

This is a “preproof” accepted article for *Evolutionary Human Sciences*.
This version may be subject to change during the production process.
10.1017/ehs.2025.8

Observational evidence that economic reciprocity pervades self-organized food co-operatives

Taylor Z. Lange^{1,2,3*} and Timothy M. Waring^{3,4}

¹Margaret Chase Smith Policy Center, University of Maine, Orono, ME 04469,
United States

²Ecology and Environmental Sciences, University of Maine, Orono, ME 04469,
United States

³School of Economics, University of Maine, Orono, ME 04469

⁴Mitchell Center for Sustainability Solutions University of Maine, Orono, ME
04469, United States

*Corresponding author: taylor.z.lange@maine.edu

January 31, 2025

This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re- use, distribution and reproduction, provided the original article is properly cited.

Abstract

Evolutionary scientists argue that prosociality has been central to human ecological success. Theoretical models and behavioral experiments have found that prosociality, and cooperation in particular, is conditional and context dependent, that individuals vary in their propensity to cooperate, and that reciprocity stabilizes these behaviors within groups. Experimental findings have had limited validation with observations of behavior in natural settings, especially in organizational contexts. Here, we report in situ measurements of collective action where reciprocity is abundant in organizations embedded in a cash economy. We study small ‘food clubs’ where members share bulk purchases and are considered to be heavily dependent on cooperation. We take advantage high-resolution data of economic interactions for 1,528 individuals across 35 clubs over a combined 107 years of operation. We develop a network method to detect different directional and temporal forms of economic reciprocity, and statistically classify individual behavioral types akin to those in experiments. We find abundant direct reciprocity, supplemented by indirect reciprocity, and that members of most clubs can be identified as consistent reciprocators. This study provides initial observational evidence that economic reciprocity may be more abundant in real-world settings, sharpening the findings of the behavioral study of cooperation and contributing to the more naturalistic study of reciprocity and prosociality.

Keywords: reciprocity, cooperation, coordination, collective action, co-operatives, food buying clubs

1 Introduction

Humans and other social organisms face a variety of social dilemmas where the interests of the group may be at odds with those of the individual, and the human ability to solve these dilemmas is theorized to be central to our success as a species (Gowdy, 2021; Gowdy and Krall, 2016; Henrich, 2015). In collective action situations, where a group works towards a common goal, individuals are often required to act prosocially and disregard the incentive to free-ride off the efforts of others (Fischbacher and Gächter, 2010; Fischbacher et al., 2001; Gintis, 2009). Theoretical models, laboratory experiments, and field experiments have revealed several mechanisms that stabilize individual prosociality by minimizing cheating, with reciprocity accumulating the most substantial amount of scholarship (Axelrod and Hamilton, 1981; Nowak, 2006; Nowak and Sigmund, 2005; Ohtsuki et al., 2006; Rand et al., 2014; Trivers, 1971). While these findings have been replicated across cultural contexts (Apicella et al., 2012; Henrich et al., 2001; Henrich and Muthukrishna, 2021; Kocher et al., 2008), observations of the patterns of reciprocity in non-experimental, contemporary contexts is lacking, leaving current findings circumscribed. Here, we present a novel study of collective action in organizational networks using observational data to compare to experimental and theoretical findings.

Prosocial behaviors are those that benefit others, and they include coordination, cooperation, and altruism (Gintis, 2009; Henrich and Muthukrishna, 2021; Jensen, 2016). Coordination is the broadest category and it is defined as an action that benefits the individual and those around them (Gintis, 2009). Cooperation is generally understood as giving assistance to others in a manner that could still benefit them, but requires a personal cost (Henrich and Muthukrishna, 2021). Altruistic behaviors constitute a subset of

cooperation and are defined as lending assistance in a manner that is not personally beneficial and incurs a personal cost (Trivers, 1971; Wilson, 2015).

Reciprocity, or giving the same as what has been given to one's self, is a potent behavioral strategy that sustains within-group prosociality across many experimental and theoretical collective action dilemmas (Janssen et al., 2011; Nowak, 2006; Pfeiffer et al., 2005; Sandroni, 2000; Schmid et al., 2021). In economic games with repeated interactions between two individuals such as the iterated prisoners' dilemma, tit-for-tat direct reciprocity, i.e. doing as one's partner did in the previous round, is the most successful strategy for sustaining cooperation and altruism specifically (Axelrod and Hamilton, 1981; Roberts, 2008). In another version of the iterated prisoners' dilemma where partnerships change each round, pay-it-forward style indirect reciprocity allows cooperation to persist when (1) players can view their partner's previous decisions and when (2) reputation accumulates over time (Nowak and Sigmund, 2005; Roberts, 2008). In public goods scenarios where individuals interact with many group members, reciprocity takes a more general form of "conditional cooperation" where one gives the same as the rest of the group does on average (Fischbacher et al., 2001). However, free-riders in public goods experiments can drag the average donation down, causing conditional cooperation to decline over time (Fischbacher et al., 2001; Thöni and Volk, 2018). Under such circumstances, cooperative behavior needs additional supporting mechanisms such as interpersonal reward or punishment to coerce others into cooperation (Fischbacher and Gächter, 2010; Rand et al., 2009).

Many studies of conditional cooperation also suggest that individuals exhibit patterns of behavior that can be classified into three broad types: conditional cooperators (who cooperate contingent on the cooperation of the rest of the group), altruists (who always cooperate), and free-riders (who always defect) (Andreozzi et al., 2020; Fischbacher et al., 2001; Frey, 2017; Volk et al., 2012). These types are descriptions of individuals' sustained strategy profiles throughout iterated games (Fischbacher et al., 2001) and they have been robustly replicated in in-person (Andreozzi et al., 2020; Deversi et al., 2020; Kocher et al., 2008) and in online video games (Frey, 2017,1). Furthermore, types appear to be consistent when they participate in multiple experiments across time (Andreozzi et al., 2020; Volk et al., 2012), and emerge in studies across cultural contexts (Frey, 2017; Kocher et al., 2008).

While these findings have contributed greatly to theories of the evolution and mechanisms of sustained cooperation (Nowak, 2006), tests of these theories with granular naturalistic data has remained elusive yet necessary, as experimental conditions often differ considerably from those in the real world. For example, experimental subjects are often drawn from large anonymous populations such as university student pools in which peers are strangers with no prior affiliation (e.g. Fischbacher et al. 2001 or Andreozzi et al. 2020) or are elicited through online services (e.g. Rand et al. 2009 or Rand et al. 2014). Furthermore, the economic games played are game theoretically derived to have clearly differentiable cooperative or selfish decisions (Gintis, 2009). In contrast, most studies of real world social dilemmas, including those surrounding public and similar goods, find that the individuals involved often have prior close affiliations, and that the derived distinction between cooperation and coordination dilemmas can become blurred (Henrich and Muthukrishna, 2021; Ostrom, 1990; Wilson, 2015; Wilson et al., 2013). These differences between experimental settings lead to questions about the transferability of experimental results, such as: Do the rates of reciprocity observed in the lab translate to real work contexts, and are behavioral types actually observable in natural dilemmas or are they merely artifacts of experimental contexts?

The move to substantiate the experimental findings on reciprocity with naturalistic data is promising. Anthropological studies of hunter-gatherer societies have found observational evidence of reciprocity and cooperation in behaviors such as food sharing (Koster and Leckie, 2014), and there are examples of inter-organizational cooperation in market economies, including firm collusion Asker and Nocke (2021). For example Mao and Song (2021) observed that reciprocity plays a key role in analysts' propensity to underwrite securities between firms. Furthermore, Frey (2017,1) measured reciprocity and detected comparable behavioral types in an online video game, though this is still an emulation of a real-world environment.

A limiting factor of the study of naturalistic cooperation has been a paucity of behavioral data that parallels experimental settings. Here, we analyze a novel, granular dataset on shared bulk purchasing interactions that are similar to some experimental public goods settings to determine how reciprocity and behavioral types emerge in a real-world setting. The data consists of individuals extending mutual aid to ensure that their needs and wants are met and that the club continues to function. We identify economic reciprocity at degrees that exceed most experimental studies. And, we derive novel behavioral types that tend towards a stable reciprocal strategy.

2 Food Clubs

Food buying clubs are semi-formal organizations that collectively arrange wholesale food purchases from national and regional distributors (Hupper, 2019; Lange, 2022; Tremblay, 2017). Clubs range in size from 5-100 members, who consist of one or two coordinators that maintain distributor accounts, and general members who only participate in purchasing. Orders are placed on a weekly or monthly basis, and items range from individual or family sized goods to bulk items with more supply than one buyer generally needs.

These clubs constitute miniature, informal consumer co-operatives, as they are formed to take advantage of the discounts afforded by wholesale bulk purchasing and gain access to specialty goods that may be unavailable in local stores (Little et al., 2010). Food from local producers and niche items such as those labeled organic or fair trade are not always supplied by the traditional food system, particularly in rural communities, because demand has not reached a profitable scale for large chain grocers (Deller et al., 2017; Little et al., 2010). However, supply can be provided by a memberowned co-operative grocery store, where motivations are not exclusively profit driven and members collectively fulfill the missing market for the demand that does exist (Deller et al., 2017; Steinman, 2020). In the absence of a formal cooperative store, food buying clubs offer a similar opportunity for motivated individuals, as well as for those who wish to begin the process of forming a cooperative grocery store (Hupper, 2019; Lange, 2022; Little et al., 2010; Tremblay and Waring, 2015).

As co-operative organizations, food buying clubs rely on cooperation much more than a typical hierarchical organization. Co-operatives are defined by their democratic decision-making, member ownership, and reliance on cooperation (Birchall, 1997), and members of cooperatives tend to identify trust and reciprocity as integral to cooperative success (Pesämaa et al., 2013). Members of cooperatives have been found to exhibit higher levels of behavioral cooperation as elicited by the dictator game (Tremblay et al., 2019), but it is unclear whether and how that cooperation is maintained. Furthermore, while formal co-operative grocery stores can potentially rely on established rules to avert or solve social

dilemmas (Waring et al., 2022), food buying clubs often lack institutions initially, making reciprocity essential for early club success (Hupper, 2019).

