Asymmetry in the Development of Cooperative and Antagonistic Relationships. A Model-Based Analysis of a Fragile Public Good Game Experiment.

Ben Loerakker[1†], Nadège Bault[1,2], Maximilian Hoyer[1] and Frans van Winden[1]

23-12-2016

[1] Center for Research in Experimental Economics and political Decision-making (CREED), Amsterdam School of Economics, University of Amsterdam, 1018WB Amsterdam, the Netherlands

[2] Center for Mind/Brain Sciences (Cimec), University of Trento,Via Delle Regole 101, 38123 Mattarello, Italy

[†] Corresponding author (B.A.Loerakker@uva.nl)

1 Introduction

Public good experiments typically display more cooperation than predicted by rational and selfish preferences. The cooperation levels, though, vary depending on the exact design, per individual, and are often diminishing over time.[1] Various explanations have been offered, among which: altruism (Andreoni and Miller, 2002), reciprocity (Falk and Fischbacher, 2006), learning (Roth and Erev, 1995), and conditional cooperation (Fischbacher et al., 2001). Recently, both direct behavioral (van Dijk et al., 2002; Bault et al., 2016) as well as neurobiological (Bault et al., 2015) evidence has been provided for an alternative explanation involving the affective tie-mechanism introduced by van Dijk and van Winden (1997). Their model is characterized by affective interpersonal ties. Simply put, a person (i) takes the welfare of another person (j) into account according to the behavior of this other person. This is a dynamic proces whereby every action of j is compared to a reference action, if this action is more beneficial to i than the reference action a positve affective impuls is generated that could change the importance i attaches to the welfare of j. This means that this process is not only dynamic but also allows for an asymmetry between the development of the weight i attaches to the welfare of j and the weight j attaches to the welfare of i.

Although this model seems rather succesful in tracking behavior in the public good games examined, there are several issues that need to be addressed. First, work on the tie mechanism so far has focused on cooperative interpersonal relationships, leaving negative (hate) relationships underexplored. This is important because there exists a lot of eveidence that negative behaviors and (hate) relationships exist.[2] . Second, it is not examined whether people react differently to positive versus negative behavior of others. The difference in direction is obvious, but how about the size of the action? Third, the tie model has not been investigated in a horse race with other models using out-of-sample predictions regarding the same game[3]. Fourth, it is not clear whether the tie model would be helpful in explaining behavior across different contexts. For example could it integrate behavior

[1]See Chaudhuri (2011) and Plott and Smith (2008) for an overview

[2]Abbink and Sadrieh (2009) and Bosman and van Winden (2002) show that people are willing to destroy other people's earnings even if it does not lead to higher earnings for themselves. In the same experiment Bosman and Van Winden also show that negative emotions are involved if money is taken away from participants and when players engage in the destruction of the others' earnings by destructing their own earnings. More recently Bolle et al. (2013) show that participants are willing to decrease the chances of others to win a prize. Furthermore they find that in a repeated game setting players retaliate harmful actions and that this retaliation is driven by negative emotions caused by these harmful actions.

[3]See Bault et al. (2016)

observed in puclic good games with behavior in repeated prisoner's dilemma games?[4]

In this paper we will address each of these four issues. To adress the first two, a novel game design is used: the fragile public good game (FPG) game. In this game there is as much room for antagonistic behavior as there is for cooperative behavior. The third issue is addressed by designing an experiment with two independent parts, which allows us to predict behavior in the second part using parameter estimates from the first part. This creates proper out-of-sample predictions for the different social preferences as well as learning models that will be explored. Finally, we show that a simple two parameter tie model can mimic observed behavioral rules like tit-for-tat and is able to explain why and how players switch rules as the parameters of a prisoner's dilemma (PD) game change. Next, the parameters estimated on the behavioral data from the FPG game are investigated to see what kind of behavior they would predict in different repeated PD settings.

Furthermore, earlier findings in terms of the relative importance of the different tie parameters are confirmed. We find evidence that people react stronger to positive behavior of others than to negative behavior. This might be one of the driving factors of repeated cooperation. Another important finding is, that the tie model predicts significantly better than other models, including of social preferences and the reinforcement learning model of Roth and Erev (1995). Finally, a tie model with just two parameters seems well able to explain results found in different repeated PD environments.

Section 2 introduces the FPG game, provides a theoretical analysis using the tie model and presents our hypothesis. Section 3 describes the experimental and estimation methods used, while section 4 presents our results. Section 5 applies the tie model to the repeated prisoner's dilemma game, and section 6 concludes.

2 Theory

2.1 Fragile Public Good Game

In order to experimentally address the questions raised regarding negative ties we designed the Fragile Public Good (FPG) game, a two-player game that allows players to financially hurt as well as help the other player. A key feature of the FPG game is that it gives as much room for destructive behavior (taking) as for constructive behavior (contributing) ragarding

[4]See for instance Dal Bó and Fréchette (2011) and Fudenberg et al. (2012)

a public good. This is achieved by having both the (standard) Nash equilibrium and the status quo i.e., the initial allocation of tokens to the common account in the middle of the action space. With, in addition, full symmetry in the marginal cost of taking and contributing, this leaves substantial leeway for the development of negative as well as positive ties. There are relatively few public good experiments with an interior Nash equilibrium (Laury and Holt, 2008), and, to the best of our knowledge, no such experiments that allow for as much destructive as cooperative behavior. This game enables us to estimate the parameter values of our model. Furthermore, by maintaining comparability with ordinary (non-linear) public goods games, we can compare our results with existing studies of such games. By using a repeated game where in the first part players interact with a fixed partner, but are then rematched randomly with a new partner for playing in the second part, we can investigate the out-of-sample predictive performance of our estimated model.

More specifically, both players in our FPG game are endowed with 7 tokens in their private account, while sharing a common account containing 14 tokens at the beginning of every round. Each token stored in the private account generates 10 MU for the player concerned, whereas a token in the common account generates 10 MU for both players. Each round, both players simultaneously decide whether to contribute tokens to the common account or to take tokens from the common account. They can transfer up to 7 tokens per round from the common account to their private account, or the other way around.

Transferring tokens in either direction comes at a marginal cost, that increases with 2 MU per token. The transfer of the first token thus costs 2 MU, transferring a second one costs an additional 4 MU (for a total of 2+4= 6 MU), transferring a third token leads to a total cost of 12 MU (2+4+6), and so forth. The effect of contributing the first token is, thus, that the other player receives 10 MU while the contributing player gets 2 MU less. By contributing a second token a player generates another 10 MU for the other player at a cost of 4 MU, etcetera, until the seventh token which earns the other player still 10 MU while it costs the contributing player 14 MU to transfer. Taking tokens has the exact same effect on the transferring player as contributing the same amount of tokens would have. For the other player, however, the effect is the exact opposite: He or she will now lose 10 MU per token instead of gaining 10 MU. Because the only difference between taking and contributing concerns the development of, respectively, a negative and a positive externality, this game allows us to study, in a clean way, whether an asymmetry exists in the impact of hurting behavior (taking) and helping behavior (contributing).

Making no transfer may be seen as a reference point as it accords with the status quo as well as the standard Nash best response. Moreover, it may easily attract a player's attention in the payoff matrix of the game (see appendix D). We will return to this below. The game is a non-linear public good game with an internal social optimum, where both players contribute either 4 or 5 tokens. The similarities with a more conventional public good game become even clearer when one sees taking seven tokens as the starting point, so one can only contribute. In that case the stage game becomes similar to a public good game with diminishing returns to contributing, albeit with an internal standard Nash equilibrium and an internal social optimum.

2.2 Model

We use an adapted version of the tie model of Bault et al. (2016). In this model α_{ijt} captures the tie that i has with j at time t, and formally expresses the weight that i attaches to the utility (payoff) of j. These ties are personal, dynamic and do not necessarily have to be symmetric.

We start from the basic model in which players have the following interdependent utility function:

$$V_{it} = U_{it} + \alpha_{ijt} U_{jt} \tag{1}$$

Here V_{it} denotes the (extended) utility function of player i at time t, while U_{it} and U_{jt} indicate the payoffs of i and j, respectively, at time t.

Players do not only take the current period into account, but also the subsequent one (one-period forward looking behavior). Empirical evidence suggests that players are rather myopic (see e.g.Bone et al. (2003) and Bone et al. (2004)). This leads us to the following simple extension of eq. (1):

$$V_{it} = U_{it} + \alpha_{ijt} U_{jt} + (U_{it+1} + \alpha_{ijt} U_{jt+1}) \tag{2}$$

For the FPG game, letting C_{it} stand for i's contribution to the common account, i's expected payoff (U^e_{it}) can be written as :

$$U^e_{it} = 210 - C^2_{it} - |C_{it}| + 10C^e_{jt}$$
$$\text{With} \tag{3}$$
$$-7 \leq C_{it} \leq 7$$

Including future periods in the player's utility function does not affect C_{it} if i does not expect to be able to influence j's next period contribution. Players may believe, however, that (some) other players are imitators or conditional cooperators (Fehr and Fischbacher, 2003). Therefore, we assume the following relationship until the last round (as there is no future left in the final round):

$$C^e_{jt+1} = \gamma_i C_{it} + (1 - \gamma_i)C^e_{jt}$$
$$\text{With } 0 \leq \gamma \leq 1 \tag{4}$$

From these equations it follows that the optimal contribution for player i depends on both the parameter γ_i, indicating how strong agent i believes he can influence agent j, and α_{ijt}, which expresses the weight i assigns to the payoff for j. We specify the latter as:

$$\alpha_{ijt} = \delta_{1i}\alpha_{ijt-1} + \delta_{2i}I_{ijt-1} \tag{5}$$

With I_{ijt-1} standing for an impulse determined by the difference between j's last round contribution and a reference contribution. In this paper, though, we will use the next, more general specification, which differentiates between positive and negative impulses:

$$\alpha_{ijt} = \begin{cases} \delta_{1i}\alpha_{ijt-1} + \delta_{2Ni}I_{ijt-1} & I_{ijt-1} \leq 0 \quad \text{(6a)} \\ \delta_{1i}\alpha_{ijt-1} + \delta_{2Pi}I_{ijt-1} & I_{ijt-1} > 0 \quad \text{(6b)} \end{cases}$$

Here, we assume that I_{ijt-1} equals the difference between the other player's contribution and the one-shot Nash equilibrium choice (0), based on the discussion above regarding the reference point (see also the estimation results in Appendix B). The tie meachanism described above can also be interpreted as a information extraction mechanism, used to determine if the other player is a friend or a foe (Bault et al., 2015). A higher positive and negative impulse parameters can then be seen as s preference for respectively cooperative or destructive behavior.

