

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

**REPEATED EXCHANGE NETWORK STRUCTURES AFFECT
FUTURE TRUST TOWARD UNMET IN-GROUP AND OUT-GROUP OTHERS**

Joseph Quinn*

University of South Carolina

Ashley Harrell

Duke University

*Corresponding author. Direct correspondence to: joe.quinn@sc.edu
911 Pickens Street, Room 307, Columbia, SC 29201

ABSTRACT

Purpose: This paper pairs insights from social exchange theory with scholarship on in-group preferences. We ask: how do the structure and diversity of the exchange network in which an actor is embedded at time 1 impact subsequent trust toward an unmet individual at time 2? Our goal is to understand whether and how the structures in which people are embedded have lasting, downstream effects on behavior, even toward those whom they might be inclined to distrust.

Methodology/Approach: We randomly assign participants to repeated exchange tasks with different structures (productive, reciprocal, or generalized) with alters who either share or do not share a salient social identity. After a period of interaction in their exchange structure, participants decide whether, and how much, to trust a new alter who either shares or does not share their social identity in a one-shot trust decision.

Findings: Participants embedded in productive exchange networks are more likely than those in generalized or reciprocal exchange networks to trust an unmet interaction partner. Moreover, while trust is higher when the trustee is an in-group member, this relationship is moderated by the form of exchange. Trust is not lower toward outgroup trustees when the truster was previously embedded in reciprocal exchange.

Social and Theoretical Implications: Our findings collectively suggest that prior exchange structures can affect the extent to which people trust unmet others from different groups. They also imply that extended exposure to, and prosocial interactions with, out-group others may not be a core prerequisite of intergroup trust.

Keywords: social exchange; trust; emotions; networks; social identity; intergroup contact.

INTRODUCTION & BACKGROUND

Social structures affect, and are affected by, the micro-level interactions between the individuals embedded within them (Blau, 1964; Lawler, Thye, & Yoon, 2009; Serpe & Stryker, 2011; Turner, 1988). Much of the scholarship on prosocial behavior (see Simpson & Willer, 2015) and economic action (e.g., Granovetter, 1985) is devoted to understanding this micro-macro dialectic, often called *the problem of social order* (Hobbes, 1994). Scholars in this area study whether and why the actions of individuals and the structural arrangements in which they

are embedded are mutually constitutive, and how individual motivations intersect with collective interests to shape patterns of prosocial behavior between actors.

One key research program in this area is the Affect Theory of Social Exchange (ATSE). The ATSE proposes that different forms of exchange structures—the patterns in which actors exchange repeatedly with each other over time— create higher or lower levels of *task jointness* (i.e., a sense of shared responsibility and non-separability of contributions). In turn, differences in task jointness cause differences in the level of emotional commitment that actors have toward the unit and its constituents (Lawler, 2001; Lawler & Thye, 2006; Lawler, Thye, & Yoon, 2008). According to the ATSE, exchange structures that induce more emotional commitment cause higher rates of emergent cooperation, greater levels of social cohesion within the structure, and stronger levels of future giving behavior within the ongoing exchange network (Lawler & Yoon, 1996, 1998; Lawler, Thye, & Yoon, 2000, 2008, 2009).

While theories like the ATSE emphasize the effects of social structure on actions like cooperation and trust, another body of literature highlights the effect of individuals' in-group preferences on the same prosocial behaviors. This line of research demonstrates that people are more likely to trust (Foddy, Platow, & Yamagishi, 2009; Romano, Balliet, Yamagishi, & Liu, 2017; Tanis & Postmes, 2005), form ties (Marsden, 1987; McPherson, Smith-Lovin, & Cook, 2001; Melamed et al., 2020; Schelling, 1971), and cooperate with (Balliet, Wu, & De Dreu, 2014; Harrell & Quinn, 2023; Melamed et al., 2020; Whitham, 2018; Yamagishi, Jin, & Kiyonari, 1999; Yamagishi & Mifune, 2009) actors from their in-group—those with whom they share a salient social category—more often than actors from a relative out-group.

Though both research programs are well-established, simultaneous randomized tests of them are rare¹. We argue that such studies are immensely important. Manipulating both the 1) structural arrangements and 2) social identities of interactants at the same time allows researchers to clearly identify and explain the structural conditions that, social scientists often argue, likely impact cooperation, trust, cohesion, inter-group conflict, and other valenced orientations toward similar and dissimilar others. Still further, we argue that the impacts of structural arrangements and identities may impact cooperation not only *during* ongoing interactions, but also *after* such structurally induced interactions occur, via lasting effects on the treatment of ingroup versus outgroup alters. Research in this vein could contribute ground-up mechanistic explanations of top-down applied programs like intergroup contact theory, a proscriptive framework for prejudice reduction (Allport, 1954) often used to motivate research about cooperation in heterogenous groups. Studies have shown that teams built with the propositions of contact theory in mind can increase rates of cooperation and reduce levels of ongoing prejudice between dissimilar people (see Pettigrew & Tropp, 2013). These effects may impact actors even after they leave the intergroup contact environment: studies of “secondary transfer effects” suggest that being embedded in a heterogeneous and pro-social context may impact an actor’s disposition toward out-group others outside of that context (see Boin et al., 2021; Kauff et al., 2023). Evidence from studies inspired by contact theory show generally positive but sometimes mixed (Pettigrew & Tropp, 2006) and even negative effects (McKeown

¹ To the best of our knowledge, simultaneous manipulations of both social identity and network structure are nonexistent except for our previous work, discussed further below (Harrell & Quinn, 2023). A handful of observational studies (Kossinets & Watts, 2009; Kao, Joyner, & Balistreri, 2019) have used longitudinal network data to explore how variation in both the structure and identity composition of networks shapes ongoing and future individual-level behaviors, but these studies encounter three substantial limitations that inhibit their ability to assess the joint effects of the structure and composition of networks as a joint exposure: (1) without intervention, cross-group ties are inherently uncommon in observational data; (2) the social ties measured are close friendships, not frequent (and prosocial) interaction partners; (3) repeated observations in the data occur one or more years apart, making it impossible to observe interaction and exchange frequencies that likely shape outcomes of interest.

& Dixon, 2017). Some scholars partially ascribe these ironic findings to a lack of deductive clarity about lower-level mediating and moderating mechanisms – the inherently structural, compositional, and affective characteristics that can impact the quality, frequency, and outcome of contact (e.g., Schäfer et al., 2021).

Mechanism-driven inquiries involving controlled manipulations of structure *and* identity are needed to more fully understand and potentially intervene on the basic social processes of inter-group behavior—in both ongoing interactions and in future interactions with unmet others. Such studies can provide insight into the relative importance of structural- versus individual-level mechanisms of prosociality between groups, help identify *why* certain exchange settings heighten or suppress the enactment of in-group preference, and even explore how persisting structural arrangements that affect ongoing inter-group cooperation might also operate as sites where people internalize updated beliefs, expectations, and affective meanings that they bring with them to future interactions with unmet out-group others. Finally, studies that manipulate both the structure of social interaction and the identities of the interactants embedded in that structure can complement and clarify research on the construction and enactment of status beliefs (Ridgeway & Correll, 2006) or meaning making about social identities (Serpe & Stryker, 2011; Heise, 2010) or prior experiences (Lersch, 2023).

