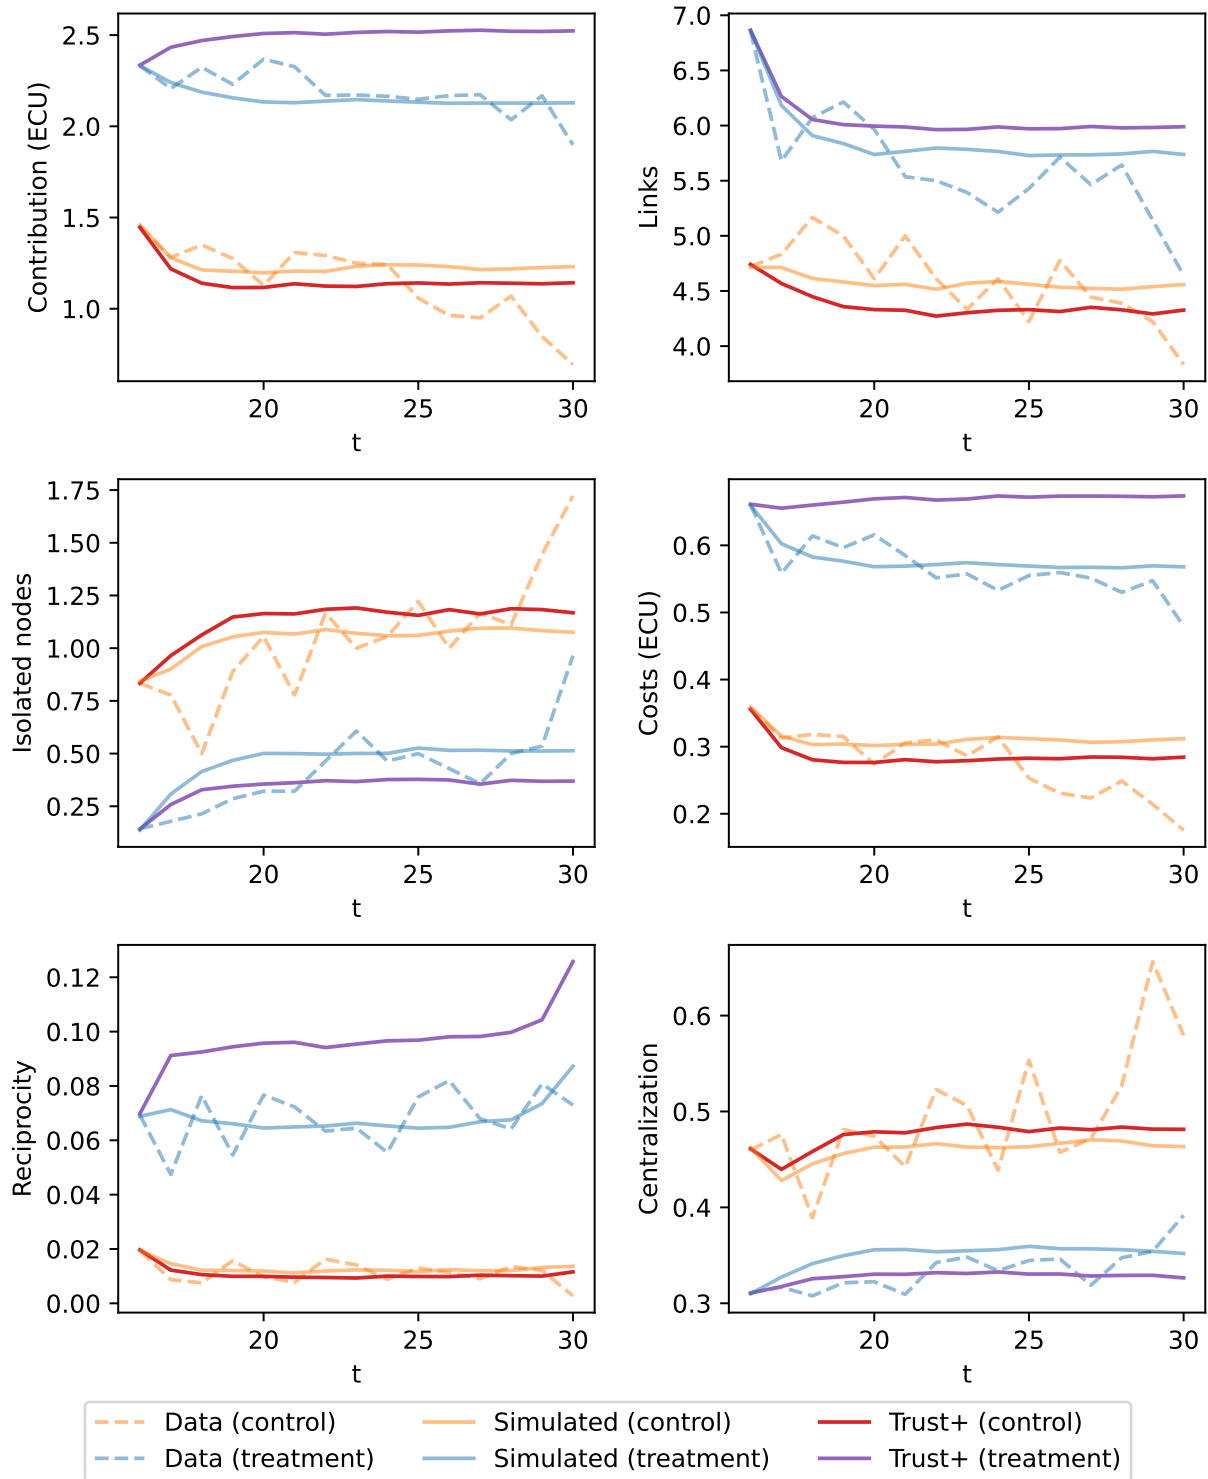


FIGURE 4. Results of counterfactual uniform increases to trust




FIGURE 5. Results of counterfactual uniform increases to overall reciprocity






FIGURE 6. Results of counterfactual uniform shifts toward positive reciprocity and away from negative reciprocity

for flexible specification of group size effects in a way designed to capture realistic incentive problems.

We characterized simple theoretical and empirical frameworks for the voluntary sharing game and used them to estimate social preferences from 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 different forms of trust and reciprocity. We find that individuals in these environments are highly effective at using information to coordinate on more efficient outcomes. While this form of collaboration is, theoretically, plagued by free-riding incentives, behavioral biases such as a preference for reciprocity can generate complementarities that help support nontrivial collaborative outcomes.

When provided with more detailed information, players exhibit a clear preference to target the positive externalities they generate toward others who have shared with them, reciprocating directly. Using our structural estimation methods, we characterize the tradeoff between altruism and different forms of reciprocity. We find that subjects tend to rely heavily on direct reciprocity when it is available, substituting away from a more generalized version of reciprocity. We also find that the subjects' reliance on different types of reciprocity are correlated strongly with heterogeneous behavioral characteristics created by a feature analysis of survey questions collected during the experiment. Using the interactions between these features and the different forms of reciprocity, we simulate counterfactual increases in different types of social preferences on the platform. The effects of these interventions vary strongly based on the information environment. Increasing trust can backfire in the baseline by increasing the incentive to free-ride, and increasing overall reciprocity has the opposite effect of driving efficiency gains in the baseline but falling short in the treatment due to increases in punishment behavior. On the other hand, interventions that focus on pushing players toward positive reciprocity (returns from mutual cooperation) and away from negative reciprocity (punishment) have a consistent effect of boosting collaboration and efficiency.