2.3 Model analysis and hypotheses

An equilibrium is defined by a situation where both players have no incentive to change their contribution. We will now discuss the conditions and nature of potential equilibria. To that purpose, we start by comparing the (expected) utility of two adjacent choices. Due to the fact that $V_{it}(\ C_{it})$ is concave, C_{it} is a best response if $V_{it}(C_{it}) \geq V_{it}(C_{it}+1)$ and $V_{it}(C_{it}) \geq V_{it}(C_{it}-1)$. Assuming here, for convenience, that both C_{it} and C^e_{jt} are greater than or equal to zero,[5] and omitting the subscripts of α and γ, the difference in utility equals:

$$V_{it}(C_{it}+1) - V_{it}(C_{it}) = 10\alpha + 10\gamma - (2C_{it}+2) - \gamma\alpha(2\gamma C_{it} + 2(1-\gamma)C^e_{jt} + \gamma + 1) \quad (7)$$

Eq. (7) shows the costs and benefits of contributing an extra token (in the positive domain). If players are not playing strategically the costs are simply $2C_{it}+2$, while the benefits are 10α. If players expect to be able to influence their counterpart the cost-benefit analysis becomes more complicated. There are benefits of 10γ, from the expected imitation or positive reciprocity by the other, as well as new costs of $\gamma\alpha(2\gamma C_{it} + 2(1-\gamma)C^e_{jt} + \gamma + 1)$, as players with $\alpha > 0$ care about the fact that the other faces a cost of reciprocating or imitating.

The following propositions can de proven: First of all, it turns out that contributions outside of the interval $[-5, 5]$ can never be part of any equilibrium for conventional values of α between 1 and -1. This result is important as it shows that the bounds of the decision space are not part of any equilibrium in that case, which is helpful for estimating the model. For instance, suppose we would like to estimate α in the myopic model, then, if a player repeatedly made boundary decisions ($C = 7$ or $C = -7$) we would only have information about respectively, the lower bound and the upper bound of the α parameter.

Proposition 1. *Contributions outside of the interval* $[-5, 5]$ *can never be part of any equilibrium if* $-1 \leq \alpha \leq 1$.

[5] In Appendix A also addresses the case where C_{it} is smaller than zero

Proof. See Appendix A.1.

Next, focusing first on symmetric equilibria ($C_{it} = C_{jt}$), we arrive at the following proposition:

Proposition 2. *Any contribution level where $C_{it} = C_{jt} \in [-5, 5]$ can be part of an equilibrium.*

Proof. See Appendix A.2.

An asymmetric equilibrium is less likely as one player is then always worse of than the other player. Theoretically, asymmetric equilibria are not impossible, though, but the parameter constraints are more restrictive than for symmetric ones. The farther the different contributions are apart the more extreme the conditions for these equilibria become. The following proposition refers to their existence:

Proposition 3. *Asymmetric equilibria exist if either $-5 \leq C_{it}, C_{jt} < 0$ or $0 \leq C_{it}, C_{jt} \leq 5$.*

Proof. See Appendix A.3.

Our last proposition establishes a parameter restriction for efficient cooperation. For convenience, we restrict ourselves here to the myopic model as this will turn out to be the most relevant model in our study. Similar restrictions including γ could be derived for the model that allows for forward looking behavior (see Appendix A.2).

Proposition 4. *For the socially optimal choices to be part of an equilibrium under the myopic model, the parameters of both players tie mechanism should satisfy the restriction: $0.2 \leq \frac{\delta_{2i}}{1-\delta_{1i}} \leq 0.25$, which is a necessary but not a sufficient condition.*

Proof. See Appendix A.4.

The intuition for this results is that if players have a $\frac{\delta_2}{1-\delta_1}$ ratio that is below 0.2 they built insufficiently strong ties. If the ratio is larger than 0.25 the opposite happens: the ties become so strong that α will grow larger than 1 implying that players will overinvest in this relationship.

Based on the propositions 1,2, and 3 our first hypothesis is:

Hypothesis 1. *If both players in a dyad do not change their contribution for multiple consecutive rounds, both contribute an equal amount and this contribution lies between -5 and 5 (inclusive)*

Our second hypothesis is motivated by the earlier mentioned work of Baumeister et al. (2001) and Baumeister and Leary (1995). They collect evidence that indicates that negative experiences that coincide with negative emotions have a stronger and longer lasting impact on someone's wellbeing and behavior than positive experiences and emotions. Furthermore, Kuhnen (2015) finds that investors weigh negative news more than positive news in an experimental setting.

Hypothesis 2. *Negative impulses have a bigger impact on the weight a player allocates to the payoff of a counterpart (the social tie) than positive ones, or in the context of our model:* $\delta_{2N} > \delta_{2P}$.

The third hypothesis concerns the performance of the model, specifically its predictive accuracy. Bault et al. (2016) already investigated the comparative performance of a ties model with $\delta_{2N} = \delta_{2P}$ within sample, where the number of parameters could be an issue. Here we apply a true out-of-sample test and compare with alternative models, now including a learning model. Therefore the final hypothesis reads:

Hypothesis 3. *When calibrated on the first FPG game our ties model gives more precise estimates of a subject's behavior in the second FPG game, as compared to competing models that are calibrated on the same data.*

Our final hypothesis is based on survey papers by Chaudhuri (2011) and Kagel et al. (1995). They find that in most public good games contribution levels are declining. If we relate these results with proposition 4, we hypothesize that most subjects will not fullfill the restrictions outlined in this proposition:

Hypothesis 4. *For the majority of the subjects* $0.2 \leq \frac{\delta_{2i}}{1-\delta_{1i}} \leq 0.25$ *will not hold.*

3 Methods

3.1 Experiment

The experiment took place in November 2012 and April 2013. It consisted of 3 sessions with 130 (65 female, 2 unreported) participants. All subjects were recruited through the recruitment system of the CREED laboratoryof the university Amsterdam. Students who had participated in previous public good experiments or power-to-take experiments (as recorded in the CREED recruitment system) were excluded.

The entire experiment was held in the CREED laboratory and completely computerized. In the experiment we used Monetary Units (MUs) to express the earnings of the participants, which were converted to euros by a rate of 700 MU to one euro. The average earnings in the experiment were €25.65 (there was no show-up fee as the theoretical minimum earnings exceded €10, no participant earned less than €15) and the sessions took about two hours.

During the experiment the participants were first asked to perform a Social Value Orientation (SVO) test (Liebrand and McClintock, 1988), where we use the version of van Dijk et al. (2002). In this test participants decide on payoff allocations between Self and an anonymous Other. MUs allocated to Other affected the earnings of a random other participant in the experiment. Participants were informed that all their choices in the SVO test remained confidential and only learned their earnings at the end of the experiment. An example of the choices made in an SVO test can be found in Appendix D.

Every question of the SVO test concerns a choice between two payoff allocations. Each allocation represents a point on a circle around the origin, where the payoff to self is on the x-axes and the payoff to Other is on the y-axes. In total the participants had to make 32 of these choices in the SVO test. An angle is constructed by aggregating all the vectors spanned up by the 32 chosen payoff allocations. An individual's distributional preferences can be expressed by this angle. For example, an angle of zero degrees means that one is completely selfish, a 45 degree angle indicates that one maximizes the sum of the payoffs to Self and Other, and an angle of 90 degrees would indicate that one only cares about the payoff of Other. The size of the vector tells us how consistent the choices are. If all choices are consistent with a certain preference the size of the vector will be 1000. In the examples given above, the tangent of the angle is always positive. However, just as with the alpha in our theoretical model also negative values are possible. The interpretation of these values is analogous to a negative α, as they indicate that a person is willing to give something up in order to decrease the payoff for Other. [6] The tangent that results from the test can be interpreted as an indication of the initial alpha (α_0) and will be used as such later on.

[6]situations where individuals prefer negative payoffs over positive payoffs for themselves are not taken into consideration here and very seldom observed.

After this SVO test the participants played 35 rounds of the Fragile Public Good (FPG) game, explained above, in a partner setting (with a different partner than the "other" from the SVO task). In the introduction of the FPG game it was made clear that both taking and contributing came at a cost. In order to check if players understood the game, they had to answer quiz questions and played three trial rounds. In these trial rounds they could also get acquainted to the feedback they would receive during the actual game. After every round they saw the choice of the other player, their own payoff in the round they just played and the payoff of the other, both of which were represented using numbers as well as bars so as to visualize the difference between the payoffs.

After the first FPG game the participants were informed that a second one would follow, again with a randomly matched partner but not the one from the previous game. Also this game consisted of 35 rounds. The final task of the experiment task was another SVO test, where the other in the test was now the same as in the final FPG game. This final test will not be used in this paper.