Clubs facilitate supply by pooling their demand to buy goods in bulk, e.g. flats of cans or 25-50lb bags of dry goods like flour. In most instances, the demand of a single buyer is less than the minimum amount required by distributors, or "case size", so multiple members are required to buy-in and "co-purchase" the item. For a co-purchase to be successful, members must combine their buy-ins to match the case size, and doing so requires a negotiation of their preferences for all available goods. This dynamic can lead to a social dilemma such that if preferred amounts don't match case sizes exactly, members need to decide who will buy more or less so that the case is fulfilled.

Previous research of these clubs has shown that club members are exceptionally cooperative in behavioral games. Members' propensity for cooperation has been measured using a dictator game (DG) and a public goods game (PGG), where club members contributed unusually large fractions of their endowment (DG: 58%, PGG: 71%) (Hupper, 2019). These contribution levels substantially exceed those reported in the literature for equivalent games, as the next largest mean contributions are DG: 48% in Paciotti et al. (2011), and PGG: 57% in Apicella et al. (2012) and 58% in Henrich et al. (2001). Qualitative responses of our sample reported in Hupper (2019) corroborate an abundance of cooperation among active clubs, and stress when cooperation is not sufficient. Coordinators of some clubs mentioned that "Members need to contribute more," while in others said "For the most part members carry their weight," or "Everybody steps up."

2.1 The Co-Purchasing Dilemma

The relationship between what members want and what they end up buying can be described as a game. Imagine a co-purchasable item as a game played between a subset m of the club's M members ($m \in M$). The item is a case of X single shares, and each share costs price p . Shares can be aggregated into quantities of x shares ($x \leq X$) with an aggregate price of $x * p$. We imagine that members might only prefer a certain number of x shares for a given item, and otherwise prefer to keep their money. To illustrate, we define members' 'utility', or hedonic value (Fehr and Rangel, 2011), as a function of x and their exogenous level of money, or 'wealth' W :

$$U_m = U_m(x, W) \tag{1}$$

At $x = 0$ (no shares purchased), members have a base level of utility endowed by just their wealth, U_m^{base} . As shares x increase, members pay $p * x$ from their wealth, and utility deviates from U_m^{base} . This trade-off allows us to formally define when members have a preference for the item: members prefer the item at share quantity x IFF the utility at that quantity is greater than or equal to their base level of utility:

$$U_m^{base} \leq U_m(x, W - p * x) \tag{2}$$

If individuals only make a purchase when they have a preference for the item, we can define the purchase decision for an individual at each subset quantity, C_{xm} , as

$$C_{xm} = \begin{cases} U_m(x, W - p * x) \geq U_m^{base} & \text{Purchase} \\ U_m(x, W - p * x) < U_m^{base} & \text{Don't Purchase} \end{cases} \tag{3}$$

This utility curve has many potential shapes (Figure ??). As previous studies of food clubs have observed members occasionally purchasing entire co-purchasable items individually (Hupper, 2019; Lange, 2022), some members may have utility functions like that in the left hand panel of Figure ?? for some goods. However, these instances are rare. We expect most utility functions to look like the middle and right hand panels of Figure ??, where there is some range from x^{min} to x^{max} that satisfies the purchase condition of Equation 3, with some optimal quantity x^* endowing the most utility. Not pictured is the final possible shape, where there is no sub-quantity that satisfies the purchase condition, i.e. $x^{min} = x^* = x^{max} = 0$, indicating the individual has no preference for the item.

Figure 1 *Three hypothetical member utility functions (blue) compared to the base utility (red) across all share x quantities. Left: The member utility strictly increases from base utility with decreasing marginal returns. As such, $x^{\square} = x^{max} = X$. Center: Utility is less than the base at low and high share quantities but exceeds it at medium quantities with an optimum at x^{\square} . Right: Utility exceeds base utility at smaller share quantities with an optimum at x^{\square} , but is less at medium and higher quantities*

Comparing possible aggregations of member demand and its relation to X , i.e. how members arrive at the case size, allows us to theoretically define behaviors across the prosociality spectrum. If all members are able to meet the case size exactly with their optimal share quantity, i.e. $\sum x_m^* = X$, the co-purchase succeeds with purely selfinterested coordination, as every individual will receive their optimal share. However, if members' optimal share quantities do not meet the case size, i.e. $\sum x_m^* \neq X$, some or all participating members will need to change their share, and accept some utility cost in order for the co-purchase to succeed. This requires cooperation or altruism.

The costs incurred by a member in completing a shared purchase defines their action as cooperation or altruism. When a member's purchase quantity deviates from x^* , but lies between x_m^{min} and x_m^{max} , it is an act of cooperation. In this case, the member still receives more utility than they would with an equivalent amount of money, but it is less than they would have if they could purchase their preferred x^* shares. Going further, if a member needs to accept a quantity that is outside of the preferred purchase range, i.e. does not satisfy Equation 3, doing so would require some degree of altruism. The utility cost is the difference between x_m and x_m^* and the added cost of forgoing the equivalent amount of money.

This simple structure shows how free riding works in food clubs. For a single copurchased item to succeed, members solicit others to accept a utility cost by purchasing more or less than they would prefer. The members who assist in this way are acting cooperatively and sometimes altruistically. Therefore, the co-purchasing game is a variable social dilemma that depends on members' preferences and item factors such as case size. Cooperators or altruists must hope to recoup their utility losses through reciprocity on other items.

In many real-world scenarios, it may not be possible to define an action as cooperative or altruistic at a given moment in time, even from the perspective of the actor (Wilson, 2015), and co-purchasing is no different (Figure ??). Individuals might be expecting a cooperative act to be reciprocated, and the act may never be intended as altruistic. If it is reciprocated, the individual may think of it as successful cooperation through reciprocity. If it is never reciprocated, the individual may come to consider it as an altruistic outlay. Club members may not know how to consider a contribution to a co-purchasing effort until long afterwards.

Figure 2 Reciprocity determines behavioral classification after the fact. Cooperative assistance in food buying clubs is only beneficial to the cooperator if it is reciprocated. Reciprocation may occur at greater lengths of time and social distance, from direct (dyadic, often short term), indirect (cyclic, often medium term), or general (from the group, possibly long term). Evidence of cooperation derives from greater occurrence of direct and indirect reciprocity.

2.2 Comparable Theoretical Dilemmas

Because co-purchases are rivalrous (i.e., members' shares cannot also be purchased by other members) but non-excludable (i.e., any member can potentially purchase a share), they are common pool resources by definition (Reiss, 2021). However, traditional common pool resource games (e.g. Bru et al. 2003) usually require participants to choose an amount to take from a resource without having to contribute to it. This does not adequately reflect the decisions made in co-purchases, as individuals are pledging an amount of money to exchange for the provision of the good. Furthermore, while there are thresholds in common pool resource games, meeting them results in the collapse and non-provision of the resource, while thresholds *must* be met for co-purchased items to be provided.

An economic game that more closely emulates the contribution requirements and threshold conditions of co-purchases is the threshold public good game with refund (TPGR) (Cadsby and Maynes, 1999; Cartwright and Stepanova, 2015). TPGR games require individuals to pledge endowed money to a common fund where it will be multiplied and redistributed equally if the common fund reaches a required threshold, but pledged money is returned to recipients if the threshold is not met (Cartwright and Stepanova, 2015). The case sizes of each co-purchase are an example of what thresholds emulate, as not meeting them results in the item not being ordered and the return of money previously pledged.

Co-purchases deviate from the TPGR model in that the goods provided are a unit price reduction and personal utility, and multiple interactions are dispersed across a network where purchasing partners can change. Successfully reaching the threshold in TPGR game results in a uniformly divided monetary payout of the common fund plus a growth factor chosen by the researcher (Cadsby and Maynes, 1999; Cartwright and Stepanova, 2015). Meanwhile, while the unit price decrease of a successful copurchasing dilemma could be considered a 'payout', total savings relative to buying the good elsewhere still differ according to members' relative pledges. Additionally, accurately measuring the full hedonic utility payout of receiving the good (which is the most important because it drives demand) would require additional data beyond what was collected for this study (Fehr and Rangel, 2011). Furthermore, while individuals do receive utility from money they would have spent on the good if a co-purchase is unsuccessful, those whose required x^m would have satisfied the purchase condition in Equation 3 incur a utility cost by the item not being purchased, whereas TPGR games refunds have no such cost. Finally, traditional TPGR games generally assume sustained interactions in each game, while the member composition of co-purchasing interactions can vary across items and orders, with the potential for old members leaving the club and new members entering every order.

While these deviations are distinct, and a full formal model of the co-purchasing dilemma remains to be built, the procedural likeness of co-purchases to TPGR games allows for a novel opportunity to compare a nuanced real-world scenarios to straightforward theoretical predictions. First, co-purchasing represents a clear collective action dilemma with real monetary and utility trade-offs that are vulnerable to free-riding,

which can be difficult to capture data on in naturalistic settings. Second, they often lack formal institutions to reinforce cooperation by lowering its cost or penalizing free riding, making them prime subjects to study how prosociality is maintained through reciprocity (or not) early in the organizational lifespan. While a full analysis of the life cycle of these clubs is beyond the scope of this study, establishing an understanding of how reciprocity unfolds will generate insights into future studies of that nature. Finally, the data we have on these clubs allows for the granular measurement of prosocial behaviors that club members acknowledge occurs and is theoretically important to early organizational success. As such, we are able to complete a novel assessment of real world prosociality that lays the groundwork for future naturalistic study.

3 Methods

3.1 Dataset

Our dataset consists of a set of purchasing records from two different software platforms used by food clubs to coordinate purchases. Platform 1 provided data for 30 clubs from late 2011 to early 2017, and platform 2 provided data for 19 clubs from early 2010 to early 2019, for a total of 107 club years. The clubs in our dataset are samples from WEIRD countries: 2 from Australia, 1 from Canada, and the rest from the United States. For our analysis, we remove all purchases of individual items as well as clubs with no bulk purchases. The final combined dataset contained 35 clubs with 1,528 individuals across 1,341 orders for a total of 10,261 co-purchases.

Data from our providers only included software usernames and purchasing details, but no personal information. As such we were unable to ascertain differences in reciprocity or cooperation by age or gender identity. Furthermore, our sample includes defunct and still functioning clubs (as of the end of the data period), though the effect of reciprocity and cooperation on the functional status of these clubs is beyond the scope of this paper and is not a part of our analyses. See supplementary materials for club level summary statistics.