In a recent study (Harrell & Quinn, 2023), we explored how both 1) the exchange structures from the ATSE and 2) the social identities of the interactants embedded in the exchange structures jointly shaped ongoing cooperation *within* the structures themselves. We found that while actors were more cooperative toward in-group beneficiaries within a given triad, the structure of exchange itself also has a large and significant effect on the degree of cooperation that emerges within *both* diverse and homogeneous triads. The result is that the

propositions of the ATSE are robust to the inclusion of categorically dissimilar (or similar) actors within exchange structures. Productive exchange (the form of exchange with the most “jointness”) yields the highest levels of cooperation, and generalized exchange (with the lowest “jointness”) yields the lowest levels of cooperation, regardless of whether the exchange networks are diverse or homogeneous (as well as when identities are unknown). Thus, people with different salient social identities are more likely to engage in cooperative social exchange when they are actively embedded in highly joint exchange structures. We describe the study methods in more detail below.

In this paper, we report the results of an additional decision our participants made during their participation in the prior study (Harrell & Quinn, 2023). Immediately *after* each phase of their randomly assigned repeated network exchange, participants were asked to make a trust decision toward either an unmet in-group or unmet out-group alter. Using these decisions, we ask (a) whether and how both the structure and the diversity of *prior* exchange networks in the ATSE affect *future* one-shot interactions toward *unmet* partners, and (b) whether any observed effects operate differently when the unmet partner either shares or does *not* share a salient social identity with the participant making the trust decision. As we have already noted, people are typically inclined to be more trusting toward those with whom they share a salient group identity. Whether (and how) these structural conditions can reduce inter-group conflict—promoting trust among those who may be less inclined to trust each other—is of critical theoretical and practical importance.

STUDY DESIGN

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

To answer these research questions, we conducted an experiment, described also by Harrell and Quinn (2023). We recruited participants from a university-affiliated research pool. They were required to have access to a laptop or desktop computer with a stable internet connection and be located in the United States. We scheduled them to participate in interactive online sessions within an interface we programmed using zTree (Fischbacher, 2007) and zTree Unleashed (Duch, Grossmann, & Lauer, 2020). Prior to the main study tasks, participants read a consent form that explained the procedures and range of payments they could expect to receive for participating. All participants received at least \$15 for completing the study but were told they could be paid up to \$20 based on their performance.

We convened nine² participants in each of 30 online study sessions. Once all participants arrived, the study began. First, as described in more detail below, participants were randomly assigned to one of two social identity categories. Then they progressed to the exchange task, which entailed a sequence of phases of repeated exchange in triads. We randomly manipulated both 1) the exchange structure and 2) whether the interactants in the triad shared a social identity, had diverse identities, or whether identities were unknown³. While a study session was randomly assigned to only one exchange structure across the entirety of the study, all participants completed three phases within the study, one in the homogeneous identities condition, one in the diverse identities condition, and one in the unknown identities condition. In each phase, participants took part in 18 rounds of repeated exchange in their triad, then proceeded to answer a set of survey questions used to measure their affective attachment to their repeated exchange

² One session contained 12 participants.

³ While the full study involved three repeated exchange phases and two post-exchange-task decisions (a trust decision and a trustworthiness decision with a different interactant) in each phase, our analysis sample in this manuscript includes only (a) trust decisions made by (b) actors who were previously embedded in a repeated exchange phase where identities were observable. Harrell and Quinn (2023) describe other analyses from this project.

structure. Then they completed the primary task described here: a Trust Game (Berg, Dickhaut, & McCabe, 1995) in which they made one-shot decisions toward an unmet interaction partner who was saliently identified as an in-group or out-group other. We describe each segment of the design (visualized in Figure 1) in more detail below.

[Figure 1 about here]

Assigning Social Identities

After providing basic demographic information (age, gender, and whether they were a university student or not), each participant completed a Klee and Kandinsky task. The task is a standard method of establishing and randomly assigning social identities in the laboratory (Aksoy, 2015, 2019; Billig & Tajfel, 1973; Ridgeway, Boyle, Kuipers, & Robinson, 1998; Ridgeway & Erickson 2000; Simpson, 2006; Yamagishi & Kiyonari, 2000) based on a neutral trait: preference for abstract paintings. Participants viewed five pairs of paintings by two artists – abstract works by the artists Paul Klee and Wassily Kandinsky – and chose the painting from each pair that they preferred more. Each participant was then told they had been classified as either a Klee or Kandinsky based on which artist they preferred relative to the other. They were also told that in the upcoming portion of the study they would complete decision-making tasks with other participants who had also been identified as Klees or Kandinskys.

In actuality, we randomly assigned whether participants were told that they were a Klee or a Kandinsky. Randomly assigning Klee and Kandinsky identities let us balance the number of participants across sessions so they could be embedded in one “homogeneous identities” triad (i.e., an exchange structure containing two fellow participants who share their social category)

and in another “diverse identities” triad (e.g., an exchange structure with one or two partners who do not share the participant’s social identity) during different phases of their study session. We used the Klee and Kandinsky task to assign minimal group identities, instead of using *a priori* social identities such as gender or race, so that participants were able to behave in ways that were uninformed by status expectations about cultural categories encountered elsewhere. Artificially established identities like those we construct in the present study have been shown to operate like “natural” categories by causing in-group favoritism in experimental settings (Balliet et al., 2014).

Manipulating Structure and Identity Composition of Exchange Networks

The computer program randomly assigned the session to the generalized, reciprocal, or productive exchange structure condition. After the Klee and Kandinsky task, a set of identical instructions informed participants that they would participate in a repeated exchange (“decision-making”) task with two random other participants, and that those others might be Klees, Kandinskys, or both. In each round, they would make a binary decision to either *transfer* or *keep* a set of points with monetary value (i.e., “monetary units,” or MUs). Decisions would be made simultaneously; only after everyone had made their decision to either transfer or keep their MUs would they know the results for the round (i.e., how much, if any, they had earned from exchange). Participants were told that the number of MUs each participant earned throughout the study would be converted into their payment at the end. They were not told that the study contained three phases of exchange, or that each phase contained 18 rounds of exchange, to suppress end-game effects. Our payoffs in the exchange task followed the exact payoffs used by Lawler et al. (2008) – though it is worth noting that different exchange structures entail different experiences. The number of exchange decisions, number of MUs obtainable via exchange, and

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

the cost of unilateral cooperation all represent key differences in the different forms of exchange. We outline these differences in Table 1 and briefly summarize them below.