3.2 Estimation

For our estimation procedure we follow Bault et al. (2016). To close the model and to enable us to estimate it we introduce a random variable ϵ_{ik}/θ_i as a noise term. Where θ_i represents the rationality or choice intensity of player i. If we now assume ϵ_{ik} to be i.i.d. and double-exponentially distributed, we arrive at a multinomial logit model. Now let π_{ikt} be the probability that a player chooses contribution k in period t, then if we multiply all these probabilities we obtain our likelihood measure:

$$\prod_t \pi_{ikt} = \prod_t \frac{e^{\theta_i V^e_{ikt}}}{\sum_h e^{\theta_i V^e_{iht}}} \qquad 0 < \theta < \infty \tag{8}$$

Estimation requires a value for α_0, the tie parameter prior to any interaction with the other player. In the estimation results shown in the subsequent sections we used the measure taken from the SVO test. For those participants with an inconsistent tie measure (the tie measure is considered inconsistent when the length of the vector is smaller than 600) we use $\alpha_0 = 0$. The model is estimated on the first FPG game.

In the group level estimations we set δ_{1i}=δ_{1j}, δ_{2i}=δ_{2j} and θ_i=θ_j for all i and j and estimate the model using Matlab's fmincon optimization procedure based on the likelihood described in equation (8). When calculating standard errors we clustered all observations from the same individual.

4 Results

4.1 Descriptive statistics

The average angle of 128 participants in the SVO test was 6.03 degrees, which corresponds to an α value of 0.11. The observations of 2 participants were lost due to technical problems, while the choices of 8 participants were considered inconsistent because their vectors were below 600 out of 1,000 in length (see Liebrand and McCLintock (1998)). The SVO tests concerning those participants are, therefore, deleted from the analyses.

A summary of the behavior during the FPG games is given below in table 1. Looking at table 1 there are some noteworthy results. First of all, we observe that average contributions are noticeably higher in the second FPG game than in the first FPG game:

Table 1: Descriptive statistics

Game	FPG1 (n=130)	FPG2 (n=130)
Average contribution	2.28	2.86*
Avg contribution first round	1.26	2.53*
Avg contribution last round	0.68	0.70
% negative contributions	11.3%	4.6%

Note: * indicates significance at the 1% level, using a Wilcoxon sign-rank test with contributions on the pair level.

This is also illustrated by figure 1, that shows the average contributions per round:

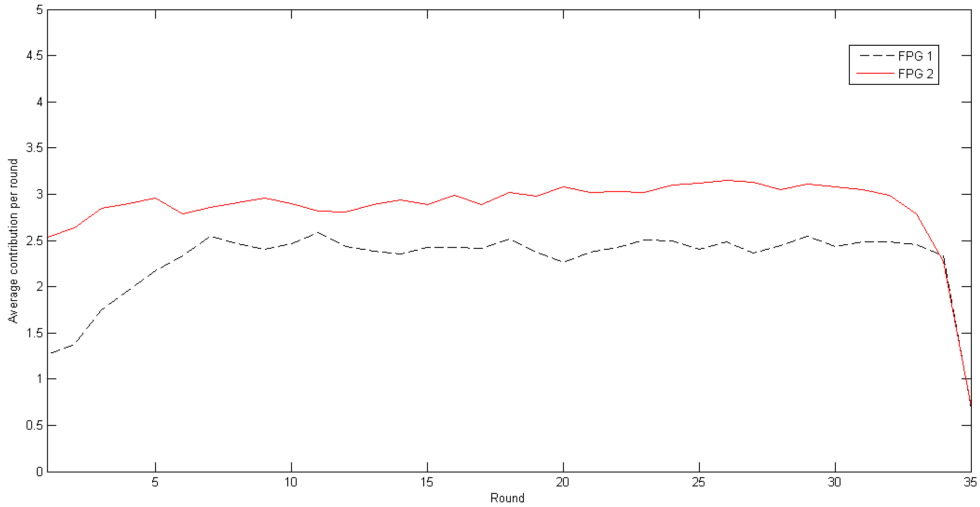

Figure 1: Average contributions per round

The figure above suggests that behavior in the second game is definitively influenced by the first game, as participants start off with much higher contributions. Furthermore, we observe that between rounds 7 and 33 the difference between the games is fairly constant at around 0.5, until the decline in the last couple of rounds (end-effect) leads to almost identical contributions in the end.

Another difference between the two games concerns the number of destructive decisions. While taking seems to play an important role in the first FPG game, its role in the second one is diminished noticeably. The first game, though, shows that destructive behavior can be relatively frequent even if there are plenty of opportunities to stay away from it. For illustration, figure 2 shows two pairs that, respectively, establish a cooperative and a sour relationship.

We find that only 1.5% (135 out of 9100) of the contributions are either larger than 5 or smaller than -5. Moreover, we do not find any instance of two players contributing an unequal but constant amount for 3 rounds or more. This confirms our first hypothesis (H1).

The rest of this section is organized as follows: Section 4.2 investigates the group level estimation results, 4.3 is devoted to individual level estimates, while section 4.4 studies the predictive performance of the model.

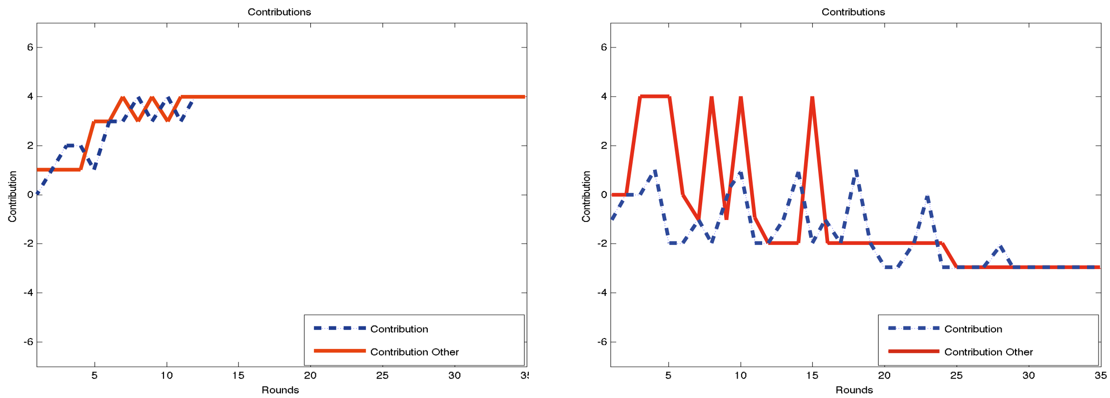

(a) Positive ties formed

(b) Negative ties formed

Figure 2: An example of a succesfully built cooperative relationship and a relationship that has turned sour

4.2 Group level results

We start by estimating the myopic version of the model (labeled Myopic), represented by eq. (1), neglecting for the moment forward looking behavior introduced in eq. (2). This leaves a model with only 3 parameters. Despite its simplicity, this model has been found quite succesful in explaining public good contributions (Bault et al., 2016).

Psychological studies (Baumeister et al. (2001)), suggest that people react to negative experiences and emotions differently than to positive ones. Therefore, in our next model, we will allow for differences between the impact of negative and positive impulses, again using the myopic version of the model (M.NP), as formulated by eqs. (6a) and (6b).

The third model that we will investigate is the forward-looking model (FL), represented by eqs. (2) and (4). And, finally, we estimate a model that allows for the aforementioned difference between the strength of negative and positive impulses as well as forward looking behavior (FL.NP).

The estimation results regarding these four models are found in table 2.

Table 2: Group level estimations

	γ	θ	δ_1	δ_2	δ_{2P}	δ_{2N}
Myopic		0.16*(0.02)	0.54*(0.09)	0.10*(0.02)		
M.NP		0.17*(0.02)	0.49*(0.10)		0.12*(0.02)	0.08*(0.02)
FL	0.06(0.06)	0.16*(0.02)	0.54*(0.10)	0.10*(0.02)		
FL. NP	0.03(0.06)	0.17*(0.03)	0.49*(0.12)		0.11*(0.02)	0.08*(0.02)

Note: standard errors are between brackets;
* indicates significance at the 1% level.

The myopic model, as well as the other models, estimates θ to be around 0.16. To give an idea about its interpretation, the predicted chance that a player with an α-value of zero contributes zero is estimated to be about 30%. The chance that a player contributes one as well as the chance that a player takes one is estimated to be about 20%, the chance of contributing and taking 2 is about 10%, while all the other contributions together take up the remaining 10% probability. For comparison, if θ would be 0 all choices have a probability of $\frac{1}{15}$ ($< 7\%$), while if θ goes to infinity the probability of a player chosing zero goes to 1. δ_1 is estimated to be close to $\frac{1}{2}$, so if the other contributes zero the valuation of the payoff of the other is halved. δ_2 is estimated to be 0.1 which means that a contribution of 5 would lead to an α of around 0.5, if the initial α-value is zero.

If we allow for a dichotomy between positive and negative impulses (M.NP) we find similar θ and δ_1 values. However, δ_{2P} seems larger than δ_{2N}. The improvement in the likelihood is significant at the 10% level (p≈0.09) even if we take just the average improvement per subject as a single observation (and significant at the 1% level if all rounds from all players are taken into account).[7]

We next consider the forward-looking model (FL), but neglect the potential difference between positive and negative impulses for the moment. We find that γ is insignificant. Only a modest improvement in the likelihood is obtained when compared with the improvement caused by an extra impulse parameter in the myopic model (p>0.50 if we take the average improvement per individual, $p \approx 0.10$ if we take all contributions into account).

[7] It should be noted that the difference between δ_{2P} and δ_{2N} reported in table 2 is a conditional result. Not all participants in our experiment where exposed to negative contributions and the ones that were exposed to them are not an exogenously chosen or created group. The players that at any point in time are faced with negative contributions are often also the ones that made negative contributions themselves. In other words, this behavior shows up pairwise and pairs that make negative contributions are likely to have different characteristics then pairs that do not make such contributions.