3.2 Co-purchasing network construction

We used a network-based method to study reciprocity between members over time (Figure ??). To detect reciprocity in co-purchases, we construct a bipartite purchasing network for each order by connecting members to the items that they purchased (all networks were generated and analyzed using R (Bates et al., 2021; Csárdi et al., 2023; R Core Team, 2022; Song et al., 2019; Wickham et al., 2017,1)). We then use this to create a unipartite co-purchasing network where members are connected to each other by their co-purchases, ensuring that multiple co-purchased items result in multiple edges per pair of members. Edge directionality is then assigned to according to the ordinal volume of co-purchase share. Arrows point to the member who purchased more in each dyad and are split into two directed edges pointing at both members when equal shares are purchased. We do this by building off the assumption that share size is directly proportional to members' relative utility, i.e. individuals who purchase more tend to derive more utility from the item, and stand to lose more utility if the bulk purchase is not completed. As such, members with lesser shares are assisting those who need to purchase more to derive their optimal utility. This method is corroborated by a negative correlation between the order in which members pledge to purchase shares of a bulk purchase, and the amount they buy (Spearman's $\rho = -.65$, $p < .001$; data only available from one of the software platforms). In

other words, individuals who join a bulk purchase later tend to buy smaller shares. We term this share-based assistance, and it allows us to delineate multiple levels of economic reciprocity.

Figure 3 Co-purchasing networks are projected to count reciprocity and ascertaining member types. *Purchase data are used to create a (A) bipartite network between members (red circles) and bulk food items (blue squares), which is projected as a (B) co-purchasing network between members that purchased the same items(s). Edge directionality is assigned according to an individual's relative share of the bulk item(s), and weight is assigned according to the number of such edges across multiple orders (not shown), to produce a (C) weighted and directed co-purchasing network, from which, (D) individual Markov transition probabilities between behaviors are derived.*

3.3 Identifying reciprocity

We organize patterns of reciprocity into separate categories based upon social and temporal proximity. We define the reciprocity of a club member, m , to other in the club by the relationship between their total in-degree, k_m^{in} , (number of times they were helped) and total out-degree k_m^{out} (number of times they have helped). For each order, we denoted the outgoing edges of m as reciprocal if they have a matching in-edge. For example, suppose two members, m_1 and m_2 , co-purchase shares of two items, A and B. If m_1 purchases a larger share of A than m_2 , and m_2 purchases a larger share of B than m_1 , we consider these co-purchases to be reciprocal in terms of relative share. Doing so allows us to categorize edges as either reciprocated or unreciprocated, and count the number of each across all orders.

We differentiate reciprocity as direct and indirect, and separate reciprocal behavior within the same order and between different orders. When calculating the different categories of reciprocity, we assume that social and temporal closeness are salient to economic transactions and should be respected when interpreting empirical patterns. This is in line with prior research findings that direct reciprocity is more consequential in sustaining cooperation when indirect reciprocity is also present (Roberts, 2008; Schmid et al., 2021), and is logically consistent. We therefore count direct reciprocity before indirect and reciprocity within orders before across orders. See our supplemental materials for a full explanation of the accounting method.

Finally, we estimate the global average rate of reciprocity per individual to compare with laboratory studies. Previous studies have measured reciprocity by regressing subjects' donations on group level average donations with subject level fixed or random effects (Croson et al., 2005; Smith et al., 2018). To compare with these studies, we specify a generalized linear model (GLM) of individuals' total out-degrees as a function of total in-degrees and group level random effects:

$$k_{m,c}^{\text{out}} = \alpha_c + \beta k_{m,c}^{\text{in}} + \epsilon_{m,c} \quad (4)$$

In (4), α_c is a random intercept term that allows us to account for club level variation such as unobserved socio-demographics, and $\epsilon_{m,c}$ is an error term. The β parameter may be interpreted as the global average percentage of in-degrees that are reciprocated across all individuals in all clubs in our sample, or the average number of out-degrees extended per in-degree. As such, it is comparable to coefficients from other studies that model average public goods donations as a function average group donation (Croson et al., 2005; Smith et al., 2018).

3.4 Identifying Cooperation and Altruism

Because the clubs in our study are small and bulk purchases are common, we expect to see some degree of reciprocity due to rational coordination; variation in preference overlap inevitably results in members buying different share sizes of many items than others in the club. We also know that club members acknowledge that cooperation is often involved in shared purchasing, and members report dissatisfaction and inequality in purchasing assistance (Hupper, 2017; Lange, 2022). We therefore investigate the strength of reciprocity as it reinforces prosociality across the spectrum, and how this varies across clubs. Significant variation in the extent of reciprocity would suggest that these clubs are solving their social dilemma in different ways, or not at all, consistent with the conditionality of prosocial behavior.

To improve inference regarding the prevalence of altruism, we use a restrictive set of criteria to identify economic interactions that are most likely to be altruistic. The co-purchases that are most likely driven by altruism are those in which one member helps another member to purchase an item which is outside of their revealed preference set (here defined as purchasing an item at least twice). To do so, we look for interactions we term "singular assistance", where individuals contribute only once to an item that has been purchased multiple times. To illustrate, suppose a member, m_1 , purchases shares of a bulk item in different orders, and member m_2 never buys that item across all orders except for a single interaction in which they bought a smaller portion of the item than m_1 . This purchasing pattern may suggest that m_1 prefers the item and m_2 does not, even after purchasing some of it in one round, so it is most plausible to conclude that m_2 was altruistically (or at least cooperatively) assisting m_1 . Reciprocity of singular assistance would then occur if m_1 assists m_2 in the same way in a later order. We estimate the regression model given by Equation 4 with members' singular assistance degrees to assess global average singular assistance reciprocity.

3.5 Stability of reciprocity

Counting of types of reciprocal edges in each order also allows us to compute their proportions over time. With this information we can look for variability in the patterns of reciprocity within groups, as well as measure their general temporal stability. Since the time between orders vary within and between clubs, a traditional time-series stability analysis such as an augmented Dickey-Fuller (Dickey and Fuller, 1981) or Philips-Perron (Phillips and Perron, 1988) tests would produce biased estimates of stability. Instead, we quantify stability as the degree of temporal variation in each reciprocity category represented by its coefficient of variation, c_v , as a simple measure of variability. Smaller c_v values indicate greater temporal stability and vice versa.

3.6 Behavioral types

We use individuals' behavior across all orders to categorize them into different behavioral types, analogous to the general behavioral types identified in laboratory circumstances (Andreozzi et al., 2020; Fischbacher et al., 2001; Frey, 2017). To do so, we calculate members' shared purchase ratio (SPR), which is the log ratio of their order level out- and in-degrees. For club member m in order t :

$$\text{SPR}_{mt} = \ln \left(\frac{1 + k_{mt}^{\text{out}}}{1 + k_{mt}^{\text{in}}} \right) \quad (5)$$

Taking the logarithm transforms the ratio into an easily interpretable magnitude: A positive SPR indicates that an individual gave more assistance than they received, a negative SPR indicates they received more than they gave, and a SPR of zero indicates perfect reciprocity. As individuals could give almost as much assistance as they received (or vice versa) and still be considered reciprocal by their peers (Molm, 2010; Trivers, 1971), a clustering approach was used to allow for data-driven flexibility in classification. For each club, members' SPRs were pooled across all orders and divide into three types using univariate k-means clustering with three means designated a-priori:

- **Helpers:** designated cluster centered at less than zero, i.e. $k_{mt}^{\text{out}} > k_{mt}^{\text{in}}$
- **Reciprocators:** cluster centered at zero, i.e. $k_{mt}^{\text{out}} = k_{mt}^{\text{in}}$
- **Beneficiaries:** cluster centered above zero, i.e. $k_{mt}^{\text{out}} < k_{mt}^{\text{in}}$

K-means clustering algorithms traditionally optimize all cluster centers (Steinley, 2006), but we modified the algorithm to fix the reciprocator center at zero, i.e., perfect reciprocity. If a SPR fell equidistant between the reciprocator center and either of the other centers, we assigned it to either the beneficiary or helper cluster to maintain conservative estimates.

To determine members' dominant behavioral type, we analyze type changes through time by calculating transition probabilities using a Markov chain (Spedicato et al., 2021). Markov chains are used to model systems where transitions occur between different states, and assume that the next state is exclusively dependent on the current state (Norris, 1997). While a Markov chain with memory would provide more precise predictions about members' evolution over time, traditional Markov chains produce good approximations of the stationary distributions of systems with memory, especially in systems with limited states such as ours (Wu and Chu, 2017). As such, we assign each member a dominant type based on which has the largest probability in the stationary distribution of their transition matrix. We report the dominant type assignments for individuals in each club, as well as a global level average transition matrix, which gives us a measurement of the stability of each role across clubs and which is the predominate behavioral type in each order (see supplemental materials Table 5 for club specific cluster centers).

3.7 Robustness Check

While accounting for assistance using relative share size is consistent with standard utility theory, it is possible that it is not always appropriate. For example, club members may reciprocate the responsibility for initiating the purchase process of mutually desired goods. Furthermore, the timing of a purchase may be another indicator of the utility endowed by an item, i.e., individuals purchase into items they want sooner in the order process than those that want less.

To account for this possibility, we checked the robustness of our results with the 13 clubs whose software provider gave us information on the timing of member purchasing. We reassigned the directionality from Figure 2C for co-purchasing networks based on *when* members commit to purchasing an item. For this purchase order definition of assistance,

members who purchase items later are assumed to be aiding those purchase sooner. From these new networks, we re-count each type of reciprocity and re-estimate our GLM.

4 Results

We find that club members are highly reciprocal in their purchasing patterns overall. The majority (60%) of reciprocity occurs direct and within orders. We also find that reciprocators are the most common and the most stable behavioral type.

4.1 Reciprocity types

In classifying different types of reciprocity, we observe within-order direct reciprocity as the most frequent form, with 60% of the average club's edges classified as directly reciprocal within orders (Figure 3A). Directly reciprocal between order edges represent a far second, followed by indirectly reciprocal edges within orders and between orders (see supplemental materials Table 2 for percentages of each). Finally, 17% of all edges remained unreciprocated as of the final observed order.