While we have demonstrated the effectiveness of our modeling paradigm and established proof-of-concept, the real power of our methodology is in its capacity to explain far more sophisticated patterns of learning and behavior. In particular, using laboratory experiments to collect small panels of network data opens the door to a new family of statistical tools that can be used to identify the structure of social preferences, and which have been understudied due to the massive size of most network data. 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 collaboration 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.
- Anderson, Lisa R, Jennifer M Mellor, and Jeffrey Milyo (2004), “Social capital and contributions in a public-goods experiment.” *American Economic Review*, 94, 373–376.
- Badev, Anton (2021), “Nash equilibria on (un) stable networks.” *Econometrica*, 89, 1179–1206.
- Boosey, Luke A (2017), “Conditional cooperation in network public goods experiments.” *Journal of Behavioral and Experimental Economics*, 69, 108–116.
- Bramoullé, Yann, Rachel Kranton, et al. (2007), “Public goods in networks.” *Journal of Economic Theory*, 135, 478–494.
- Brown, Christopher L (2024), “Team production in endogenous networks.” *Journal of Economic Behavior & Organization*, 217, 560–580.
- Camerer, Colin F (2003), “Behavioural studies of strategic thinking in games.” *Trends in cognitive sciences*, 7, 225–231.
- Chandrasekhar, Arun (2016), “Econometrics of network formation.” *The Oxford Handbook of the Economics of Networks*, 303–357.
- Choi, Ginny Seung and Virgil Henry Storr (2022), “The market as a process for the discovery of whom not to trust.” *Journal of Institutional Economics*, 18, 467–482.
- Cox, James C (2004), “How to identify trust and reciprocity.” *Games and economic behavior*, 46, 260–281.
- Dasaratha, Krishna (2020), “Distributions of centrality on networks.” *Games and Economic Behavior*, 122, 1–27.
- Dasaratha, Krishna (2023), “Innovation and strategic network formation.” *The Review of Economic Studies*, 90, 229–260.
- Elliott, Matthew and Benjamin Golub (2019), “A network approach to public goods.” *Journal of Political Economy*, 127, 730–776.
- Fehr, Ernst, Urs Fischbacher, and Michael Kosfeld (2005), “Neuroeconomic foundations of trust and social preferences: initial evidence.” *American Economic Review*, 95, 346–351.
- Fehr, Ernst and Simon Gächter (1998), “Reciprocity and economics: The economic implications of homo reciprocans.” *European economic review*, 42, 845–859.
- Fehr, Ernst and Simon Gächter (2000), “Fairness and retaliation: The economics of reciprocity.” *Journal of economic perspectives*, 14, 159–182.
- 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.
- Glaeser, Edward L, David I Laibson, Jose A Scheinkman, and Christine L Soutter (2000), “Measuring trust.” *The quarterly journal of economics*, 115, 811–846.
- Golub, Benjamin and Evan Sadler (2021), “Games on endogenous networks.” *arXiv preprint arXiv:2102.01587*.
- 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 (2022), “Mixed logit models and network formation.” *Journal of Complex Networks*, 10, cnac045.
- Hiller, Timo (2022), “A simple model of network formation with competition effects.” *Journal of Mathematical Economics*, 99, 102611.
- 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.
- 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.
- Parise, Francesca and Asuman Ozdaglar (2023), “Graphon games: A statistical framework for network games and interventions.” *Econometrica*, 91, 191–225.
- Rand, David G, Samuel Arbesman, and Nicholas A Christakis (2011), “Dynamic social networks promote cooperation in experiments with humans.” *Proceedings of the National Academy of Sciences*, 108, 19193–19198.
- Solimine, Philip and R Mark Isaac (2023), “Reputation and market structure in experimental platforms.” *Journal of Economic Behavior & Organization*, 205, 528–559.
- Wooldridge, Jeffrey M (2010), *Econometric analysis of cross section and panel data*. MIT press.

## 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.

The screenshot shows a decision interface for 'Round 1 of 15'. The user is identified as 'Player W' and has 20 tokens to allocate. The interface is divided into several sections:

- Top Left:** 'You are Player W' and 'You have 20 tokens to allocate'. Below this is a text input field for 'How much would you like to allocate to Account B?' with a small '1' inside, and a red 'OK' button.
- Top Right:** 'Please select your designated recipients (if any) below:'. This section contains a list of player names: 'Player W' (highlighted in gray), 'Player X', 'Player Y', and 'Player Z'.
- Bottom:** A table titled 'History of your decisions and earnings by round:' with columns for 'Round', 'Account B Allocation', 'Designated Recipients', and 'Round earnings:'. The table is currently empty.

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.