The other parameters, θ, δ_1 and δ_2 are very similar to the values found for the myopic model.

The full model (FL.NP) shows results that can be seen as a combination of the results found for M.NP and FL. Positive impulses are again stronger in impact than negative impulses, with parameter values similar to the ones of the myopic model, while δ_1 is again around 0.5.

An interesting result is that δ_1 is estimated to be close to 0.5 in all model specifications, which is consistent with earlier findings for two-player public good games (Bault et al., 2016). The fact that this finding is replicated could indicate that in this type of environment people, at least, weigh their history about as much as new information. It is also noteworthy that at the group level $0.2 < \frac{\delta_{2P}}{1-\delta_1} < 0.25$ always holds, while $0.2 < \frac{\delta_{2N}}{1-\delta_1} < 0.25$ does not. This indicates that our fourth hypothesis (H4) might not be correct and suggests that (most, not all) people are able to form stable cooperative relationships, but do not sustain long destructive relationships.

Furthermore, it is worth noting that we find that positive impulses seem to have a stronger effect on the α parameter than negative impulses. This seems to contrast with the findings summarized by Baumeister et al. (2001). However, there are studies that find that the influence of a positive signal might weigh stronger than that of a negative signal, see for instance King-Casas et al. (2005) and Rand et al. (2009). Another aspect could be the earlier mentioned finding that people are not only more affected by negative experiences but that they are also more motivated to get out of a negative situation. This result suggests that we should reject our second hypothesis (H2).

Note, furthermore, that the difference between δ_{2P} and δ_{2N} could be influenced by the fact that the myopic model does not allow for any forward-looking behavior. If (some) players are in fact forward looking, this might be partially captured by the δ_2-parameter(s). If this was the case it would lead to a bias in the δ_2-parameter(s), where the effect on δ_{2N} would be negative while the effect on δ_{2P} would be positive. The reason is that if a player is forward looking he or she wants to contribute more than if a player is not (as it is assumed that the other will positively react). This positive effect on the contributions will lead to δ_{2P} being higher, while δ_{2N} will be estimated to be lower (so that the effect will be less negative). However, we see that also in the full model (FL.NP) this difference between positive and negative impulses still exists. This directly opposes our second hypothesis (H2).

Moreover, the parameter γ, which is to capture the forward-looking behavior, is never significantly different from zero at the 5% level.

4.3 Individual level results

From the individual level estimates we can see how many of our participants are able to maintain stable cooperative or destructive relationships. We start by evaluating the myopic model. We find that 104 out of 130 participants meet the conditions mentioned in Proposition 4. This means that 80% of our subjects are able to build sufficient ties to sustain cooperation. Of the remaining 26 participants, 10 have a ties mechanism that is too strong to be efficient (they are not able to sustain an efficient cooperative relationship), while the other 16 have an insufficiently strong tie mechanism. We find that in total 72 (36 pairs) out of these 104 participants are in fact cooperating efficiently in the final 10 rounds (that is, in more than six out of the last 10 rounds both contribute equally and either 4 or 5 tokens). This is in line with proposition 4, stating that $0.2 < \frac{\delta_{2i}}{1-\delta_{1i}} < 0.25$ is a necessary but not a sufficient condition for stable and efficient cooperation. It, however, contradicts our fourth hypothesis (H4).

When we allow for different parameters for positive and negative impulses we find -as we did on the group level- for most participants the positive impulse parameter is larger than the negative one: for 37 out of 130 participants $\delta_{2P}>\delta_{2N}$, for 15 participants $\delta_{2N}>\delta_{2P}$, and for two participants $\delta_{2N}=\delta_{2P}=0$. Two participants did not receive any impulses, one received only negative impulses, and the remaining 73 participants encountered just positive impulses.

Shifting our attention to the forward looking-model now, we first investigate how many individuals have an estimate of γ that is significantly different from zero. It turns out that 79 out of 130 participants are indeed forward looking (γ positive at the 5% level). This seems to contrast with our previous finding at the group level, which might suggest a large heterogeneity among subjects in their forward-looking behavior.

When we compare the results of the full model, however, the same pattern as found for the group level is observed. Now, only 47 out of the 130 participants show forward looking behavior [8]. Moreover, we still find that for 88 out of the 130 participants $0.2 < \frac{\delta_{2Pi}}{1-\delta_{1i}} < 0.25$, while most players do not seem to be able to sustain negative relationships, as for only three participants we find that $0.2 < \frac{\delta_{2Ni}}{1-\delta_{1i}} < 0.25$. There are also ten participants that build excessively strong negative ties (i.e., α values smaller than -1). These findings might explain why we observe some prolonged intervals of negative contributions and the existence of sour relationships, as illustrated in figure 2[9]. Looking at the evidence presented at the group as well as at the individual level we must reject H4.

4.4 Predictive performance out-of-sample and model comparison

Now that we have estimated the Ties model both at the group level and the individual level, we put it to a more difficult test. We investigate if our model is not only able to explain behavior after the fact, but also to predict behavior in independent future rounds. Moreover, we will compare its predictive performance with the performance of three other models: the inequity-aversion model by Fehr and Schmidt (1999), the reinforcement learning model of Roth and Erev (1995), and a model with a fixed weight attached to the payoff of the other player (i.e., a fixed α).

In order to get truely out-of-sample predictions, we use the following procedure. We first estimate the myopic model at the group-level, allowing for different positive and negative impulse parameters. We choose the myopic model because the other models we compare our model with do not allow for forward-looking behavior either. Moreover, this does not affect the performance of our model too much as the additional parameter capturing this effect is insignificant. First we estimate the model on the group-level, then we take the contributions of the new other in the second game to calculate the α-values (for α_0 the values from the SVO test are used again), using the estimated parameters of the first game. The predicted action is the choice that generates the highest likelihood according to (8). Note that we do not re-estimate the model after every round and also do not readjust α on the basis of choices made by the participants themselves in the second game. This allows predictions to run away from the realized values. The fact that contributions in the second

[8] Not all participants encounter many negative impulses though, so this result is conditional on encountering enough negative impulses

[9] For more on this see Hoyer et al. (2014), where the occurrence of such relationships is further analyzed with different experimental designs. The results of all the estimations mentioned above are available online.

game were generally higher in the second game should make forecasting harder.

We use this procedure not only for the ties model but also for the other models of social preferences and the basic reinforcement learning model referred to above. To make forecasts for these models we use a similar procedure as described for the ties model.

For the fixed alpha model this means that we estimate the α parameter at the group-level on the behavioral data of the the first FPG game and then use this estimates this to predict the choices made in the second FPG game.

For the Fehr-Schmidt model we estimate, again at the group-level, the α and β parameters of the following expression for the expected utility of a particular choice k) V^e_{ikt}, using the behavioral data of the first FPG game and (8):

$$V^e_{ikt} = X^e_{ikt} - \alpha(X^e_{jt} - X^e_{it})_{X_{it}<X^e_{jt}} - \beta(X^e_{it} - X^e_{jt})_{X_{it}>X^e_{jt}} \tag{9}$$

Where $X^e_{ht}(h = i, j)$ denotes the expected payoffs calculated using either the expected contribution of the other in the same round (participants were asked for this after every choice made by themselves) or by the actual contribution of the other in the previous round. When estimating this model, we find β to be larger than α for both specifications of $X^e_{ht}(h = i, j)$, which is contrary to the predictions and findings by Fehr and Schmidt, but more often found in the literature (Yang et al., 2012). What is more problematic, though, is that both α and β are estimated to be larger than 1, again violating the assumptions of the model. Note that $\beta>1$ implies that one would prefer to throw away a dollar to diminish inequality with one dollar. Because of the clear lack of support for this model we will not further consider it below.

In the Roth and Erev model of reinforcement learning players learn the value of certain actions by playing them. The higher the payoff after playing a certain action the more this action gets reinforced, meaning that the probability of choosing this (or a similar) action increases. In the three parameters version of the model used here: s denotes the strength (or speed) of learning, indicating how much the chosen action is reinforced, ϕ stands for a decay effect that captures the speed by which the attraction of an action diminishes over time, and E denotes an experimentation effect that represents the reinforcement of adjacent choices. To get to estimates we use a similar procedure as explained in Erev and Roth (1998), meaning that all probabilities of an individual i choosing an action k at time

t (π_{ikt}) are inititially the same, as the attraction of each choice q_{ikt} is assumed to be the same q at the beginning of the game. After a choice is made (zero is chosen as the first prediction in this exercise) the distance (R) between the realized payoff and the minimal payoff, combined with the effect of the parameters, determine the new attraction of choices and thereby the choice probability distribution in the round thereafter. More precisely:

$$\begin{aligned}
\pi_{ikt} &= q_{ikt} / \sum_h q_{iht} \\
q_{ik1} &= q_{ih1} = q \\
q_{ikt+1} &= (1-\phi_i) q_{ikt} + E_{ikt} \\
E_{ikt} &= s_i R_{jt} (1-\epsilon_i) \text{ if } k = j \\
E_{ikt} &= s_i R_{jt} (\epsilon_i / 2) \text{ if } k = j \pm 1 \\
E_{ikt} &= 0 \text{ otherwise}
\end{aligned} \tag{10}$$

Another interesting model for explaining behavior in dynamic settings was introduced by Camerer and Ho (1999). Their Experience-Weighted Attraction (EWA) learning model not only allows for learning via payoffs, but also that players may learn over time what other players are likely to do. Although this kind of belief learning is undoubtedly important in many economic settings, it should not affect behavior in our game, as in our game the net return on a contribution vis-à-vis that of another contribution of a player does not depend on the contribution of the other player but only on his or her own contribution.