Globally, our GLM indicates that members reciprocate 88% of the co-purchasing assistance they receive ($\beta = .88$, $p < 0.001$). This can be considered an upper bound reciprocity that covers the whole suite of prosocial behaviors, including altruism, cooperation, and coordination. Using the more restrictive singular assistance measure we find that individuals reciprocate 46% of co-purchasing help they receive on average ($\beta = .46$, $p < 0.001$). As this metric most closely measures reciprocated altruistic cooperation, this can be considered the lower bound estimate of reciprocity. See supplemental materials Table 4 for full model results.

4.2 Temporal variation

We find that the variability of each type of reciprocity corresponds inversely to the pattern of prevalence (Figure 3B). For example, the proportion of directly reciprocal edges within orders have the least temporal variation (and are therefore the most stable), both temporal forms of direct reciprocity vary less over time than indirect reciprocity, and within-order reciprocity is less variable than reciprocity between orders.

Figure 4 Rapid, direct reciprocity is most common and most stable. (A) More than half of all bulk purchasing is reciprocated directly and rapidly (within the same order). Diamonds represent mean club abundance of each reciprocity type. (B) Within order, direct reciprocity is also the most stable. Diamonds represent mean coefficient of variation, c_v , by club.

4.3 Behavioral types

66.1% of club members were classified as dominant reciprocators by the stationary distribution of their Markov chains. Across all clubs, reciprocators are the most abundant member type, with an average proportion of 65.3%, followed by helpers (29.4%) and beneficiaries (12.4%) (Figure 4). Reciprocator is also the absorbing state of the global average transition, where the probability for staying reciprocator in the next order is 51%, and the probabilities for becoming a reciprocator from helper and beneficiary are 55.7% and 56.2% respectively. Figure 5 shows the member type composition of each club in a stacked

bar graph, and Figure 6 shows the global average transition probabilities. See supplemental materials Table 5 for SPR cluster centers, and Table 6 for member type counts.

Figure 5 Reciprocators are the most abundant member type. Club composition by member type, arranged in descending order of proportion of reciprocators. Reciprocators are the most abundant type, followed by helpers and beneficiaries, who remain consistently below 25% of members.

Figure 6 Reciprocator is the most stable behavioral type within orders. Reciprocator type is the average absorbing state, with a greater than 50% chance for each role to be a reciprocator in the next order, globally. Circle size is proportional to the sum of all incoming transition probabilities.

Figure 7 Changing the criteria of assistance induces changes in the average timing of reciprocity. By defining assistance by purchase order, there is a substantial increase in the amount of reciprocity occurring between orders, directly and indirectly. There is also a large decrease in within order direct reciprocity. Most bulk purchasing is reciprocated directly (by the beneficiary) and over time (between orders). Global means are shown as black diamonds.

4.4 Robustness Check

Defining assistance by purchase order for a subset of clubs induces changes in the patterns of observed reciprocity. Our random effects estimation of the altered networks estimates that individuals reciprocate approximately 63% of the assistance they are given ($\beta = .63$, $p < 0.001$). We also find that edges classified as directly reciprocal edge *between* orders are most abundant with 32.9%, followed by indirectly reciprocated between orders (25.7%), and within order direct (18.7%), and indirect (8.9%).

5 Discussion

The observed patterns of economic cooperation in small food clubs provide some of the first estimates of reciprocity rates and the emergence of behavioral types in a real world organizational scenario. The degree of economic reciprocity exhibited in these clubs is, to our knowledge, unprecedented in the empirical literature, and it supports all prosocial behavior. Our findings are supported by multiple lines of evidence, include a high mean level of reciprocity across clubs, the existence of multiple types of reciprocity (direct and indirect, short-term and longer-term), the predominance of the reciprocator behavioral type within clubs, and a robustness check.

High levels of reciprocity. Our GLM estimated that 88% of all edges are reciprocated, as are 46% of singular assistance edges. The all edge model represents a ceiling value that combines interactions ranging from coordination to altruism. The singular assistance model used a highly restrictive definition of cooperative edges to approximate the floor value that would be indicative of just cooperation and altruism, as individuals would likely have bought the co-purchased items multiple times if they bore any utility.

We find evidence that we are measuring the full spectrum of prosociality by comparing our model estimates with the average donations of economic games played with these groups in previous research (Hupper, 2019). The results of our total edge model correspond with the results of a public goods game ($\beta = .88$, PGG mean donation = 71%) and the

singular assistance results correspond with the results of the dictator game ($\beta = .46$, DG mean donation = 58%). Dictator games elicit the propensity for unenforced fairness and altruism towards another individual (Engel, 2011), and our singular assistance edges were restricted to identify similar behavior in a real world context, as they can only link to a single person, the purchase initiator. Concurrently, public goods games measure how cooperative individuals are towards an entire group (Zelmer, 2003), and our total edge model captured the reciprocity directed at multiple members of co-purchases. This apparent alignment between parameters and game donation amounts may imply that our singular assistance model measures actual reciprocal altruism, while our total edge model represents individuals being prosocial in general to support the club. However, the inevitable presence of coordination based reciprocity in our total edge model makes absolute comparison imperfect.

Further evidence of our claim comes in the comparison of our estimates to previous studies of conditional cooperation that compare individuals' donation in public goods games to the mean group donation or subjects perception thereof. As club members' in-degrees indicate the total amount of assistance given to them from their club mates, they are tantamount to subjects' payouts in experimental public goods games and are a measure of average group level donation as experienced by the club member. Members' out-degrees are consequently the amount of assistance they've given and is comparable donations in experiments. In all cases, our total edge model exceeds previous estimates of conditional cooperation, and our single assistance model falls within their range.

In anonymous public goods games, Croson et al. (2005) found that people gave an average of 40% of the observed group donation (Table 1 Voluntary contribution mechanism in compared study). In another anonymous public goods games, Fischbacher and Gächter (2010) found that individuals gave 67% of what they believed the rest of the group gave (Table 2 model 3 in compared study). Using an instrumental variable approach in an anonymous public goods games, Smith (2013) found that individuals gave between 45% and 57% of what they believed others gave in the previous round (Table 2 in compared study), depending on model specification. Finally, in public goods games with Hazda hunter-gatherers, Smith et al. (2018) found that people donated 55% of the average group donation (Table S4, Model 1 of compared study). All of these studies utilize a linear public goods setup, which is specifically designed to elicit behavior that incurs a personal cost, i.e. cooperation and altruism. Since our singular assistance reciprocity estimate falls well within these studies, we believe we are likely observing reciprocal altruism, or cooperation at the very least. Furthermore, the outlay of our total edge model from previous studies suggests that we are observing reciprocity support the full spectrum of prosocial behaviors, including coordination.

One possible explanation for the our observed rates of reciprocity is that our groups hail from WEIRD or Western, Educated, Industrialized, Rich, Democratic, countries (Henrich, 2020). WEIRD societies have greater degrees of market integration, which has been theorized to increase interpersonal cooperation (Henrich, 2020). Cross-cultural studies that compare group average cooperation in economic games and various measures of social development and market integration have been employed to test this hypothesis (Henrich et al., 2001). For example, Henrich et al. (2010) found that individuals from countries with more market integration tend to exhibit higher levels of cooperation in economic games. Furthermore, individuals from firm-type organizations within WEIRD societies, including managers and general employees, perceive that prosociality is fundamental to

organizational success (Koster and Sanders, 2006). Since our clubs are *facilitating* market integration by supplying a missing market, higher prosociality may be required.

Predominance of direct reciprocity. Direct reciprocity is generally more effective at supporting cooperation than indirect reciprocity, and does so in greater amounts when they are allowed to co-evolve (Roberts, 2008; Schmid et al., 2021). When assessing directional assistance by share size, our findings support this assertion, as we found that most reciprocity occurred between pairs of individuals and on a short time scale (about 1-4 weeks). 72% of all reciprocal interactions and 60% of all co-purchasing was attributable to within-order direct reciprocity. Though assessing directional assistance through purchase order in our robustness check decreases the amount of direct reciprocity within orders, direct reciprocity is still most prevalent. This suggests that co-purchasing reciprocity supports prosocial co-purchasing behavior, at least in the short term, and club members resolve this as immediately and directly as possible. If buying clubs did not require prosociality, and cooperation or altruism specifically, between members, we would expect less reciprocity of all types, and especially less direct reciprocity.

Our results also show how dyadic reciprocal interactions can undergird the cooperative success of group level social dilemmas. This parallels research in which experimental subjects alternated between public goods games, and dyadic prisoners dilemmas. A study by Rand et al. (2009) found that cooperators who were rewarded with cooperation (and free riders who were punished via defection) in the prisoner's dilemma rounds tended to cooperate more in the public goods rounds. This same dynamic could be unfolding in the buying clubs we study, as members who aid in buying larger items that require extra individuals could be rewarded with assistance on smaller items that they initiate, or could be punished for not doing so by having their items go unfulfilled. However, testing such an assertion would require analysis that is beyond the scope of the current work.

More socially and temporally distant forms of reciprocity are also detected when assessing assistance by share. Theory and simulation reveal that both direct and indirect reciprocity evolve simultaneously to support cooperation, with direct reciprocity playing a more prominent role (Boyd and Richerson, 1989; Nowak and Sigmund, 2005; Roberts, 2008), and our results suggest that our sample corroborates this theory. If only direct reciprocity were at play, we would only see co-purchases where many individuals purchase the same amount of an item, or individuals trading off on who buys the lion's share of an item. Similarly, if only indirect reciprocity were at play, we might mostly find instances where individuals buy the largest portion of one item with one person and the smallest portion of another item with someone else. Since we detect combinations of both (among other reciprocal patterns) our results suggest that the theoretical predictions of how direct and indirect reciprocity reinforce cooperation play out in this naturalistic setting.

The differences in reciprocity patterns when using a purchase order-based may reveal more about clubs' purchase processes. To begin, the average global reciprocity rate is 25 percentage points less than the share-based criterion. This could be due to arbitrary circumstances that delay individuals' ability to pledge, such as delays to their ability to log into the software. As the utility assumptions are consistent across both criteria (individuals who want or need an item more should hypothetically buy more of it and sooner), these reasons could also explain why the correlation between share size and purchase order isn't stronger. Additionally, the increase in between order reciprocity could indicate that individuals are trading off responsibility for initiating purchases. If opposite individuals are initiating items each order, we would expect to find more between order reciprocity rather than within order, as we do. Despite this, we still find that direct reciprocity is more

abundant that indirect, which is further evidence that direct reciprocity plays a key role in successful co-purchasing.