Participants then received further instructions that differed based on the exchange structure condition. Those assigned to *productive* exchange sessions were told that they would make a decision to transfer the MUs allocated to each of them at the beginning of the round to a “joint fund” shared by the triad. All MUs in the joint fund would be multiplied by a factor corresponding to the number of others who also contributed (see Table 1), and then divided equally among all three members of the triad. As in real-world examples of productive exchange (e.g., co-authoring a paper), participants benefited from their own cooperative efforts, but these benefits were greater when all others also cooperated – akin to a n-person assurance dilemma, where mutual cooperation yields the largest payoffs.

In *reciprocal* exchange sessions, participants were told that they could decide whether to transfer MUs to one, both, or neither of their interaction partners. Likewise, they were able to receive MUs from one, both, or neither of the others. All MUs sent to the other were multiplied, making MUs that were sent more valuable than MUs that were kept; participants only kept MUs if they opted to give to neither of their interaction partners.

Finally, in each round of *generalized* exchange sessions, participants could decide whether to send MUs to one specific other and could receive points from a different other (see Table 1). As in the other two conditions, MUs transferred to an other increased in value. However, by necessity participants in this condition earned fewer MUs per round compared to those randomly assigned to reciprocal or productive exchange, because they could only benefit from exchange from one other, rather than two others. Following Lawler et al. (2008), to account for this difference between the original generalized exchange condition (which we call the LTY

version in Table 1) and the other two forms of exchange, we ran several “points-adjusted” (PA) versions of the generalized exchange condition. The PA version adjusted the payoffs such that the maximum number of points a participant could receive from the one other who could give to them was identical to the maximum number of points a participant could receive from both others in the other two conditions.⁴

[Table 1 about here]

In all conditions, participants were able to see their own profits at the end of each round. Then they proceeded to the next round within the same triad until the end of the phase. When reading the study instructions for the repeated exchange task, participants saw a black-and-white exchange structure corresponding to the condition of their session with nodes with exemplar labels of A, B, and C (see Table 1 for the examples we showed in the study). They were told that, in the actual task, they would see a similar image, but that the image would contain information about the actual others with whom they were assigned to interact: both their actual randomly assigned letter identifier (which all respondents were given at the start of the study), and a color indicating the identity of each interactant as a Klee or Kandinsky, as shown in Figure 1. In the two “known identities” conditions that we analyze in the present study⁵, participants saw both letter identifiers *and* colors distinguishing each node. Nodes of actors with Klee identities were blue, and with Kandinsky identities were green.

⁴ We controlled for whether the generalized exchange condition was LTY or PA in all cooperation models, and conducted robustness tests demonstrating that the two generalized exchange conditions did not significantly differ from each other.

⁵ The analysis sample for the present study excludes decisions made by participants following the “unknown identities” condition.

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

When the exchange task began, participants in the *homogeneous identities* condition were placed in an exchange structure with two others who shared their Klee (or Kandinsky) identity, and were resultingly embedded in a triad with either all Klee (all blue) or all Kandinsky (all green) nodes. Participants in the diverse identities condition were embedded in an exchange structure with a mix of blue and green nodes to denote a triad with both Klee and Kandinsky participants⁶.

Each study session took place entirely online. If a participant was unable to progress through the session due to a technical issue or went idle for more than two minutes during the real-time decision-making task, the study program filled in for them by deploying a tit-for-tat exchange strategy for that round onward. This approach prevented the study from being cancelled mid-session. At the end of the phase (i.e., after 18 rounds of decision-making), participants completed two additional activities. First, they answered questions about their “experience with the two others to whom [they] were connected” to measure affective attachment toward the triad (Lawler et al., 2008). There were five items (*bad-good*, *detached-attached*, *disloyal-loyal*, *disconnected-connected*, and *negative-positive*) each measured on a nine-point scale.

Second, they participated in a one-shot Trust Game with participants that they were (accurately) told were *not* the two participants with whom they had just interacted in the previous decision-making (i.e., exchange) task. We describe the mechanics of this phase-ending task below.

⁶ Our study used two identities (Klee or Kandinsky), but participants interacted in triads, and we did not use deception in the sense that people were interacting with real others in real time. Thus, by necessity (and as shown in Figure 1), participants in diverse identities triads were paired with either one or two different-identity alters. This is because, in any given diverse triad, one participant was paired with two different-identity others (thus, ego was in the “minority” in the triad), while, from the other two participants’ perspectives, they were paired with one same-identity and one different-identity other (in the “majority”). Importantly, all participants interacted, in some way, with at least one different-identity other.

Manipulating Partner Identity in the One-Shot Trust Game

After completing the repeated exchange task and affective attachment measures in each phase, participants completed a one-shot Trust Game with an unmet interaction partner who possessed either an in-group or out-group social identity, signaled by node color (Figure 1). They were placed into the role as “truster”⁷: participants were allocated 300 MUs, and asked how many of them they wished to send to a “trustee” (called “returner” in the study instructions; the truster role was referred to as the “sender”). All MUs sent to the trustee tripled in value. The trustee would decide what proportion of MUs they would like to return to the truster. Participants were not informed of the amount their partner returned until the study was complete.

After three phases of repeated exchange (one each in the *homogeneous*, *diverse*, and *unknown identities* within-subject conditions) and one-shot trust decisions, the study was complete. Participants provided responses to open-ended questions about their perceptions of the study, including whether they suspected deception on the part of the experimenters⁸. Participants then read a debriefing screen, and were sent an Amazon e-gift card ranging in value from \$15-\$20 based on the number of MUs they accumulated⁹.

RESULTS

Analytic Approach

⁷ After the trust decision was made, the participant would then be placed in the trustee role, and be tasked with making a *trustworthiness decision* by deciding on a proportion of MUs sent to them by a different truster (i.e., another new interaction partner), which had tripled in value, to return. For simplicity, analyses in the present study focus entirely on trust decisions made by participants.

⁸ Out of the 240 participants in our sample, 28 (11.7 percent) reported some form of suspicion about whether they were truly interacting with real others (none reported suspicion that the Klee and Kandinsky feedback was randomized).

⁹ We standardized payments across conditions, so that those in the standard version of generalized exchange did not earn less than those in the other exchange structure conditions.

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

Studies leveraging the ATSE typically conceive of exchange structure as an independent variable that kickstarts an endogenously related set of outcomes: exchange frequency, affective attachment, and future behavior toward actors within the exchange structure (see Lawler et al., 2008, 2009, 2015). We build on this well-established finding by situating the entire endogenous process as our independent variable and measuring its impact on future behavior toward unmet others in non-repeated interactions. Put differently, we evaluate whether characteristics of prior repeated exchange can influence trust *beyond* the exchange structures themselves.