Finally, we mention the 'types' model of Levine (1998), a social preference model that may appear similar behaviorally, but is conceptually quite different from the ties model. In this model the weight an individual i attaches to the utility (u_j) of another individual j is dependent on one's own (constant) altruism parameter (α_i), the belief about the altruism parameter of the other (α_j) and a parameter λ that weighs both, such that i's utility (u_i) gets transformed into an extended utility, v_i:

$$v_i = u_i + \frac{\alpha_i + \lambda \alpha_j}{1 + \lambda} u_j \tag{11}$$

Although this model assumes unexplained fixed altruistic parameters, there is some similarity with the ties model. If the contributions of the other player are seen as signals of that player's altruism level, these signals would then change the belief of the other's altruism level and thereby the weight one attaches to his or her utility (α). Because the model does not specify the belief updating process, let alone how to apply it in our setting, we do not further consider it here.

Our discussion of social preferences models and the learning model of Roth and Erev, reflects that only a few social preference models available are able to make predictions for the dynamic behavior in our games. This is not surprising, as most theoretical models are not designed to explain dynamics.

For the predictions regarding the second game, we again use the choices with the highest likelihood of being chosen, given the parameters estimated on the behavioral data of the first game. Table 3 presents the mean absolute error and the mean squared error of the predictions:

Table 3: Out-of-sample prediction with group-level estimates

Model	Mean Absolute Error	Mean Squared Error
Ties Model	0.51 (0.53)	1.65 (2.39)
Fixed Alpha	1.92 (2.41)	6.97(14.05)
Roth and Erev	1.64 (1.07)	4.39 (4.59)

Note: standard errors using average errors per individual between brackets

From our results in table 3 we can conclude that the Ties model seems to perform best, supporting our fourth hypothesis. This is confirmed by Wilcoxon signed rank tests with the average error (for both squared and absolute errors) per individual as observations. These tests show that this model outperforms the other models when it comes to predicting ($p<0.01$ for all tests). Furthermore it is interesting to note that especially the reinforcement learning model by Roth and Erev does not do a good job when it comes to predicting behavior out of sample as it's mean absolute prediction error is more than three times higher than that of the Ties model.

As an alternative test we check how individual-level estimations perform. We use the same procedure as with the group-level estimates, but now each player's predicted choice is calculated using individual estimates. Table 4 shows the results.

Table 4: Out-of-sample prediction with individual estimates

Model	Mean Absolute Error	Mean Squared Error
Ties Model	1.00 (1.11)	3.58 (5.61)
Fixed Alpha	1.69 (1.72)	6.11(12.68)
Roth and Erev	3.01 (1.47)	11.63 (7.21)

Note: standard errors using average errors per individual between brackets

Again it turns out that the Ties model performs significantly better than the learning model and the fixed alpha model, supporting our fourth and final hypothesis. Note, though, that with this specification both the Ties and the reinforcement learning model, and especially the reinforcement learning model, performs worse than when group-level estimates are used. This may seem surprising, but is caused by the fact that some individuals experience very little variation in impulses in the first FPG game, making it difficult to estimate their individual parameters precise.

5 Applying the Ties Model to the Repeated Prisoner's Dilemma

The Ties model gives an explanation for the development of cooperation or antagonism that is quite different from the rest of the literature, as it focuses on changes in social preferences generated by interaction experiences rather than on given (fixed) social preferences or simple heuristics represented by automata. In this section we will explore a connection to another strand of research focusing on the evolution of behavior in repeated games, specifically studies by Dal Bó and Fréchette (2011) and Fudenberg et al. (2012). To understand the strategies people use when placed in environments that are either well- or ill-suited to generate cooperation, they have subjects play multiple repeated Prisoner's Dilemma (PD) games with continuation probabilities between 1/2 and 7/8. Using maximum likelihood estimation procedures, they estimate the share of a series of simple strategies, or automata, such as tit-for-tat (TFT), always defect (AD), and tit-for-two-tats (TF2T). Fudenberg et al. (2012) find that, if cooperation becomes more profitable, people become 'slower to anger' and 'faster to forgive', this is, they are more willing to allow a defection and pick up cooperation after only a few cooperative choices of the counterpart. This is reflected in the presence of a strategy like TF2T in these environments. This section serves to illustrate how different parameter combinations of the Ties model, and their estimates, can generate (or mimic) these strategies as well as a the shift towards more lenient and forgiving strategies as cooperation becomes more attractive. In Appendix C these arguments are worked out in more mathematical detail.

We start our Ties model-based analysis of the PD game by introducing other-regarding preferences. We do this by adding the α-weighted payoff of the other to a player's payoff. Starting from a general representation of a PD game without any other-regarding payoffs (see Table 5a), we apply these other-regarding preferences to two specific games with benefit/cost (b/c) ratios of 2 (Table 5b) and 4 (Table 5c). These examples are chosen for comparability with the games found in Fudenberg et al. (2012). What stands out from these new payoff matrices is that defecting is now no longer necessarily the dominant action. If α is lager than 1/2 (in Table 5b) or 1/4 (in Table 5c) cooperation becomes dominant. If we now define the impulse generated by a cooperative choice to be of size one and the impulse from defection by the other to be of size zero (as this is the Nash equilibrium action of the stage game), we can apply a similar model as the one we introduced for the public good game (see below).

Table 5: Prisoner's Dilemma (with other regarding preferences)

(a) b/c

	C	D
C	b-c	-c
D	b	0

(b) b/c=2

	C	D
C	1+1α	-1+2α
D	2-α	0

(c) b/c=4

	C	D
C	3+3α	-1+4α
D	4-α	0

Note: Table 5a gives the actual payoffs of player 1, while 5b and 5c give the valuation of these payoffs by a player that also cares about the other player. C stands for cooperation, and D for defection.

Both Dal Bó and Fréchette (2011) and Fudenberg et al. (2012) find experimental evidence that many subjects in their experiments use either a tit-for-tat (TFT) or a tit-for-two-tats (TF2T) strategy although these strategies are often not evolutionary stable (in a evolutionary game theory context).[10]

[10]TFT requires a player to start with choosing C and, subsequently, to choose whatever his opponent did in the previous round. Thus, after observing C (D) the player chooses C (D). If a player starts with D first instead of C, the strategy is labeled DTFT. TF2T requires a player to always choose C, unless his counterpart chose D in the previous two periods. If a player starts with D first instead of C, the strategy is labeled DTF2T. Dal Bó and Fréchette estimate that for their games with a continuation probability of 3/4 between 35% (if b/c$\approx$2) and 56% (if b/c $\approx$ 4) of subjects choose TFT. Fudenberg et al., who consider many more strategies and introduce noise, find between 19%, for b/c=1.5, and 7%, for b/c=4, of players

Below we will show that the simple and neurobiologically underpinned Ties model (Bault et al. (2015)) can help explain the behavior observed in these experiments.

The previously mentioned studies combine the simplicity and descriptive power of strategies like tit-for-tat with sophisticated estimation procedures that illustrate the popularity of these strategies among experimental subjects. They do not, however, explain why and when exactly players switch to different strategies when the cost/benefit ratio in a PD game environment changes. Using the Ties model we can fill this gap and predict different behavior for different b/c ratios. lt also allows us to test if the behavior of subjects is consistent between different specifications of the same game. While the previously mentioned studies do not attempt to explain why subjects use different strategies within the same game environment, our method does not attach a single strategy to an individual or even to an individual in a particular interaction. Another advantage is that applying the estimated Ties model allows us to see if the tie mechanism and resulting strategies are consistent across different, albeit related, games regarding social dilemmas.

An example of a strategy that is easily generated by the tie mechanism is the wellknown TFT strategy. According to this strategy, a player starts with cooperating (choosing C) and, subsequently, simply chooses whatever his opponent did in the previous round. Thus, after observing C (D) the player chooses C (D). If play starts with D first instead of C, the strategy is labeled DTFT. For the tie mechanism to generate such behavior, the following is required: First, a player should have a strong enough impulse parameter δ_2 (the exact size depends on the b/c ratio). Secondly, the memory of this player must not be too strong, as otherwise a strong tie can be built up that tolerates deviations by the other player. Hence, the tie-persistence parameter δ_1 must be sufficiently small. Finally, α_0 determines whether play starts with C (for TFT) or D (for DTFT).

If we allow players to start with $\alpha_0 \neq 0$ we find in appendix C that, for certain parameter values of δ_1 and δ_2, play starts to mimic often reported strategies like AD, TFT, and TF2T. However, a much more challenging task is to use the parameter value estimates of this paper and see to what strategies these parameter estimates correspond, a task to which we turn next. Besides the before-mentioned strategies we also investigate a modified strategy: 'qualified' tit-for-two-tats (QTF2T). This strategy is similar to TF2T in all but one respect:

choosing TFT in contrast 20% players choose TF2T when b/c was 4, only 5% of players chose TF2T when b/c was 1.5 as can be found in table 7. Dal Bó and Fréchette do not consider TF2T.

it requires more than one cooperative choice by the other before defection (D) is forgiven. In terms of the Ties model this means that first the value of α has to be significantly built up; in this case, described in appendix C.3, until the theoretical maximum ($\frac{\delta_2}{1-\delta_1}$), but any value corresponding to any number of consecutive cooperative decisions can be chosen. The intuition behind this 'qualification' is that players using TF2T are vulnerable to exploitation. Other players could exploit them by alternating between C and D. Since it seems unlikely that players would accept such exploitation we require the other player to show good intentions for a longer period, before these strategies become 'forgiving'. In Appendix C the case for $b/c = 4$ is worked out.