Behavioral types. In the initial work of Fischbacher et al. (2001), individuals were classified based on the proportion of their endowments given: free-riders kept all or most of their endowment, altruists gave all or most of their endowment, and conditional cooperators gave around the group average from the previous round. While intuitive, using these classifications to describe real world behavior is limited because experiments oversimplify the dynamics of naturalistic settings. Studies of strategies and behaviors in the real world further highlight that complexity rarely allows for straightforward behavioral classification (Efferson et al., 2023; Mesoudi et al., 2016), leading to the possibility that behavioral types are experimental contrivances.

While there is logical correspondence between the behaviors exhibited in our groups and those found in experiments, the differences are notable in definition and emergence. Conditional cooperators and reciprocators have the most equivalency, as conditional cooperators do what the group does on average and reciprocators are generally giving what has been given to them. Free-riders in experiments line up with beneficiaries as both gain more than they give, and helpers line up with altruists for the opposite reason. However, despite having a usual type, many members still alternate from order to order, making their behavior less straightforward than would be expected from experiments. Specifically, beneficiaries still render assistance (albeit less often) making their behavior less opportunistic in absolute terms than traditional free-riders. Similarly, helpers still benefit from assistance in most orders, indicating that their role is less self-sacrificial than their altruistic counterparts in experiments.

This leads to the conclusion that in naturalistic settings such as food clubs, cooperation and reciprocity appear as probabilistic patterns rather than behavioral types, rather than definitive classifications, at least in our case. Experiments have analyzed the stability of behavioral types across time by retesting previous subjects, and found that individuals tend to repeat their behavior (Andreozzi et al., 2020; Kocher et al., 2008). While we also observe repeat behaviors, we more often observe members fluctuate among the roles with a tenancy to default to one more often than the others. Furthermore, reciprocity appears to be the most stable strategy of them all, as more individuals across all clubs tend to default to this role, and it acts as the absorbing state in the global transition matrix.

As conditional cooperators and reciprocators have the most in common, it is useful to compare our observed proportion of reciprocators with those of conditional cooperators in the experimental literature. 68.5% of all individuals were classified as consistent reciprocators, which is one of the highest proportions found in the literature to date. Fischbacher et al. (2001) were the first to identify behavioral types by correlating subjects' donations with the average donation of the group with the focal individual's donation removed, and classified 50% of their subjects as conditional cooperators. In a review of 17 replications of Fischbacher et al. (2001), Thöni and Volk (2018) found proportions of conditional cooperators ranging from 40% to 77%, with an average of 62%; of the studies reviewed, only 2 studies exceeded our proportion. In an organizational setting, Deversi et al. (2020) found approximately 41% conditional cooperators (called "matchers" in text) in a company-wide experiment using a clustering classification method. Finally Frey (2017) found that 39% of individuals playing an online video game with a threshold public goods mechanism could be classified as conditional cooperators by correlating individual's effort with the effort of others in the game.

These comparisons indicate that members of our clubs are more likely than most individuals in public goods experiments to be reciprocators/conditional cooperators, and that conditional cooperation and reciprocity may be higher in naturalistic public goods situations among consistent peers with pre-established relationships, or where individuals interact outside of the collective action dilemma. Furthermore, we may observe more reciprocators because cooperation and coordination are intermixed in these dilemmas. Experimental conditional cooperators are *explicitly* cooperating due to the experimental setup, whereas reciprocators are potentially reciprocating coordination and cooperation. This could indicate that when real-world dilemmas contain a mix of coordination and cooperation dilemmas, reciprocity tends to be even more of an evolutionarily stable strategy in the long term facilitation of general prosociality.

Group variation. Finally, we find that buying clubs vary dramatically in their reciprocal economic behavior. Clubs vary in the amount of reciprocity they exhibit on average (from 91.7% to 59.7%), and in their member behavioral type composition (16.7% to 100% reciprocators, 0%-33.3% beneficiaries, and 0%-81.3% helpers). Group level variation is a natural phenomenon in all social contexts, however, research suggests that human cultural (and organizational) evolution may be largely driven by the ability of groups to maintain effective patterns of group cooperation toward collective goals (Richerson et al., 16ed; Waring et al., 2022). Thus, to the extent that the differences we observe in economic reciprocity are indicative of underlying patterns of cooperation, they may be consequential in the survival of these small clubs. Further research is required to investigate the consequences of reciprocity on club persistence.

The wide variation we observe could be evidence of multiple means of solving these clubs' social dilemma. Clubs with a higher abundance of beneficiaries and helpers could be structured that way because helpers gain utility both from the items they purchase and assisting other members of the club, i.e., they have a social preference that is fulfilled by their club's success. This is in line with survey evidence where some members have indicated that they enjoy completing bulk purchases that would otherwise go unfilled (Hupper, 2017,1). Moreover, clubs could establish norms and institutions as they age to lessen the cost of their social dilemma or eliminate it entirely.

Proof of concept. The greatest implication of our research is that patterns of cooperation in real world situations are increasingly observable thanks to digital records of economic interactions. Human cooperative behavior has long matched qualitative and ethnographic descriptions of human behavior (Henrich and Muthukrishna, 2021). However, behavioral games that measure cooperation, such as the public goods game or the dictator game, can now be more directly calibrated to true economic cooperation. Furthermore, our results highlight that altruism, cooperation, and coordination operate within a spectrum of prosocial behavior, and reciprocity is an effective strategy at maintaining them all. As such, researchers' intent on understanding human cooperation in a naturalistic setting must further integrate real world scenarios into models of human prosociality. Our analysis shows that this is not only possible, but necessary for the study of cooperation.

The differences between the structure of co-purchases and experimental public goods settings are instructive in this regard. To begin, members of these clubs are part of a group with clear, narrow goals that substantially overlap with one another, that is, an organization. In experimental settings, individuals are often strangers with no previous affiliation (though the findings of Smith et al. (2018) come from affiliated tribe members of hunter-gatherers). Many studies of cohesive groups that manage resources have qualitatively deduced high degrees of prosociality (Ostrom, 1990; Wilson et al., 2013), and survey evidence reveals that

members of these clubs are similarly nonrandom and self-selected Hupper (2019); Lange (2022); Tremblay and Waring (2015), so we would expect them to be more cooperative and reciprocal than strangers in a game because of their stake in the club itself and their familiarity with their club-mates. Consequently, our results serve as a quantitative verification that individuals with established relationships in organizational contexts express greater levels of reciprocity and cooperation than would be expected based solely on laboratory experiments with strangers.

That these clubs are made of up of non-random individuals prosocially combining their purchase power to supply a missing market reveals how our results likely generalize most to situations where traditional market mechanisms are missing. The history of consumer co-operatives is rife with examples where individuals pool their resources to make markets, all the way to the beginning of the co-operative movement (Waring et al., 2022). Adjacent organizations committed to mutual-aid have been supplying services and information to undeserved populations for decades, including providing medical and mental health services (Archibald, 2007), and forming temporary hubs to reduce food waste and address food insecurity during COVID-19 (Lofton et al., 2022). Informal and pre-formal organizations and networks that rely on prosociality likely rely on high rates of reciprocity as well, at least until formal rules can be established.

A final key diversion between food buying clubs and experimental settings is that individuals in public goods games receive tangible monetary benefits through uniform payoffs (plus the amount they kept during donation). In food clubs, individuals do not receive a uniform share payoff, but gain a discount on the goods they purchased in bulk and access to a greater selection of goods that may not be found in traditional market settings (i.e., the grocery store) by being in the club. This reveals that in this real-world setting, individuals act reciprocally even though the benefits are not necessarily monetary (access to the club and intangible utility from items) and are heterogeneous across individuals (as discounts are different depending on what items are bought).

Limitations and future work. There are a few limitations in our study. First, a nuanced time-series analysis would provide a better understanding of stability, though many of these models are infeasible when club orders occur at irregular intervals. There are also aspects of reciprocity and cooperation that our share-based and purchase order-based criteria of assistance do not capture. For example, there may be situations where demand exceeds the threshold for an item but is insufficient to purchase two or more cases. While the first case will be fractionally Pareto-optimal, it still does not reach the *actual* optimum. In this circumstance, we may observe reciprocity based on share if some individuals stepped up to fill more than one case, but we would not observe it if individuals have accepted less than their ideal amount to allow access to all who wanted some. Furthermore, as stated before, our purchase order-based criterion is potentially confounded by unobservable factors such as internet access and individual schedule that may arbitrarily affect when individuals are able to pledge their share.

A potential drawback of our method of calculating reciprocity is that it may not fully account for the utility costs associated with how prosociality can be conceptualized within these groups. To illustrate, we can imagine a situation where an individual's lions share of a co-purchase is more (or less) than their optimum, while those with smaller shares are purchasing at their optimum. Though it is still justifiable to say that those who buy less are still helping the member with the most to avoid failing to purchase, we would not pick up the cooperative nature of utility cost being assumed by the principle buyer so the secondary buyers do not have to deviate from their optimum. This situation would result in us *under-*

counting the true amount of reciprocity and the degree of cooperation involved. Our results may therefore represent a conservative estimate of reciprocity.

The example above highlights how difficult it can be to truly disentangle prosocial behaviors from each other in naturalistic settings, despite quality data. While we are confident that our singular assistance edge definition comes close to matching the cooperation and altruism elicited by laboratory studies, further study that incorporates other factors would refine our estimates. Future studies might fit dynamical models that measure individual preferences over time to determine how much of the observed reciprocity is experienced as altruistic. Additionally, future surveys of members could elicit how often they buy more or less of an item than preferred on behalf of others, as well as how often they feel their shares sacrifice their ability to purchase goods they might have otherwise preferred. This would allow us to better quantify how much reciprocity can be portioned into sustaining coordination, cooperation, and altruism.