As such, our results begin with a descriptive summary of how cooperation decisions and downstream affective attachment varied across prior repeated exchange network contexts that are randomized in terms of their structure and identity composition. These two manipulations functionally assigned participants to pre-trust-game contexts with (1) different degrees of ongoing cooperation and positive emotion, and (2) differences in the opportunity to directly benefit from out-group exchange partners.

Respondents in different exchange structures had different number of exchange decisions to make each round (in reciprocal exchange, participants made two distinct decisions to exchange with each of their two alters; in productive and generalized exchange, participants made one decision to transfer MUs to the joint fund or to the one other with whom they could exchange, respectively). Thus we identify a single “focal beneficiary” (see also Harrell & Quinn, 2023) and measure cooperation as the mean of a set of dichotomous decisions: whether a participant gave to the focal beneficiary in each round of a phase of repeated exchange preceding a trust game. The focal beneficiary is the joint fund in productive exchange structures; the only

other to whom the participant was able to give in generalized exchange; and the other situated counterclockwise relative to a participant in reciprocal exchange¹⁰.

We then turn to a set of questions about the impact of prior network exchange conditions on subsequent trusting behavior toward similar and dissimilar partners. We measured trust after each 18-round repeated network exchange phase of the study. During this end-of-phase trust game, participants made a trust decision (i.e., opting to send any amount between 0 to 300 MUs to a “returner” partner) toward an unmet partner. The distribution of our dependent variable informs our approach. Participants could send any integer quantity of MUs between 0 and 300, but they often sent MU quantities that ended in zero. Of the 463 trust decisions in the analysis sample, the nine most common MU quantities sent (300, 150, 100, 200, 0, 50, 250, 75, and 10, respectively) accounted for 435 (93.9%) of the observed values. Moreover, a plurality of trust decisions ($n = 221$) involved sending all 300 MUs to the unmet partner, and a non-trivial quantity of trust decisions ($n = 22$) involved sending 0 MUs to the unmet partner. The resulting outcome – proportion of MUs sent – is both quasi-categorical and zero- one-inflated. This distribution makes a conventional mixed effects GLM modeling strategy an inappropriate fit. To account for these distributional properties, we construct a set of beta inflated regression models that estimate the proportion of MUs sent as a function of our treatment conditions and allow for inferences about the mean of these proportions. We estimate the effects of the type of exchange structure at time 1, whether the exchange structure was diverse or homogeneous at time 1, whether there was

¹⁰ This decision follows that of Harrell and Quinn (2023). It is worth noting that it results in the removal of half of the cooperation decisions made within reciprocal exchange structures but allows us to analyze one binary outcome in each round across all three exchange structure types. These analyses are less relevant to the current project, however, which is centered on decisions made in the Trust Game. Cooperation decisions results are only briefly summarized here.

a shared salient identity between the trust game interaction partner at time 2, and interactions between these variables on trust decisions.

Cooperation and Emotional Commitment Following Repeated Exchange

We first describe the differences in cooperation *within* the exchange structure at time 1, caused by our manipulations of exchange structure and diversity. Our treatments are endogenously generated contexts produced via random assignment: the two experimental factor manipulations expose participants to different sets of endogenous conditions that cause variation in the degree of emergent pre-trust-game cooperation and positive emotion – and partner behavior observability – within each structure. We used multilevel generalized linear models to predict the likelihood of exchanging (versus not exchanging) with the focal beneficiary by giving (versus keeping) one’s MUs (see also Harrell & Quinn, 2023). The models included random intercepts to account for the dependencies present in the data: rounds were nested in multiple phases, nested in participants in our study sessions.

Results revealed that, all else equal, productive exchange structures yielded the highest levels of exchange, followed by reciprocal exchange, followed by reciprocal exchange, followed by generalized exchange. These findings are in line with theoretical predictions and empirical tests of the ATSE (Lawler, 2001; Lawler et al., 2008). Supporting the interdisciplinary scholarship on ingroup biases, this model also suggested that (all else equal) participants were more likely to transfer their resources in homogeneous identities triads. The diversity of the triad did not interact with the exchange structure, suggesting that both structure and ingroup biases operate simultaneously and independently of each other to predict the rate of successful

exchange in the network. Figure 2 shows cooperation patterns during the exchange task, by our experimental conditions.

[Figure 2 about here]

Also as predicted by the ATSE, we observed similar patterns for affective attachment. Emotional commitments to the unit were strongest in productive exchange, followed by reciprocal and generalized exchange. They were also higher in homogeneous identities triads, and exchange structure and triad diversity did not interact with each other. Controlling for the proportion of giving that happened during the exchange task made the effects of exchange structure and triad diversity nonsignificant, suggesting, as does the ATSE, that exchange structures yield different levels of cooperation which, in turn, generate different levels of emotional commitment to the unit.

Prior Repeated Exchange Conditions, Future Trust, and Out-Group Partners

We next ask whether the degree to which a participant trusts an unmet partner in a one-shot trust decision at time 2 depends on the structure of a participant's repeated exchange network at time 1, the extent to which the participant was exposed to diverse others within the time 1 exchange network, and whether the (previously unmet) time 2 partner is a member of the ingroup or outgroup. Just like the cooperation decisions within the exchange network, the data from the Trust Game is nested within participants. They completed a trust decision after both the homogeneous identities *and* the diverse identities phases of the study. We account for these

dependencies, as with the cooperation decisions, by including random intercepts at the participant level in our models for Trust Game behavior.

Table 2 reports the results of the beta-inflated regression analysis. The models estimate the effect of the treatment factors on the log odds of change in the expected proportion of MUs a participant sent to their partner in a trust decision. Here, significant coefficient values greater than or less than 0 indicate an increase or decrease in this outcome based on a one-unit change in the parameter – or a change in treatment status relative to the reference condition – respectively. The reference category for prior exchange structure is the generalized exchange condition; the reference category for prior repeated exchange diversity is the homogeneous identities condition.

[Table 2 about here]

We begin with the simplest model (Model 1), which contains only the main effects of our three manipulations: prior exchange structure (generalized, reciprocal, or productive), diversity of the prior exchange structure (homogeneous or diverse), and the unmet trustee's identity in relation to the participant's (ingroup or outgroup). The model also controls for the phase in which the decision was made: first, second, or third¹¹. The goal is to assess whether any of our study conditions had main effects on downstream trust at all, and if so, in what direction, before probing how our conditions interacted. Results show, first, that the *diversity* of prior exchange structures did not have a significant impact on subsequent trust toward an unmet partner.

However, both 1) an actor's prior exchange structure and 2) the identity of their downstream

¹¹ While there were three phases in the initial study, we only analyze behavior from two, but whether the homogeneous identities phase and diverse identities phase came first, second, or third was randomly assigned within sessions.

partner, impacted trust. We discuss each of these effects in turn. Then, we consider interactions that might qualify the main effects (Model 2).