For a sensible comparison between the parameter estimates found in this study and those relevant for a PD environment we need to normalize the impulse, I. For, note that if we multiply the impulse by a factor i, the estimate of δ_2 will change with factor $1/i$. Therefore, we normalize by assuming a cooperative action in a PD game to be equivalent with a fully Pareto efficient action (C=4) in our FPG game and defining the impulse in that event to be equal to $I_n \equiv \frac{C_j - C_J^{ref}}{C^{eff} - C_J^{ref}}$. Thus, in order to translate the values we found for δ_2 to values suitable for a PD game environment, where choosing C (the efficient choice) is valued as 1, we multiply δ_2 by 4. As before α_0 will stand for the starting value of α.

Below there are two graphs, the first for $b/c = 2$ and the second for $b/c = 4$, which show which parameter values of the tie mechanism (δ_1 and δ_2) correspond to which strategies[11]. The lines mark the conditions for which the Ties model predicts the behavior of the strategies mentioned earlier. The crosses in the graph represent the different individuals in our experimental study, using normalized δ_{2P} (as only positive impulses are possible in this environment) and δ_1 values, estimated with the myopic model that allows for a dichotomy between positive and negative impulses.

[11] for characteristics of these strategies and a more elaborate analysis, see Appendix C.

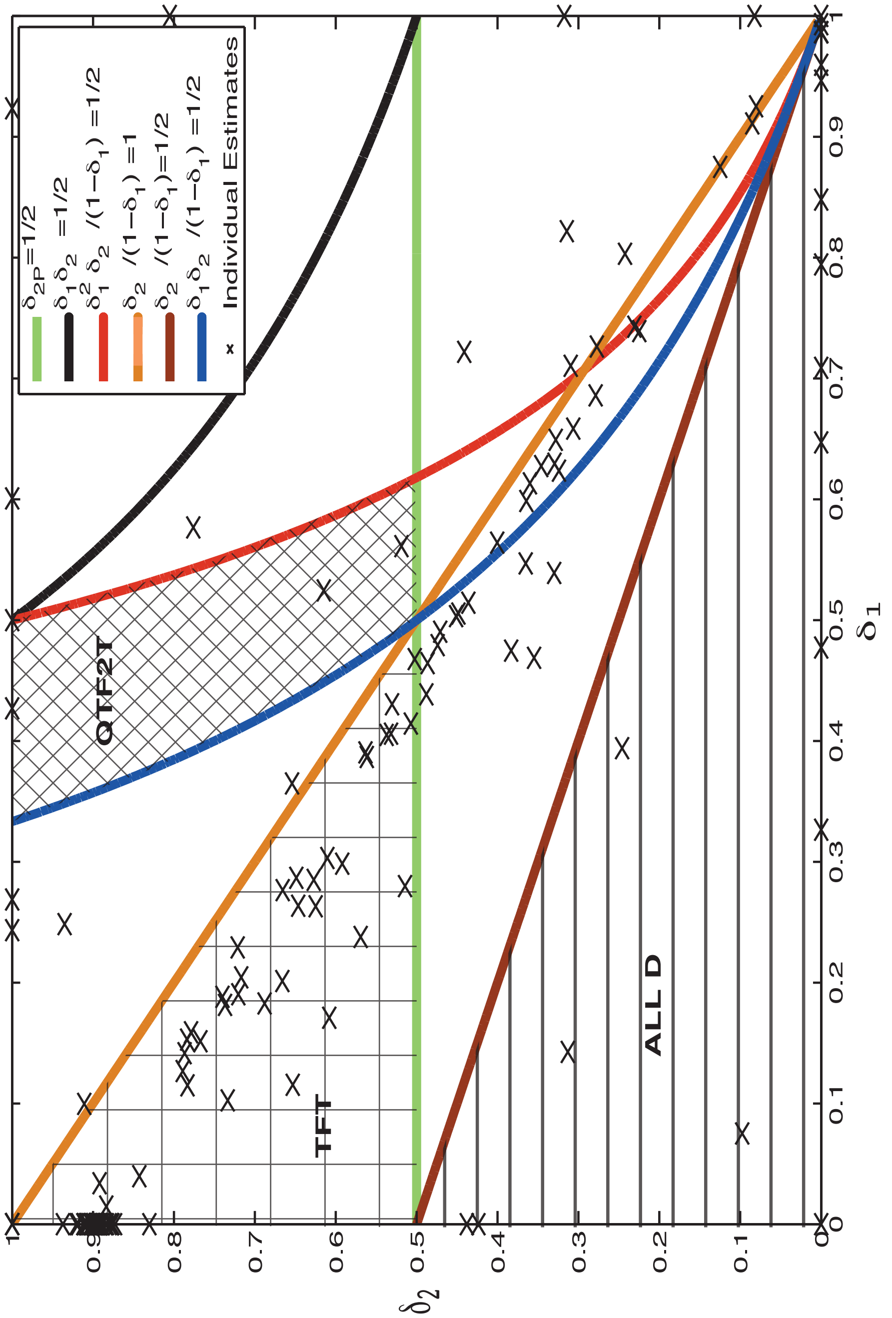

copy.pdf

Figure 3: Parameter estimates and strategies for b/c=2

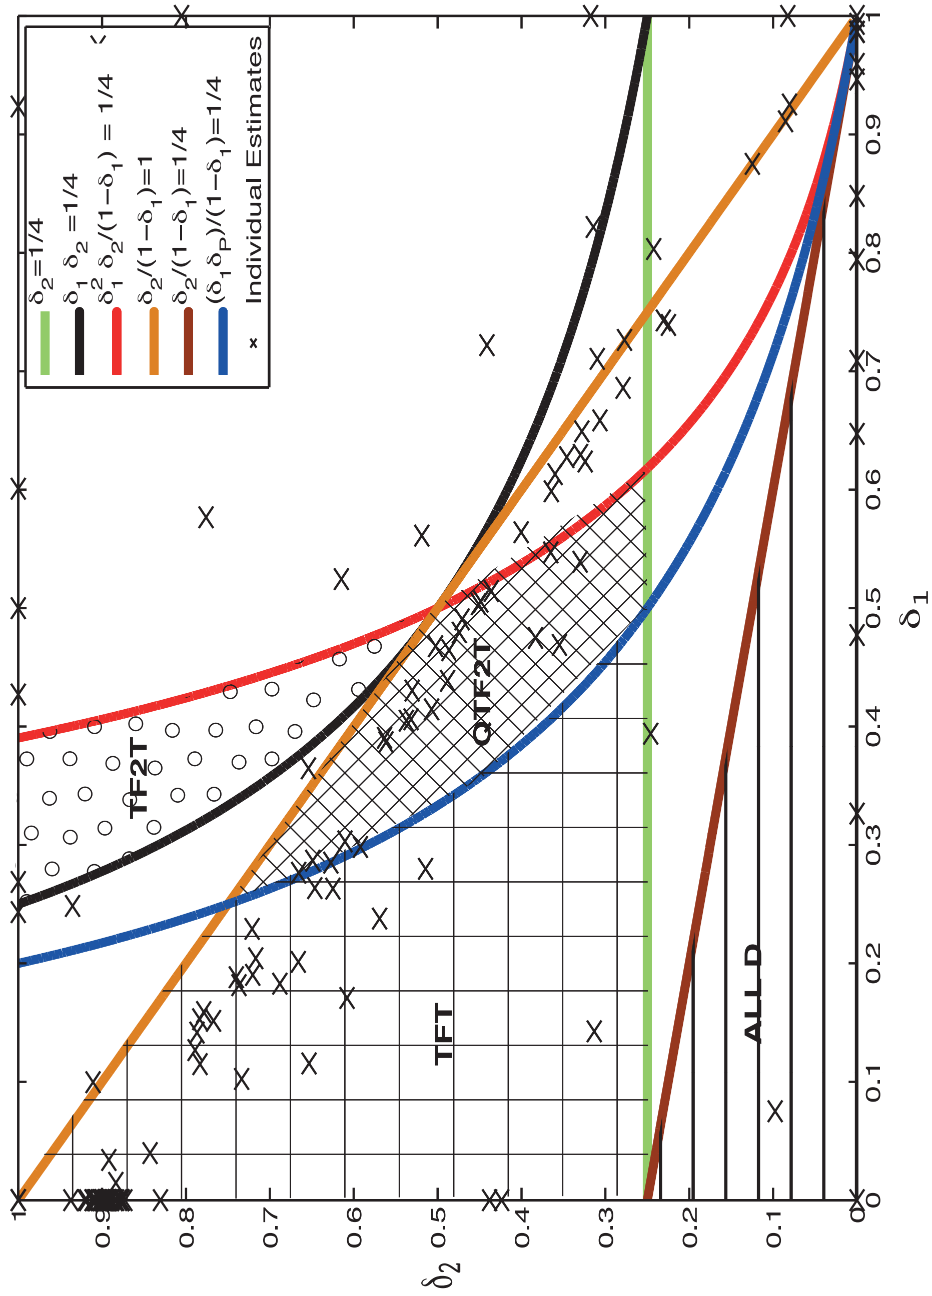

copy.pdf

Figure 4: Parameter estimates and strategies for b/c=4

From these figures it becomes clear that for a lower b/c ratio the same players 'switch' from cooperative strategies to strategies which imply more defection. It also shows that in the $b/c = 4$ setting, (D)TFT, TF2T and related strategies are commonly found, as in Fudenberg et al. (2012). There is however also a noticeable difference, the lack of players playing the AD strategy. A potential reason for this might be the possibility in our Fragile Public Good Game to destroy the public good, which could make players reluctant to not contribute, out of fear for punishment. This is highlighted by the fact that if we use the estimates from the myopic model that does not allow for a difference in impulse impact there are some more AD players, as δ_2 is typically estimated to be lower in this case. Observations that are to the right of the top-left bottom-right diagonal represent players with a very strong tie mechanism as continuous cooperation by the other player would lead to an α-value greater than one.