5.1 Summary

We have used a novel approach to detect and describe reciprocity in the economic networks of small food buying clubs in which economic cooperation is thought to be necessary. Our observational data validate three key findings of the prior experimental literature on cooperation and extend the domain in which we can expect cooperation and reciprocity to be applicable. We find higher levels of reciprocity compared to previous studies that support likely altruism, cooperation, and coordination. We find similar fractions of behavioral types, with reciprocators or conditional altruists being most common. We find that rapid dyadic reciprocity is much more common than indirect reciprocity and that short-term reciprocity is more common than longer term patterns. The results of our study imply that a notable fraction of economic reciprocity in buying clubs is cooperative, and even altruistic, rather than merely coordinative because the patterns match those of experimental studies designed to elicit cooperation. Further research might use similar methods to explore the evolution of cooperation in different organizational contexts.

Acknowledgments

The authors would like to thank Marco Smolla for his continued insight throughout this process, Afton Hupper and Ethan Tremblay for their prior work on food clubs and their thoughtful input on analysis. We thank Abby Leidenfrost for her assistance with data cleaning, and our software providers for allowing us to access the data from their clubs. We would also like to thank Christine Beitzl, Andrew Crawley, Caroline Noblet, and Bryce Morsky, Charles Efferson, and two anonymous reviewers for their insightful reviews of this manuscript.

Financial Support

Funded by National Science Foundation CAREER award SES-1352361, and USDA National Institute of Food and Agriculture, Hatch project ME022008. Data collection was approved by the University of Maine's Human Subjects Review Board protocol 20160312.

Author contributions

TL and TW conceptualized the article, TL conducted data analysis guided by TW. TL & TW wrote the manuscript.

Conflict of Interests

The authors declare that they have no competing interests.

Research Transparency and Reproducibility availability

All data and code used in this paper is available at <https://github.com/tzlang16/Reciprocity-in-Real-World-Networks>.

Figure 8 Three hypothetical member utility functions (blue) compared to the base utility (red) across all share x quantities. Left: The member utility strictly increases from base utility with decreasing marginal returns. As such, $x^{\square} = x^{\max} = X$. Center: Utility is less than the base at low and high share quantities but exceeds it at medium quantities with an optimum at x^{\square} . Right: Utility exceeds base utility at smaller share quantities with an optimum at x^{\square} , but is less at medium and higher quantities




Figure 9 Reciprocity determines behavioral classification after the fact. Cooperative assistance in food buying clubs is only beneficial to the cooperator if it is reciprocated. Reciprocation may occur at greater lengths of time and social distance, from direct (dyadic, often short term), indirect (cyclic, often medium term), or general (from the group, possibly long term). Evidence of cooperation derives from greater occurrence of direct and indirect reciprocity.

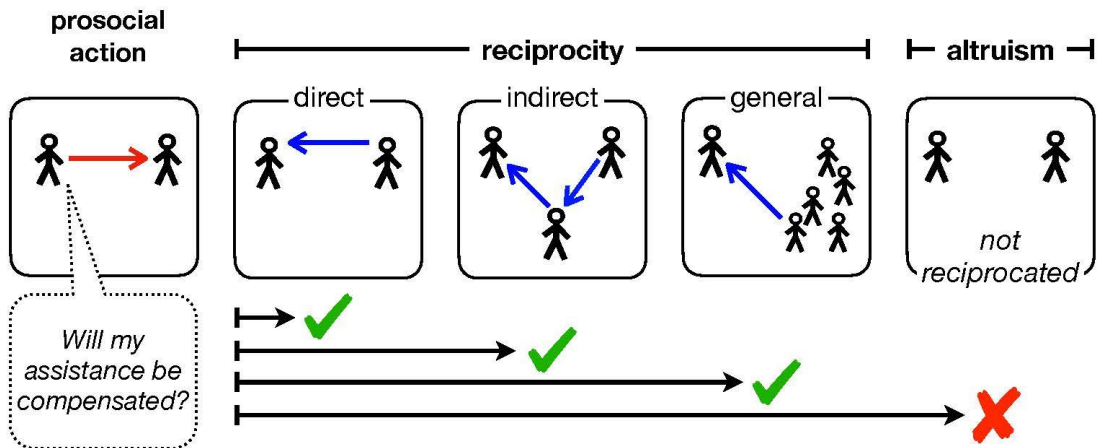


Figure 10 Co-purchasing networks are projected to count reciprocity and ascertaining member types. Purchase data are used to create a (A) bipartite network between members (red circles) and bulk food items (blue squares), which is projected as a (B) co-purchasing network between members that purchased the same item(s). Edge directionality is assigned according to an individual's relative share of the bulk item(s), and weight is assigned according to the number of such edges across multiple orders (not shown), to produce a (C) weighted and directed co-purchasing network, from which, (D) individual Markov transition probabilities between behaviors are derived.

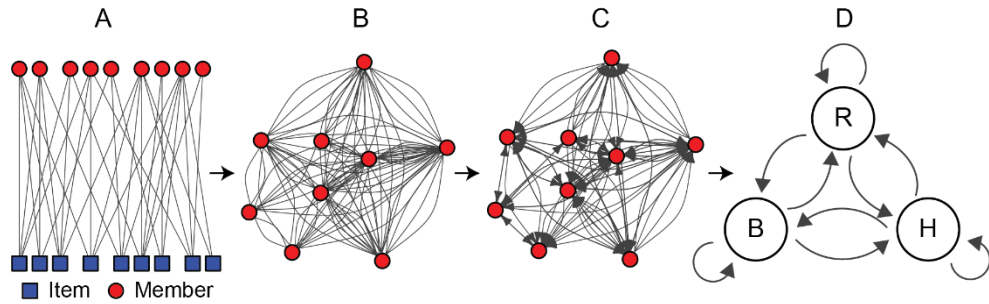


Figure 11 Rapid, direct reciprocity is most common and most stable. (A) More than half of all bulk purchasing is reciprocated directly and rapidly (within the same order). Diamonds represent mean club abundance of each reciprocity type. (B) Within order, direct reciprocity is also the most stable. Diamonds represent mean coefficient of variation, c_v , by club.

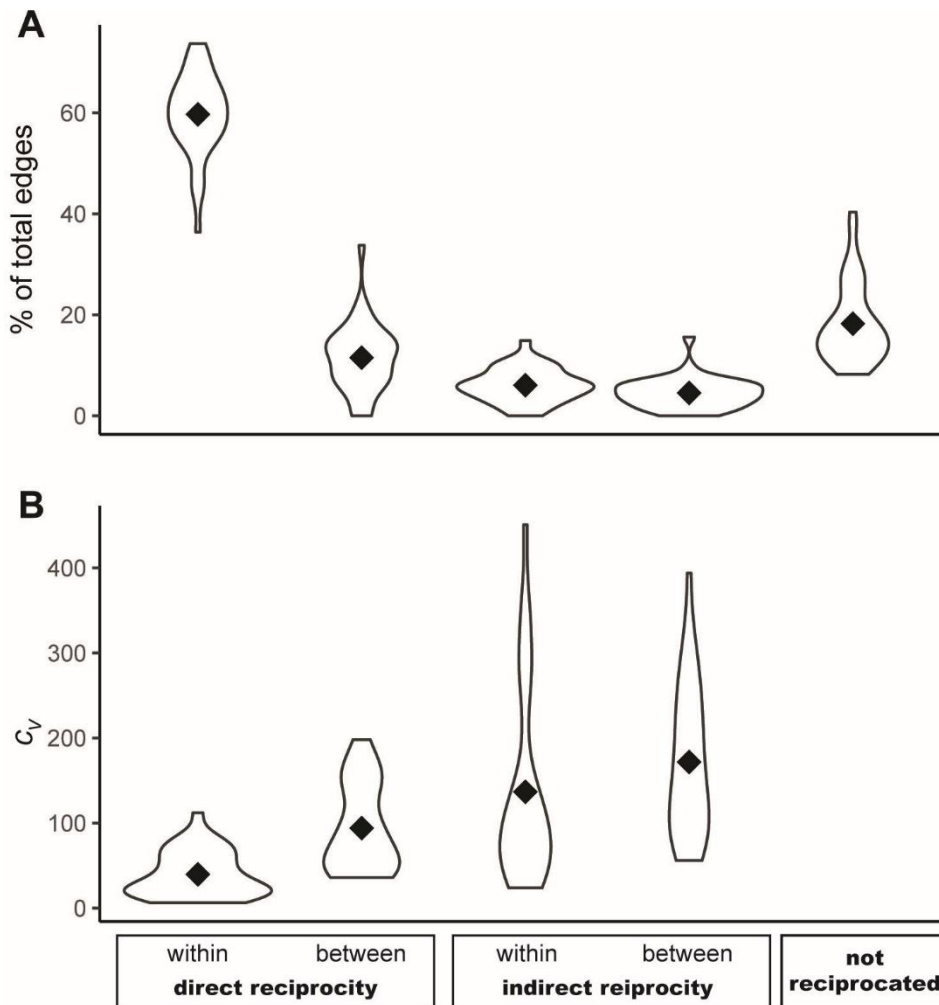


Figure 12 Reciprocators are the most abundant member type. Club composition by member type, arranged in descending order of proportion of reciprocators. Reciprocators are the most abundant type, followed by helpers and beneficiaries, who remain consistently below 25% of members.

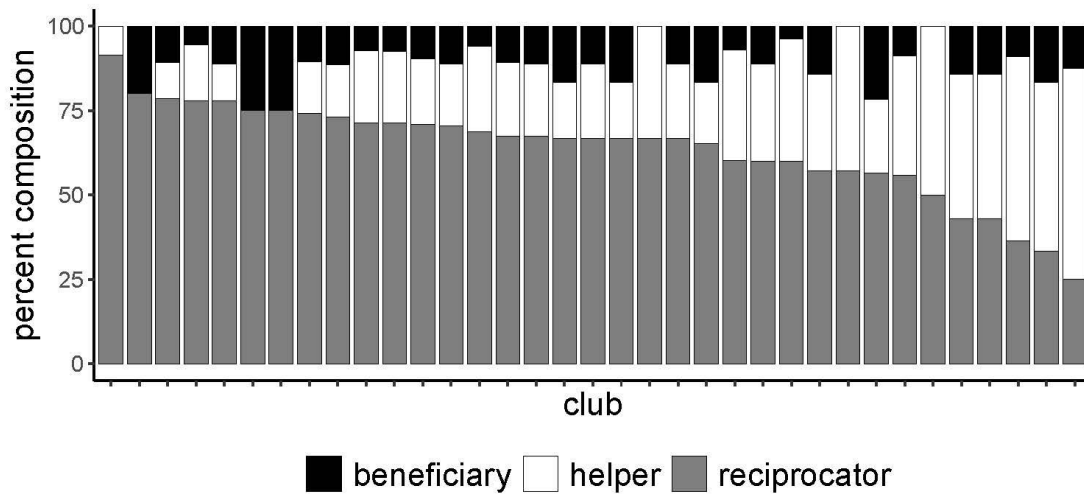


Figure 13 Reciprocator is the most stable behavioral type within orders. Reciprocator type is the average absorbing state, with a greater than 50% chance for each role to be a reciprocator in the next order, globally. Circle size is proportional to the sum of all incoming transition probabilities.