First, the diversity of prior exchange structures did not appear to impact subsequent trust toward an unmet partner ($B = .072, p = .231$). One possibility is that diversity of the previous exchange structure only mattered for trust toward unmet partners who were in the *outgroup*, but not for those unmet partners in the *ingroup*. However, in a follow-up model (not shown) we interacted diversity of the prior exchange structure with whether the trust game partner was in the *outgroup* (vs. *ingroup*), and the interaction was also not significant ($B = -.02, p = .83$). Because diversity of the prior exchange structure did not appear to impact trust, we drop it from further discussion here, though we retained it as a control variable.

We did find evidence that whether the time 2 trustee was a member of the focal actor's *ingroup* or *outgroup* mattered for trust. In line with the previous scholarship on *ingroup* biases, participants extended less trust toward *outgroup* trustees, though the results were marginally significant in Model 1 ($B = -.110, p = .068$). We also found evidence that prior repeated exchange *structure* impacts subsequent one-shot trust behavior. Model 1 specifically suggests that participants trusted more when they had previously been embedded in a reciprocal exchange structure, compared to a generalized exchange structure ($B = .465, p < .001$). Trust following embeddedness in productive exchange fell in between that of reciprocal and generalized exchange, and did not significantly differ from trust following generalized exchange ($B = .099, p = .147$). This initial finding seems to stand in contrast with our results about the effect of exchange structure on ongoing cooperation rates within the networks that preceded the trust decisions in our experiment. As a result, we probed whether the different forms of exchange structure *differentially* impacted subsequent trust when the trustee was in the *ingroup* or *outgroup*

– or whether certain exchange structures may attenuate or exacerbate out-group bias in trust decisions.

The results (Model 2) are also shown in Table 2. First, it is worth noting that this model produces a better fit (based on the reduction in the AIC). Additionally, this model reveals that behavior toward outgroup versus ingroup trustees differs based on prior exchange structure (or, put differently, that reduced levels of trust to outgroup trustees are moderated by the structure of one's prior repeated exchange network). First consider the main effects. There is now a main effect for productive exchange, suggesting that, when the trustee is in one's ingroup, those participants who had previously experienced productive exchange were more trusting ($B = .224$, $p = .019$). The main effect on trust for reciprocal exchange, while positive, becomes less than the main effect of productive exchange and is not significantly different from generalized exchange ($B = .120$, $p = .224$). The main effect of the trustee's identity becomes significant ($B = -.171$, $p = .045$; in Model 1, it was marginally significant), suggesting that, for generalized exchange, participants extended less trust to outgroup trustees.

Now consider the interactions terms. Compared to generalized exchange, participants' trust toward outgroup trustees declines (compared to ingroup trustees) when they were previously embedded in a productive exchange structure ($B = -.278$, $p = .031$). Thus, while there are benefits to having been in a productive exchange task when the trustee is an ingroup member, these benefits are nullified when the trustee is in the outgroup. Most surprisingly, actors are *more* trusting of outgroup trustees when they are embedded in a reciprocal exchange structure prior to their one-shot trust decision ($B = .870$, $p < .001$). We estimated the predicted proportion of MUs sent by prior exchange structure and trustee out-group identity status from Model 2 and visualize

this result in Figure 3 (after fixing prior exchange structure diversity to zero¹²) to illustrate the interaction effect. Figure 3 emphasizes the impact of prior exchange structure as a moderator of trust toward outgroup trustees in downstream one-shot decisions. When paired with a previously unmet outgroup trustee after being previously embedded in reciprocal exchange, participants trust their outgroup trustee far more than those previously embedded in generalized or productive exchange structures.

[Figure 3 about here]

One possibility for this surprising result is that reciprocal exchange structures mirror the direct and reciprocal structure of a dyadic one-shot trust game. As such, participants previously embedded within reciprocal exchange structures not only obtain the heightened positive emotions relative to those in generalized exchange (Lawler et al., 2008; Harrell & Quinn, 2023), but they are also able to directly observe the behavior of specific others in ways that benefit them, and might come to anticipate repeated exchange with fairly high frequencies of success when making dyadic decisions about whether to exchange with others – regardless of their social identity. Moreover, note again that when the trustee is an *ingroup* member, those in productive exchange structures are more trusting (main effect of productive exchange; $B = .224, p = .019$). The result is that, while productive exchange may yield the highest levels of cooperation with unmet ingroup trustees, productive exchange also leads to the largest degree of ingroup bias – greater trust toward ingroup members compared to the other two forms of exchange, and lower trust toward outgroup members compared to the other two forms of exchange. This finding both

¹² Setting this value to one instead of zero did not meaningfully alter the resulting estimates reported in our tables or visualizations.

complements and challenges the logic that may be implied by the ATSE: while our results show that exchange structures do impact trust as a form of future giving behavior (here, toward an unmet alter), the effect size does not fully correspond to the degree of task jointness embodied by each prior exchange structure. We discuss this and other issues in more detail next.

CONCLUSION & DISCUSSION

In this paper, we assess whether the structure and diversity of prior repeated exchange settings affect future behavior in the form of decisions to trust *outside* of the structure, toward unmet interactants. We also ask whether the effects of prior repeated exchange qualities on future one-shot trust behavior are different when actors share, or do not share, a salient social category with their interaction partner. Our results show that the structure of an actor's prior repeated exchange network has both main effects on subsequent trust, and interaction effects when their decision involves an in-group versus out-group trustee. Previously experiencing productive exchange tends to heighten trust, but it also heightens in-group preference in trust decisions by increasing the difference in the amount of MUs allocated to in-group versus out-group trustees. By contrast, previously experiencing reciprocal exchange appears to increase trust toward unmet out-group trustees. In our study, participants who were previously embedded in reciprocal exchange structures – regardless of whether this prior structure contained out-group exchange partners – trusted unmet out-group partners in future one-shot decisions significantly more than actors previously embedded in prior productive or generalized repeated exchange structures.

Our findings suggest that the structure of an actor's prior relational context can impact their future behavior outside of that structure, and that the structure of previous interactions can increase or reduce out-group biases in the form of trusting a new interaction partner. To some

extent, this result is surprising in the context of the ATSE. After all, the ATSE predicts that the degree of task jointness involved in an exchange structure affects the extent to which actors participate in future giving behavior – but in prior work, this future behavior was operationalized as prosocial action taken toward others that were a *part of* the exchange structure in which the actor of interest had been embedded (e.g., Lawler et al., 2008). By operationalizing future behavior as behavior enacted toward (1) not-previously-met trustees who are (2) either in-group or out-group actors, we learn that productive exchange still produces the most positive main effect on the giving behavior of future trust, but that the story changes when the truster is categorically dissimilar to the trustee. We also learn that out-group bias in trust is attenuated when an actor was previously embedded in repeated reciprocal exchange.