Finally, it is interesting to note that the Ties model could also explain the repeated PD finding of Breitmoser (2015) that if one player chooses C and the other D, both show an about equal probability of playing C in the next round. This is presented as evidence against the existence of TFT. If one thinks in terms of a tie mechanism this finding may not be so surprising. After all, if a player played C in the previous round his or her α value must have been relatively high, while if a player played D this value must have been relatively low. Now, because the tie (α) of the former player will decay (as it gets multiplied by δ_1) the chance that this player chooses C declines. In contrast, the tie of the other player will be reinforced (with δ_2) by counterpart's cooperative action in the previous round. Consequently, the α values will move towards each other. In short, as one tie is initially relatively strong (reflecting a higher probability of playing C) and becomes weaker, while the other tie is relatively weak (reflecting a higher probability of playing D) and becomes stronger, the chances to play C for both players converge, making the finding of Breitmoser explicable by a tie mechanism.

6 Conclusion

We conclude with a summary of our main findings. First of all it turns out that the estimation results of our Fragile Public Good Game are very much in line with earlier studies of public goods games using the ties model. More specifically, we also observe that the memory component of the tie mechanism, represented by the tie persistence parameter δ_1, is about equally important as the impulse component, represented by the parameter δ_2, with the scalefree δ_1 being estimated to be close to 0.50.

In contrast to our original hypothesis we find that positive impulses have a stronger impact than negative ones. Apperently, the bad is not stronger than the good in this context cf. (Baumeister et al., 2001). This asymmetry in the tie mechanism is helpfull in getting cooperation going, while not rendering cooperators defenseless against people that are just trying to benefit from them. We also find that players do not seem to be very much forward-looking.

Our out-of-sample predictions show that the Ties model significantly outperforms both a well-known reinforcement learning model as well as a model with constant social preferences. For both the learning model as well as the Ties model we find that the predictive power improves if we use group-level instead of individual-level estimates. This appears to be due to the lack of behavioral variability for some of our subjects.

The Ties model also generated insights for a (repeated) Prisoner's Dilemma (PD) game context. Strategies observed in experiments can be understood with the help of the Ties model. Moreover, using the estimated parameters from our public good game, the model helps explain why people switch to different strategies when faced with a different cost-benefit ratio in the PD game. Our alternative explanation for the behavior in repeated PD games does not require people to switch strategies in a seemingly ad hoc way.

References

Abbink, K. and Sadrieh, A. (2009). The pleasure of being nasty. *Economics Letters*, 105(3):306–308.

Andreoni, J. and Miller, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, pages 737–753.

Bault, N., Fahrenfort, J. J., Pelloux, B., Ridderinkhof, K. R., and Van Winden, F. (2016). An affective social tie mechanism: Theory, evidence and implications. *Unpublished manuscript. from http://www1. feb. uva. nl/creed/pdffiles/Second_Behavioral_Paper_15-11-13. pdf.*

Bault, N., Pelloux, B., Fahrenfort, J. J., Ridderinkhof, K. R., and van Winden, F. (2015). Neural dynamics of social tie formation in economic decision-making. *Social cognitive and affective neuroscience*, 10(6):877–884.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., and Vohs, K. D. (2001). Bad is stronger than good. *Review of general psychology*, 5(4):323.

Baumeister, R. F. and Leary, M. R. (1995). The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological bulletin*, 117(3):497.

Bolle, F., Tan, J., and Zizzo, D. (2013). Vendettas. *American Economic Journal: Microeconomics (forthcoming).*

Bone, J., Hey, J., and Suckling, J. (2004). A simple risk-sharing experiment. *Journal of Risk and Uncertainty*, 28(1):23–38.

Bone, J. D., Hey, J. D., and Suckling, J. R. (2003). Do people plan ahead? *Applied Economics Letters*, 10(5):277–280.

Bosman, R. and van Winden, F. (2002). Emotional hazard in a power-to-take experiment. *The Economic Journal*, 112(476):147–169.

Breitmoser, Y. (2015). Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma. *American Economic Review*, 105(9):2882–2910.

Camerer, C. and Ho, T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.

Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics*, 14(1):47–83.

Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review*, 101(1):411–429.

Erev, I. and Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881.

Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and economic behavior*, 54(2):293–315.

Fehr, E. and Fischbacher, U. (2003). The nature of human altruism. *Nature*, 425(6960):785–791.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics Letters*, 71(3):397–404.

Fudenberg, D., Rand, D. G., and Dreber, A. (2012). Slow to anger and fast to forgive: cooperation in an uncertain world. *The American Economic Review*, 102(2):720–749.

Hoyer, M., Bault, N., Loerakker, B., and Van Winden, F. (2014). Destructive behavior in a fragile public good game. *Economics Letters*, 123(3):295–299.

Kagel, J. H., Roth, A. E., and Hey, J. D. (1995). *The handbook of experimental economics.* Princeton university press Princeton, NJ.

King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., and Montague, P. R. (2005). Getting to know you: reputation and trust in a two-person economic exchange. *Science*, 308(5718):78–83.

Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, 9:2029–2061.

Laury, S. K. and Holt, C. A. (2008). Voluntary provision of public goods: experimental results with interior nash equilibria. *Handbook of experimental economics results*, 1:792–801.

Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of economic dynamics*, 1(3):593–622.

Liebrand, W. B. and McClintock, C. G. (1988). The ring measure of social values: A computerized procedure for assessing individual differences in information processing and social value orientation. *European journal of personality*, 2(3):217–230.

Plott, C. R. and Smith, V. L. (2008). *Handbook of experimental economics results*, volume 1. Elsevier.

Rand, D. G., Dreber, A., Ellingsen, T., Fudenberg, D., and Nowak, M. A. (2009). Positive interactions promote public cooperation. *Science*, 325(5945):1272–1275.

Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8(1):164–212.

van Dijk, F., Sonnemans, J., and van Winden, F. (2002). Social ties in a public good experiment. *Journal of Public Economics*, 85(2):275–299.

van Dijk, F. and van Winden, F. (1997). Dynamics of social ties and local public good provision. *Journal of Public Economics*, 64(3):323–341.

Yang, Y., Onderstal, S., and Schram, A. (2012). Inequity aversion revisited.

Appendices

A Proof of Propositions

A.1 Proposition 1: Contributions outside of $-5 \leq C_{it} \leq 5$ can never be part of any equilibrium if $-1 \leq \alpha \leq 1$.

We use the fact that our agents can only change their decision in discrete steps. Subtracting $V(C_{it})$ from $V(C_{it}+1)$ we get:

$$V(C_{it}) - V(C_{it}+1) = 2C_{it} + 2 - 10\alpha + \gamma(-10 + \alpha(2\gamma C_{it} + 2(1-\gamma)C^e_{jt} + \gamma + 1)) \quad (12)$$

Where γ is between 0 and 1. For the proposition to be true this equation must be positive. We reformulate this condition to:

$$2C_{it} + 2 + 2\alpha\gamma\big(\gamma C_{it} + (1-\gamma)C^e_{jt} + \frac{1}{2}\gamma + \frac{1}{2}\big) > 10(\alpha + \gamma) \quad (13)$$

We begin by only looking at equilibria with symmetric contributions ($C_{it} = C_{jt}$).

$$2C_{it} + 2 + 2\alpha\gamma(C_{it} + \frac{1}{2}\gamma + \frac{1}{2}) > 10(\alpha + \gamma) \tag{14}$$

So at $C_{it} = 5$ we have:

$$\begin{gathered} 12 + 2\alpha\gamma(5 + \frac{1}{2}\gamma + \frac{1}{2}) > 10(\alpha + \gamma) => \\ 12 + \alpha\big(\gamma(11 + \gamma) - 10) - 10\gamma > 0 \end{gathered} \tag{15}$$

This last statement is always true for -1<α<1, since if α is one we have:

$$12 + 11\gamma + \gamma^2 > 10(\gamma + 1) \tag{16}$$

Which is always the case. If α is -1 we have:

$$12 - 11\gamma - \gamma^2 > 10(\gamma - 1) \tag{17}$$

Now since (15) is a monotone function in α these results hold for the entire interval.

We can use the same method to show that $V(C_{it} > V(C_{it} - 1)$ always holds when $C_{it} \leq -5$. At $C_{it} = -5$ the equivalent of (15) is:

$$-12 + 2\alpha\gamma(-6\,\frac{1}{2} + \frac{1}{2}\gamma) > 10(\alpha + \gamma) \tag{18}$$

This is never true for positive α's. For α is -1 we get:

$$-12 + \gamma(13 - \gamma) > 10(\gamma - 1) \tag{19}$$

This cannot be for true for γ between zero and one either.

For the asymmetric equilibria we have to go back to (13). The left side is increasing in C^e_{jt} for $\alpha > 0$ and decreasing when $\alpha < 0$. To see if there are instances where contributing 6 is preferred to contributing less we therefore only need to check for $C^e_{jt} = 1$. This gives:

$$12 + 2\alpha\gamma(5\gamma + (1 - \gamma) + \frac{1}{2}\gamma + \frac{1}{2}) > 10(\alpha + \gamma) \tag{20}$$

Again this is always true for $0 \leq \alpha \leq 1$ (and if the other contributes positively, α cannot be negative in an equilibrium). A similar procedure can be used to show that no choice more negative than -5 can be part of an equilibrium.

A.2 Propositon 2.2: All symmetric equilibria with $-5 \leq C_{it} \leq 5$ are possible.