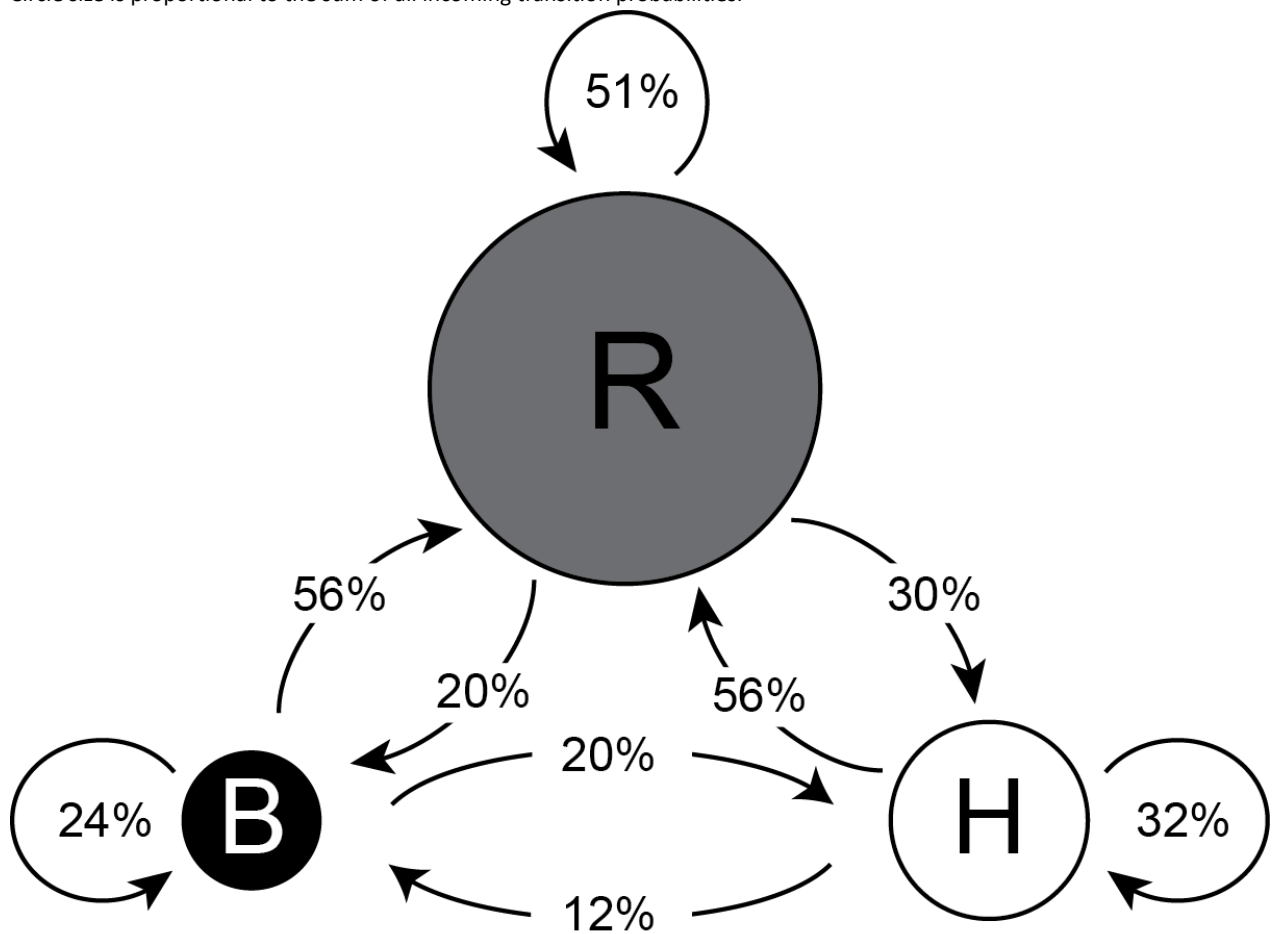
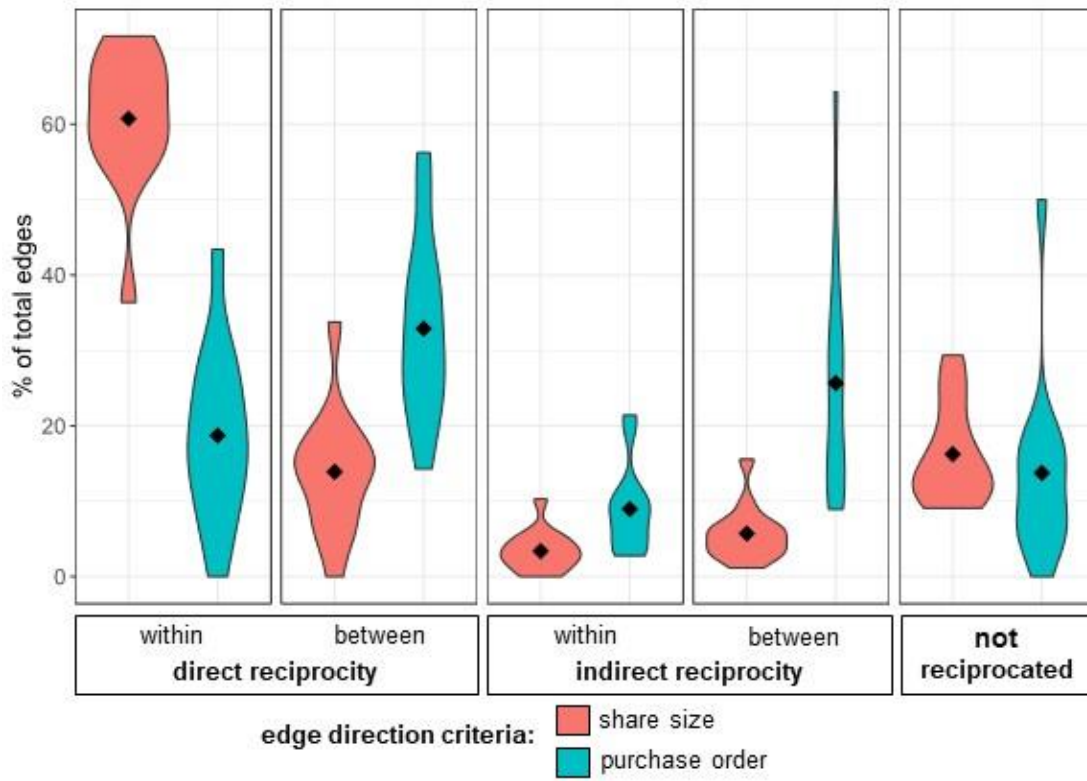


Figure 14 Changing the criteria of assistance induces changes in the average timing of reciprocity. By defining assistance by purchase order, there is a substantial increase in the amount of reciprocity occurring between orders, directly and indirectly. There is also a large decrease in within order direct reciprocity. Most bulk purchasing is reciprocated directly (by the beneficiary) and over time (between orders). Global means are shown as black diamonds.



References

- Andreozzi, L., Ploner, M., and Saral, A. S. (2020). The stability of conditional cooperation: Beliefs alone cannot explain the decline of cooperation in social dilemmas. *Scientific Reports*, 10(1):13610.
- Apicella, C. L., Marlowe, F. W., Fowler, J. H., and Christakis, N. A. (2012). Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382):497–501.
- Archibald, M. E. (2007). An Organizational Ecology of National Self-Help/Mutual-Aid Organizations. *Nonprofit and Voluntary Sector Quarterly*, 36(4):598–621.
- Asker, J. and Nocke, V. (2021). Chapter 12 - Collusion, mergers, and related antitrust issues. In Ho, K., Hortaçsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization*, volume 5 of *Handbook of Industrial Organization, Volume 5*, pages 177–279. Elsevier.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science (New York, N.Y.)*, 211(4489):1390–1396.
- Bates, D., Maechler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., Dai, B., Scheipl, F., Grothendieck, G., Green, P., Fox, J., and Bauer, A. (2021). lme4: Linear Mixed-Effects Models using Eigen and S4.
- Birchall, J. (1997). *The International Co-operative Movement*. Manchester University Press.
- Boyd, R. and Richerson, P. J. (1989). The evolution of indirect reciprocity. *Social Networks*, 11(3):213–236.
- Bru, L., Cabrera, S., Capra, C. M., and Gomez, R. (2003). A common pool resource game with sequential decisions and experimental evidence. *Experimental Economics*, 6(1):91–114.
- Cadsby, C. B. and Maynes, E. (1999). Voluntary provision of threshold public goods with continuous contributions: Experimental evidence. *Journal of Public Economics*, 71(1):53–73.
- Cartwright, E. and Stepanova, A. (2015). The consequences of a refund in threshold public good games. *Economics Letters*, 134:29–33.
- Croson, R., Fatas, E., and Neugebauer, T. (2005). Reciprocity, matching and conditional cooperation in two public goods games. *Economics Letters*, 87(1):95–101.
- Csárdi, G., , T., Traag, V., Horvát, S., Zanini, F., Noom, D., and Müller, K. (2023). Igraph: Network Analysis and Visualization.
- Deller, S. C., Lamie, D., and Stickel, M. (2017). Local foods systems and community economic development. *Community Development*, 48(5):612–638.
- Deversi, M., Kocher, M. G., and Schwioren, C. (2020). Cooperation in a Company: A Large-Scale Experiment. SSRN Scholarly Paper ID 3570293, Social Science Research Network, Rochester, NY.
- Dickey, D. A. and Fuller, W. A. (1981). Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root. *Econometrica : journal of the Econometric Society*, 49(4):1057–1072.
- Efferson, C., Vogt, S., and von Flüe, L. (2023). Activating cultural evolution for good when people differ from each other. In Tehrani, J. J., Kendal, J., and Kendal, R., editors, *The Oxford Handbook of Cultural Evolution*. Oxford University Press.
- Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, 14(4):583–610.
- Fehr, E. and Rangel, A. (2011). Neuroeconomic foundations of economic choice—recent advances. *Journal of Economic Perspectives*, 25(4):3–30.
- Fischbacher, U. and Gächter, S. (2010). Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments. *American Economic Review*, 100(1):541–556.
- Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3):397–404.
- Frey, U. J. (2017). Cooperative strategies outside the laboratory — Evidence from a long-term large-N-study in five countries. *Evolution and Human Behavior*, 38(1):109–116.
- Frey, U. J. (2019). Long-term evidence on cooperation and cultural differences in public goods dilemmas. *Biology Letters*, 15(4):20190143.