More evidence is needed to both replicate and identify the mechanisms underpinning our novel results. We suspect that the mechanism lies in both the positive emotions that prior exchange structures can generate and the capacity for an actor to benefit in direct and observable ways from others in their prior exchange structure. Repeated reciprocal exchange induces more positive emotions (compared to generalized exchange) *and* might allow participants to link these positive emotions to the notion of a dyadic exchange partner, compared to productive exchange. In productive exchange, emotions are also positive—indeed, *more* positive than in reciprocal exchange according to the ATSE—but they are also generated toward a *unit* rather than any given individual. The positive emotions and dyadic exchange that occur in reciprocal exchange might combine in such a way that people may attribute these emotions to their one-shot trust game dyadic partner. After all, both reciprocal exchange decisions and decisions to extend trust toward another actor entail direct and implicitly dyadic decisions. Note that, while we set out to examine prosocial behavior at time 2 in a different context compared to any of the time 1

exchange structures, it is possible that reciprocal exchange structures yielded higher outgroup trust merely because the structure of these decisions are more similar. It is possible that if our time 2 behavior was, instead, cooperation in a public goods dilemma with new interaction partners, we would have observed that productive exchange (rather than reciprocal exchange) yielded higher levels of cooperative time 2 behavior.

At the same time, it is noteworthy that prior reciprocal exchange only increased future trust relative to other exchange structures when the trust decision was made toward an *outgroup* alter, and that this was the case regardless of the diversity or homogeneity of the prior exchange structure. We hope to reproduce and explore the social psychological mechanisms that underlie this pro-outgroup effect of reciprocal exchange in future work. We suspect that our manipulation of social identities, and the duration of our repeated exchange phases, may have not been potent or long enough (respectively) to produce very pronounced effects. We plan to explore these follow-up questions about mechanism identification and manipulation strength and encourage other researchers to expand on these exploratory experimental findings.

In doing so, we hope to motivate scholars who study the structural determinants of status- or identity-based social inequality, or the causes and consequences of inter-group interactions, emotions, and attitudes, to develop research questions and construct study designs that involve the manipulation of (or at least leveraging theoretically meaningful variation in) *both* the structure and identity composition of repeated interaction between individuals. Our prior work shows that structural arrangements clearly affect the degree of cooperation that actors will direct toward people who are categorically dissimilar to themselves – implying that diverse teams can be constructed in ways that produce higher or lower levels of intergroup cooperation and cohesion (Harrell & Quinn, 2023). In this study, we find that prior exchange structure also affects

the degree of *future* trust behavior toward new out-group others encountered immediately after a phase of repeated network exchange – that the structure of prior environments can shape future intergroup behavior. Our findings imply that orientations toward categorical groups are malleable, at least in proximity to a manipulation like repeated exchange. Such an implication is of consequence to researchers interested in prejudice reduction (e.g., Pettigrew, Tropp, Wagner, & Christ, 2011) or affective meaning change (e.g., Quinn, Freeland, Rogers, Hoey, Smith-Lovin, 2023) about social identities, the role of organizations in creating structures that heighten or minimize the legitimacy and enactment of gendered or racialized inequalities (Kanter, 1977; Ray, 2019), and the impact of prior structural circumstances on future interactions between groups (Boin et al., 2021; Kao et al., 2019).

REFERENCES

- Aksoy, O. (2015). Effects of heterogeneity and homophily on cooperation. *Social Psychology Quarterly*, 78(4), 324–44.
- Aksoy, O. (2019). Crosscutting circles in a social dilemma: Effects of social identity and inequality on cooperation. *Social Science Research*, 82, 148–63.
- Allport, G. W. (1954). *The Nature of Prejudice*. Cambridge: Addison-Wesley.
- Balliet, D., Wu, J., & De Dreu, C. K. W. (2014). Ingroup favoritism in cooperation: A meta-analysis. *Psychological Bulletin*, 140(6), 1556–81.

- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10, 122-142.
- Billig, M., & Tajfel, H. (1973). Social categorization and similarity in intergroup behaviour. *European Journal of Social Psychology*, 3(1), 27–52.
- Blau, P. M. (1964.) *Exchange and Power in Social Life*. New York, NY: Wiley.
- Boin, J., Rugar, M., Graf, S., Neji, S., Spiegler, O., & Swart, H. (2021). The generalization of intergroup contact effects: Emerging research, policy relevance, and future directions. *Journal of Social Issues*, 77(1), 105–31.
- Duch, M. L., Grossmann, M. R. P., and Lauer, T. (2020). Z-Tree Unleashed: A Novel Client-Integrating Architecture for Conducting z-Tree Experiments over the Internet. *Journal of Behavioral and Experimental Finance*, 28, 100400.
- Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2):171–78.
- Foddy, M., Platow, M. J., & Yamagishi, T. (2008). Group-based trust in strangers. *Psychological Science*, 20(4), 419–22.
- Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481–510.
- Harrell, A., & Quinn, J. M. (2023). Shared identities and the structure of exchange distinctly shape cooperation. *Social Forces*, soad011.
- Heise, D. R. (2010). *Surveying Cultures: Discovering Shared Conceptions and Sentiments*. Hoboken, NJ: Wiley.
- Hobbes, T. (1994). *Leviathan: With Selected Variants from the Latin Edition of 1668*. Hackett Publishing.
- Kanter, R. M. (1977). *Men and Women of the Corporation*. New York, NY: BasicBooks.
- Kao, G., Joyner, K., & Balistreri, K. S. (2019). *The Company We Keep: Interracial Friendships and Romantic Relations from Adolescence to Childhood*. New York, NY: Russell Sage Foundation.
- Kauff, M., Kotzur, P.F., Van Assche, J., Schäfer, S. J., Zalk, M. H. W., & Wagner, U. (2023). A longitudinal test of secondary transfer effects of negative intergroup contact and mediating processes. *European Journal of Social Psychology*, 53(6), 1172–90. doi: 10.1002/ejsp.2966.

REPEATED EXCHANGE NETWORK STRUCTURES AFFECT FUTURE TRUST

- Kossinets, G., & Watts, D. J. (2009). Origins of homophily in an evolving social network. *American Journal of Sociology*, 115, 405-450.
- Lawler, E. J. (2001). An affect theory of social exchange. *American Journal of Sociology*, 107(2), 321–52. doi: 10.1086/324071.
- Lawler, E. J., Thye, S. R., & Yoon, J. (2000). Emotion and group cohesion in productive exchange. *American Journal of Sociology*, 106, 616-657.
- Lawler, E. J., Thye, S. R., & Yoon, J. (2006). Commitment in structurally enabled and induced exchange relations. *Social Psychology Quarterly*, 69(2), 183–200. doi: 10.1177/019027250606900204.
- Lawler, E. J., Thye, S. R., & Yoon, J. (2008). Social exchange and micro social order. *American Sociological Review* 73(4), 519–42. doi: 10.1177/000312240807300401.
- Lawler, E. J., Thye, S. R., & Yoon, J. (2009). Social commitments in a depersonalized world. doi: 10.1354/books/unregistered/9781610446600.
- Lawler, E. J., Thye, S. R., & Yoon, J. (2015). *Order on the Edge of Chaos: Social Psychology and the Problem of Social*. New York, NY: Cambridge University Press.
- Lawler, E. J., & Yoon, J. (1996). Commitment in exchange relations: test of a theory of relational cohesion. *American Sociological Review*, 61(1), 89. doi: 10.2307/2096408.
- Lawler, E. J., & Yoon, J. (1998). Network structure and emotion in exchange relations. *American Sociological Review*, 871-894.
- Lersch, P. M. (2023). Change in personal culture over the life course. *American Sociological Review*, 88(2), 220–51. doi: 10.1177/00031224231156456.
- Marsden, P. V. (1987). Core discussion networks of americans. *American Sociological Review*, 52(1), 122. doi: 10.2307/2095397.
- McKeown, S., & Dixon, J. (2017). The ‘contact hypothesis’: critical reflections and future directions. *Social and Personality Psychology Compass*, 11(1). doi: 10.1111/spc3.12295.
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: homophily in social networks. *Annual Review of Sociology*, 27(1), 415–44. doi: 10.1146/annurev.soc.27.1.415.
- Melamed, D., Simpson, B., Harrell, A., Munn, C. W., Abernathy, J. Z., & Sweitzer, M. (2020). Homophily and segregation in cooperative networks. *American Journal of Sociology*, 125(4), 1084–1127. doi: 10.1086/708142.
- Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5), 751.

- Pettigrew, T. F., & Tropp, L. R. (2013). Does intergroup contact reduce prejudice? Recent meta-analytic findings. In *Reducing Prejudice and Discrimination*. New York, NY: Psychology Press. doi: 10.4324/9781410605634.
- Pettigrew, T. F., Tropp, L. R., Wagner, U., & Christ, O. (2011). Recent advances in intergroup contact theory. *International Journal of Intercultural Relations*, 35(3), 271–80. doi: 10.1016/j.ijintrel.2011.03.001.
- Quinn, J. M., Freeland, R. E., Rogers, K. B., Hoey, J., & Smith-Lovin, L. (2023). How cultural meanings of occupations in the U.S. changed during the Covid-19 pandemic. *American Behavioral Scientist*, 67(1), 125-147. doi: 10.1177/00027642211066041
- Ray, V. (2019). A theory of racialized organizations. *American Sociological Review*, 84(1), 26–53. doi: 10.1177/0003122418822335.
- Ridgeway, C. L., Boyle, E. H., Kuipers, K. J., & Robinson, D. T. (1998). How do status beliefs develop? The role of resources and interactional experience. *American Sociological Review* 63(3):331. doi: 10.2307/2657553.
- Ridgeway, C. L., & Correll, S. J. (2006). Consensus and the creation of status beliefs. *Social Forces*, 85(1), 431–53. doi: 10.1353/sof.2006.0139.
- Ridgeway, C. L., & Erickson, K. G. (2000). Creating and spreading status beliefs. *American Journal of Sociology*, 106(3), 579–615. doi: 10.1086/318966.
- Romano, A., Balliet, D., Yamagishi, T., & Liu, J. H.. (2017). Parochial trust and cooperation across 17 societies. *Proceedings of the National Academy of Sciences*, 114(48), 12702–7. doi: 10.1073/pnas.1712921114.
- Schäfer, S. J., Kauff, M., Prati, F., Kros, M., Lang, T., & Christ, O. (2021). Does negative contact undermine attempts to improve intergroup relations? Deepening the understanding of negative contact and its consequences for intergroup contact research and interventions. *Journal of Social Issues*, 77(1), 197–216. doi: 10.1111/josi.12422.
- Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–86. doi: 10.1080/0022250x.1971.9989794.
- Serpe, R. T., & Stryker, S. (2011). The symbolic interactionist perspective and identity theory. In *Handbook of Identity Theory and Research*, 225–48. New York, NY: Springer New York. doi: 10.1007/978-1-4419-7988-9_10.
- Simpson, B. (2006). Social identity and cooperation in social dilemmas. *Rationality and Society*, 18(4), 443–70.
- Simpson, B., & Willer, R. (2015). Beyond altruism: sociological foundations of cooperation and prosocial behavior. *Annual Review of Sociology*, 41(1), 1–21.

Tanis, M., & Postmes, T. (2005). A social identity approach to trust: Interpersonal perception, group membership and trusting behaviour. *European Journal of Social Psychology*, 35(3), 413–24.

Turner, J. H. (1988). *A Theory of Social Interaction*. Stanford, CA: Stanford University Press.

Whitham, M. M. (2018). Paying it forward and getting it back: The benefits of shared social identity in generalized exchange. *Sociological Perspectives*, 61(1), 81–98.

Yamagishi, T., Jin, N., & Kiyonari, T. (1999). Bounded generalized reciprocity: Ingroup boasting and ingroup favoritism. *Advances in Group Processes*, 16(1), 161–97.

Yamagishi, T., & Kiyonari, T. (2000). The group as the container of generalized reciprocity. *Social Psychology Quarterly*, 63(2), 116.

Yamagishi, T., & Mifune, N. (2009). Social exchange and solidarity: In-group love or out-group hate? *Evolution and Human Behavior*, 30(4), 229–37.

FIGURES

Figure 1. Variation of the three manipulations within each phase of the study.