- Gintis, H. (2009). *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction Second Edition*. Princeton University Press.
- Gowdy, J. M. (2021). *Ultrasocial: The Evolution of Human Nature and the Quest for a Sustainable Future*. Cambridge University Press, Cambridge.
- Gowdy, J. M. and Krall, L. (2016). The economic origins of ultrasociality. *Behavioral and Brain Sciences*, 39.
- Henrich, J. (2015). *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press, Princeton.
- Henrich, J. (2020). *The WEIRDest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous*. Farrar, Straus and Giroux, New York.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., and McElreath, R. (2001). In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. *American Economic Review*, 91(2):73–78.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., and Ziker, J. (2010). Markets, Religion, Community Size, and the Evolution of Fairness and Punishment. *Science*, 327(5972):1480–1484.
- Henrich, J. and Muthukrishna, M. (2021). The Origins and Psychology of Human Cooperation. *Annual Review of Psychology*, 72:207–240.
- Hupper, A. (2017). *The Role of Cooperation and Prosocial Behavior in Food Buying Clubs: An Exploratory Study*. Honors Thesis, The University of Maine, Orono, Maine.
- Hupper, A. (2019). Local Food Policy & Consumer Food Cooperatives: Evolutionary Case Studies. Master's thesis, The University of Maine, Orono, Maine.
- Janssen, M. A., Anderies, J. M., and Joshi, S. R. (2011). Coordination and cooperation in asymmetric commons dilemmas. *Experimental Economics*, 14(4):547–566.
- Jensen, K. (2016). Prosociality. *Current Biology*, 26(16):R748–R752.
- Kocher, M. G., Cherry, T., Kroll, S., Netzer, R. J., and Sutter, M. (2008). Conditional cooperation on three continents. *Economics Letters*, 101(3):175–178.
- Koster, F. and Sanders, K. (2006). Organisational citizens or reciprocal relationships? An empirical comparison. *Personnel Review*, 35(5):519–537.
- Koster, J. M. and Leckie, G. (2014). Food sharing networks in lowland Nicaragua: An application of the social relations model to count data. *Social Networks*, 38:100–110.
- Lange, T. Z. (2022). *The Organizational Evolution of Small Food Buying Clubs*. PhD thesis, University of Maine, Orono, Maine.
- Little, R., Maye, D., and Ilbery, B. (2010). Collective purchase: Moving local and organic foods beyond the niche market. *Environment and Planning A*, 42(8):1797–1813.
- Lofton, S., Kersten, M., Simonovich, S. D., and Martin, A. (2022). Mutual aid organisations and their role in reducing food insecurity in Chicago's urban communities during COVID-19. *Public Health Nutrition*, 25(1):119–122.
- Mao, C. X. and Song, W.-L. (2021). Does Reciprocity Affect Analysts' Incentives to Release Timely Information? Evidence from Syndication Relationships in Securities Underwriting. *Management Science*, 67(1):617–639.
- Mesoudi, A., Chang, L., Dall, S. R. X., and Thornton, A. (2016). The evolution of individual and cultural variation in social learning. *Trends in Ecology and Evolution*, 31(3):215–225.
- Molm, L. D. (2010). The Structure of Reciprocity. *Social Psychology Quarterly*, 73(2):119–131.
- Norris, J. R. (1997). *Markov Chains*. Cambridge University Press, Cambridge, UK.
- Nowak, M. A. (2006). Five Rules for the Evolution of Cooperation. *Science (New York, N.Y.)*, 314(5805):1560–1563.
- Nowak, M. A. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437(7063):1291–1298.
- Ohtsuki, H., Hauert, C., Lieberman, E., and Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505.
- Ostrom, E. (1990). *Governing the Commons*. Cambridge University Press. Google-Books-ID: daKNCgAAQBAJ.

- Paciotti, B., Richerson, P., Baum, B., Lubell, M., Waring, T., McElreath, R., Efferson, C., and Edsten, E. (2011). Are Religious Individuals More Generous, Trusting, and Cooperative? An Experimental Test of the Effect of Religion on Prosociality. In *The Economics of Religion: Anthropological Approaches*, volume 31, pages 267–305. Emerald Group Publishing Limited.
- Pesämaa, O., Pieper, T., Vinhas da Silva, R., Black, W. C., and Hair, J. F. (2013). Trust and reciprocity in building inter-personal and inter-organizational commitment in small business co-operatives. *Journal of Co-operative Organization and Management*, 1(2):81–92.
- Pfeiffer, T., Rutte, C., Killingback, T., Taborsky, M., and Bonhoeffer, S. (2005). Evolution of cooperation by generalized reciprocity. *Proceedings of the Royal Society B: Biological Sciences*, 272(1568):1115–1120.
- Phillips, P. C. B. and Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2):335–346.
- R Core Team (2022). R: A language and environment for statistical computing. R Foundation for Statistical Computing.
- Rand, D. G., Dreber, A., Ellingsen, T., Fudenberg, D., and Nowak, M. A. (2009). Positive Interactions Promote Public Cooperation. *Science (New York, N.Y.)*, 325(5945):1272–1275.
- Rand, D. G., Nowak, M. A., Fowler, J. H., and Christakis, N. A. (2014). Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences*, 111(48):17093–17098.
- Reiss, J. (2021). Public Goods. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition.
- Richerson, P., Baldini, R., Bell, A. V., Demps, K., Frost, K., Hillis, V., Mathew, S., Newton, E. K., Naar, N., Newson, L., Ross, C., Smaldino, P. E., Waring, T. M., and Zefferman, M. (2016/ed). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 39.
- Roberts, G. (2008). Evolution of direct and indirect reciprocity. *Proceedings of the Royal Society B: Biological Sciences*, 275(1631):173–179.
- Sandroni, A. (2000). Reciprocity and Cooperation in Repeated Coordination Games: The Principled-Player Approach. *Games and Economic Behavior*, 32(2):157–182.
- Schmid, L., Chatterjee, K., Hilbe, C., and Nowak, M. A. (2021). A unified framework of direct and indirect reciprocity. *Nature Human Behaviour*, 5(10):1292–1302.
- Smith, A. (2013). Estimating the causal effect of beliefs on contributions in repeated public good games. *Experimental Economics*, 16(3):414–425.
- Smith, K. M., Larroucau, T., Mabulla, I. A., and Apicella, C. L. (2018). Hunter-Gatherers Maintain Assortativity in Cooperation despite High Levels of Residential Change and Mixing. *Current Biology*, 28(19):3152–3157.e4.
- Song, J., Zhong, H., and Wang, H. (2019). Ckmeans.1d.dp: Optimal, Fast, and Reproducible Univariate Clustering.
- Spedicato, G. A., Kang, T. S., Yalamanchi, S. B., Thoralf, M., Yadav, D., Cerdón, I., Jain, V., Giorgino, T., Bilderbeek, R. J. C., Ebbert, D., and Maheshwari, S. (2021). Markovchain: Easy Handling Discrete Time Markov Chains.
- Steinley, Douglas. (2006). K-means clustering: A half-century synthesis. *British Journal of Mathematical and Statistical Psychology*, 59(1):1–34.
- Steinman, J. (2020). Consumer Food Co-ops in the Age of Grocery Giants. *American Journal of Economics and Sociology*, 79(3):833–875.
- Thöni, C. and Volk, S. (2018). Conditional cooperation: Review and refinement. *Economics Letters*, 171:37–40.
- Tremblay, E. (2017). *Splitting Together: The Evolution of Cooperation in Food Buying Clubs*. PhD thesis, The University of Maine.
- Tremblay, E., Hupper, A., and Waring, T. M. (2019). Co-operatives exhibit greater behavioral cooperation than comparable businesses: Experimental evidence. *Journal of Co-operative Organization and Management*, 7(2):100092.
- Tremblay, E. and Waring, T. M. (2015). A smiling face is half the meal: The role of cooperation in sustaining maine’s local food industry. *Maine Policy Review*, 23(2):43–50.
- Trivers, R. L. (1971). The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(1):35–57.

- Volk, S., Thöni, C., and Ruigrok, W. (2012). Temporal stability and psychological foundations of cooperation preferences. *Journal of Economic Behavior & Organization*, 81(2):664–676.
- Waring, T., Lange, T. Z., and Chakraborty, S. (2022). Institutional adaptation in the evolution of the 'Co-Operative principles'. *Journal of Evolutionary Economics*, 32(1):333–365.
- Wickham, H., Francois, R., Henry, L., and Müller, K. (2017). *Dplyr: A Grammar of Data Manipulation*.
- Wickham, H., Henry, L., and RStudio (2018). *Tidyr: Easily Tidy Data with 'spread()' and 'gather()' Functions*.
- Wilson, D. S. (2015). *Does Altruism Exist?: Culture, Genes, and the Welfare of Others*. Yale University Press.
- Wilson, D. S., Ostrom, E., and Cox, M. E. (2013). Generalizing the core design principles for the efficacy of groups. *Journal of Economic Behavior & Organization*, 90:S21–S32.
- Wu, S.-J. and Chu, M. T. (2017). Markov chains with memory, tensor formulation, and the dynamics of power iteration. *Applied Mathematics and Computation*, 303:226–239.
- Zelmer, J. (2003). Linear Public Goods Experiments: A Meta-Analysis. *Experimental Economics*, 6(3):299–310.