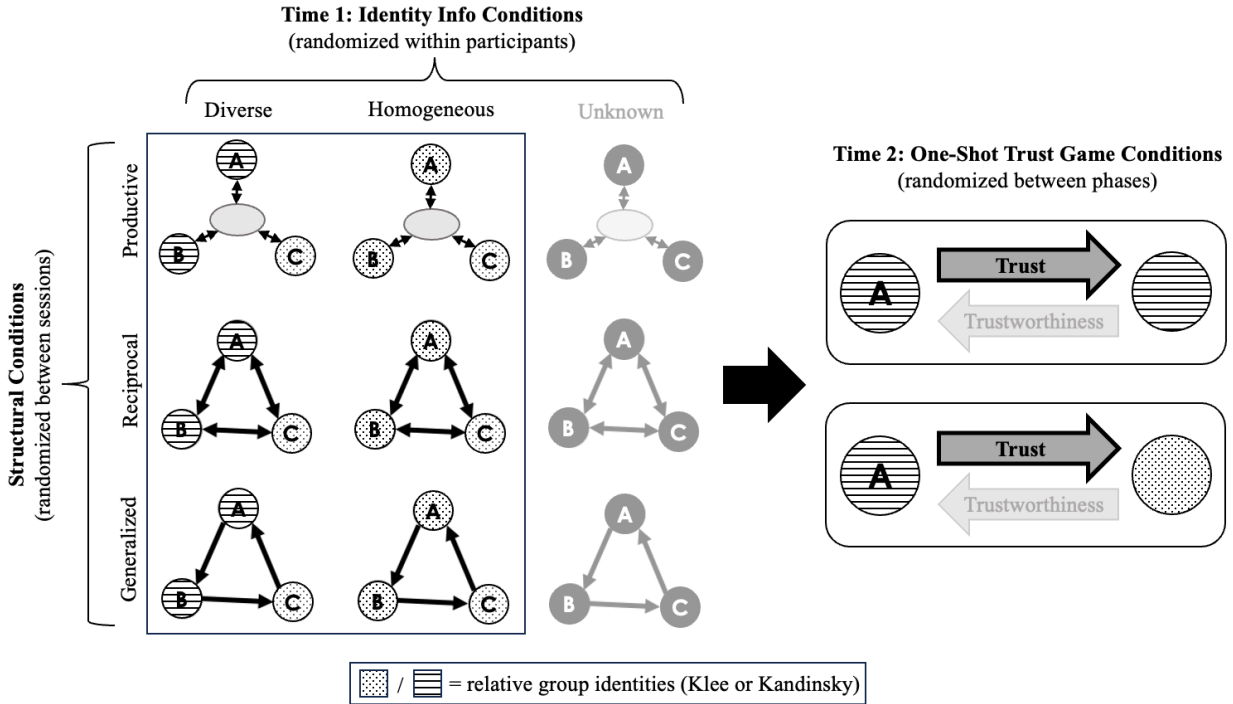


Figure 2. Giving to the focal beneficiary over time, by the key study manipulations. Reproduced from Harrell and Quinn 2023.

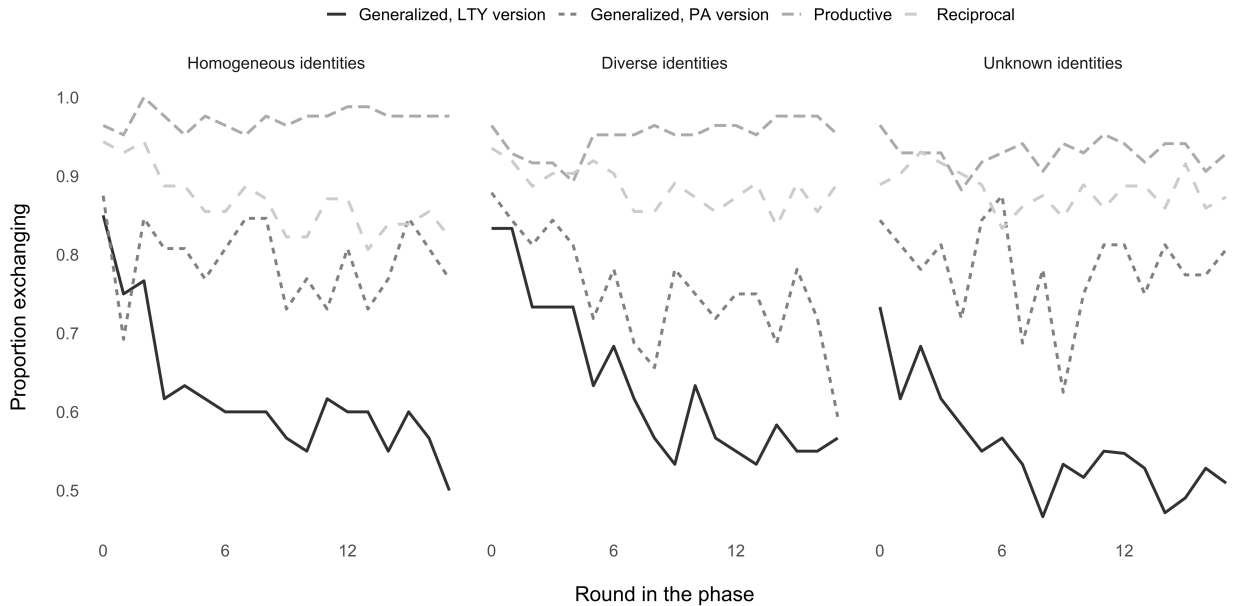



Figure 3. Expected Proportion of MUs Sent to Trust Partner at Time 2 by Exchange Structure at Time 1



TABLES

Table 1. Decision-making in productive, reciprocal, and generalized exchange in Lawler, Thye, and Yoon (2008) and the current study. Reproduced from Harrell and Quinn 2023.

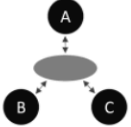
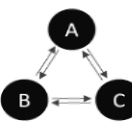
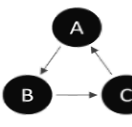
		Productive	Reciprocal	Generalized
				
When ego gives, who is benefited?		Ego + 2 alters	1 or 2 alters	1 alter
From where may ego receive benefits?		The joint fund	1 or 2 alters	1 alter
Ego's profit if...	Ego gives, both alters don't	267 (800/3 from the joint fund)	0	0
	No one gives	300	300	300
	Ego doesn't give, both alters do	1100 (2400/3 from the joint fund + 300)	1900 (800 from each of 2 alters + 300)	1900 (<i>PA version</i> , 1600 from 1 alter + 300) or 1100 (<i>LTY version</i> , 800 from 1 alter + 300)
	All three give	1600	1600	1600 (<i>PA version</i>) or 800 (<i>LTY version</i>)

Table 2. Beta Inflated μ Model Estimating Log Odds of the Mean Proportion of MUs Sent

μ Model Coefficients	Model 1 (Manipulation Main Effects)			Model 2 (Main + Interaction Effects)		
	β	SE	P-Value	β	SE	P-Value
Prior Exchange: Prod.	0.099	0.068	0.147	0.224	0.095	0.019 *
Prior Exchange: Recip.	0.465	0.080	<0.001 ***	0.120	0.099	0.224
Prior Exchange: Diverse IDs	0.072	0.060	0.231	0.061	0.057	0.287
Trust Game: Out-Group Partner	-0.110	0.060	0.068 .	-0.171	0.085	0.045 *
Phase 2	-0.282	0.073	<0.001 ***	-0.247	0.069	<0.001 ***
Phase 3	-0.402	0.075	<0.001 ***	-0.379	0.071	<0.001 ***
TG: Out-Group x PE: Prod.				-0.278	0.128	0.031 *
TG: Out-Group x PE: Recip.				0.870	0.153	<0.001 ***
(Intercept)	-0.1153	0.0721	0.111	-0.0893	0.0806	0.2684
	Global					
	Deviance	AIC	SBC			
Model 1:	343.36	597.69	1123.87			
Model 2:	316.99	578.65	1119.97			

Notes: $p < .001$: ***; $p < .01$ = **; $p < .05$ = *; $p < .1$ = .; both models include a random intercept to account for variance between phases within sessions. Coefficients represent the log odds of positive change in the expected proportion of MUs a participant sent to their partner when making a trust decision.