For a stable situation we need a value for α such that $V(C_{it} - 1) < V(C_{it}) > V(C_{it} + 1)$ holds and we need that after (infinitely) repeated play of C_{it} this still holds. We start by investigating the case in which here agents are not forward looking (γ is zero).

A.2.1 Myopic Agents

We first look at the case in which both contributions are positive. In this situation equation (14) simplifies to:

$$2C_{it} + 2 > 10\alpha \tag{21}$$

From (21) we with every increase of α by 0.2 the contribution that gives the highest value shifts one up. For an equilibrium to be sustainable we need the α-value to be stable (in a steady state) for a the given contribution. So we use (5), and look for:

$$\alpha = \delta_1 \alpha + \delta_2 I \tag{22}$$

For I we use the Nash equilibrium as a reference point as we did throughout the paper. This leads to:

$$\begin{gathered} \alpha = \delta_1 \alpha + \delta_2 C \\ \text{or} \\ \alpha = \frac{\delta_2 C}{1 - \delta_1} \end{gathered} \tag{23}$$

From (21) we know that:

$$0.2C < \alpha < 0.2(C + 1) \tag{24}$$

Combining (23) and (24) we obtain:

$$0.2 < \frac{\delta_2}{1-\delta_1} < 0.2 + (0.2/C) \tag{25}$$

If we look at the same situation (γ=0) for negative values (more precise for $C \leq$-2, we will discuss the situations in which C is 0 or -1 later) we change (21) into:

$$2C > 10\alpha \tag{26}$$

So also in the equation above we see that the best response changes with every increase (or drop) in α of 0.2. Following an analogous procedure to the one we used for a positive C we obtain the following condition:

$$0.2 > \frac{\delta_2}{1-\delta_1} > 0.2 - \frac{0.2}{C} \tag{27}$$

If C=0, then the stimulus is zero. This will lead to the value of α moving gradually towards zero as well. As the best response to an α-value of zero is to play 0 we have that the [0,0] equilibrium can always exist regardless of the δ-parameters. To the entire range of α-values wherefor a contribution of zero is a vest response we use (26) and observe that as long as α>-0.2 the value of playing zero is bigger then the value of -1. This gives us the lower bound α=-0.2. Now for the higher bound we have to see when playing 1 is more attractive then playing 0. From (21) we find that this boundary is 0.2.

A.2.2 Forward Looking Agents

If γ is unequal to zero all values of C_{it} are still part of symmetric equilibria, but the condition on δ_1 and δ_2 becomes stricter. Just as in the previous case we start from (14) and fill in (23):

$$\begin{aligned} 2C + 2 + 2\frac{\delta_2 C}{1-\delta_1}\gamma(C + \frac{1}{2}\gamma + \frac{1}{2}) &> 10(\frac{\delta_2 C}{1-\delta_1} + \gamma) \\ 2C + 2 &> (10 - 2\gamma(C + \frac{1}{2}\gamma + \frac{1}{2}))\frac{\delta_2 C}{1-\delta_1} + 10\gamma \end{aligned} \tag{28}$$

$$\frac{2C+1-10\gamma}{(10-2\gamma(C+\frac{1}{2}\gamma+\frac{1}{2}))C} > \frac{\delta_2}{1-\delta_1} \tag{29}$$

If $10-2\gamma(C+\frac{1}{2}\gamma+\frac{1}{2} < 0$ the inequality changes direction.

We also fill in the lower bound we obtain:

$$\begin{aligned} \frac{2C-10\gamma}{(10-2\gamma(C-1+\frac{1}{2}\gamma+\frac{1}{2}))C} &> \frac{\delta_2}{1-\delta_1} \\ \frac{2C-10\gamma}{(10-2\gamma(C-1+\frac{1}{2}\gamma+1/2))C} < \frac{\delta_2}{1-\delta_1} &< \frac{2C+2-10\gamma}{(10-2\gamma(C+\frac{1}{2}\gamma+\frac{1}{2}))C} \end{aligned} \tag{30}$$

We can repeat this procedure in the negative domain and get the following condition:

$$\frac{2C-2-10\gamma}{(10-2\gamma(C-1+\frac{1}{2}\gamma-\frac{1}{2}))C} > \frac{\delta_2}{1-\delta_1} > \frac{2C-10\gamma}{(10-2\gamma(C+\frac{1}{2}\gamma-\frac{1}{2}))C} \tag{31}$$

A.3 Proposition 3: Asymmetric equilibria exist if $C_iC_j > 0$ and $|C_i| \leq 5$ and $|C_j| \leq 5$

For simplicity we restrict ourselves to myopic agents. This changes (23) into:

$$\alpha = \frac{\delta_2 C_j}{1-\delta_1} \tag{32}$$

And (24) changes into:

$$0.2C_i < \alpha < 0.2(C_i+1) \tag{33}$$

Leading to:

$$0.2\frac{C_i}{C_j} < \frac{\delta_{2i}}{1-\delta_{1i}} < 0.2\frac{C_i}{C_j} + \frac{0.2}{C_j} \tag{34}$$

In order for this situation to be an equilibrium we also need:

$$0.2\frac{C_j}{C_i} < \frac{\delta_{2j}}{1-\delta_{1j}} < 0.2\frac{C_j}{C_i} + \frac{0.2}{C_i} \tag{35}$$

Looking at the extreme case of a [1,5] equilibrium this implies:

$$0.04 < \frac{\delta_{2i}}{1-\delta_{1i}} < 0.08 \tag{36}$$

And:

$$1 < \frac{\delta_{2j}}{1-\delta_{1j}} < 2 \tag{37}$$

While such an equilibrium is mathematically possible, for it to be maintained the two players have to be quite different.

There are no equilibria where one player contributes a negative amount while the other contributes a positive amount. Constant negative contributions eventually create a negative α in the other and contributing positively can not be an optimal choice under a negative α.

A.4 Proposition 4: For the socially optimal choices to be a stable equilibrium under the myopic model, both players satisfying $0.2 < \frac{\delta_{2i}}{1-\delta_{1i}} < 0.25$ is a necessary, but not a sufficient condition.

This result is a directly visible in (25) if we plug in 4 as the contribution level. (25) also shows that if a player has the characteristics to be in a socially optimal equilibrium he or she is also willing to conform with any other (non negative) symmetric equilibrium with lower contributions.

B Reference Point

In this part we will evaluate the model fit for different reference points in our model. In this section we restrict ourselves to the myopic version of the model, allowing for different positive and negative impulse parameters (as this was the best predicting model). We have for the size of the impuls:

$$I_{ijt} = C_{jt} - C_i^{ref} \tag{38}$$

The definition of the reference contribution C^{ref} is not trivial. Several points are however appealing from a theoretical standpoint. The first candidate that we consider is the point that we chose in the main paper, the contribution that an agent chooses in a one-shot Nash equilibrium ($C^{ref} = 0$). Another static option is the use of the Pareto optimal contribution as a reference point ($C^{ref} = 4$). It is also possible that the reference point is not static and depends on either an agent's own behavior or the previous behavior of the other. We test the predictive performance for two such reference points: $C_t^{ref} = C_{it}$ if an agent's own contribution is used and ($C_t^{ref} = C_{jt-1}$) if one looks at the contribution of the other in the previous round. In the last two cases we initialize the system using $C_1^{ref} = 0$. In the table below the parameters from estimating the model using different reference points are shown.

Table 6: Estimates for different reference points

Reference point	δ_1	δ_{2P}	δ_{2N}	$\sum LL$
Nash	0.490	0.115	0.080	-8545
Pareto	0.975	0.108	0.000	-11265
Own Contribution	1	0.003	-0.006	-11598
Contribution other	0.989	0.179	0.160	-10826

From the table we see that using the Nash solution as a reference point produces to the highest likelihood. It is also interesting to note that with dynamic reference points δ_1 is estimated to be very close to 1. A reason for this might be that if two players are in a positive symmetric equilibrium (where C_{it}=C_{jt} for multiple rounds) the value of α goes to zero (or might even go negative in case of the pareto optimum being the reference point) as all the impulses are zero. We observe these equlibria quite regularly in our dataset. The only way for the models with these particular reference point specifications to keep α high, which is necessary for positive contributions to occur, is for δ_1 to approach zero. A side effect of this is that with $\delta_1 \approx 1$ players basically have an infinite memory and early impulses have the same effect as new ones. Judging on the basis of the likelihood, though, this is not the case.

Also the Pareto optimum does not perform well. This can be due to the fact that using this reference point even positive contributions might lead to a negative α and thus to the strange situation that if $\delta_{2N} > 0$ positive contributions by one player would lead to (expected) negative contributions by the other. This would lead to a negative spiral, that we hardly ever observe in the data.

We thus conclude that, if we use (6a), (6b) and (38) to model the tie mechanism, then $C^{ref} = 0$ is the best rule to use for the reference point.

It is interesting to see that if we focus on $\Delta\alpha$=(α_t−α_{t-1}) in a positive and symmetric equilibrium, we get into a situation where the change in α is basically determined by the change in contributions since the initial value of α diminishes over time. If we start from the basic tie mechanism described in (5) we have (with $I = C_{jt}$):

$$\begin{aligned}\alpha_{ij2} = \alpha_{ij1}\delta_{1i} + \delta_{2i}C_{j1} \quad => \quad \alpha_{ij3} = \alpha_{ij1}\delta_{1i}^{2} + \delta_{1i}\delta_{2i}C_{j1} + \delta_{2i}C_{j2} \quad => \\ \alpha_{ijt} = \alpha_{ij1}\delta_{1i}^{t-1} + \delta_{1i}^{t-2}\delta_{2i}C_{j1} + ... + \delta_{1i}\delta_{2i}C_{jt-2} + \delta_{2i}C_{jt-1} \quad => \\ \alpha_{ijt} \approx \frac{\delta_{2i}C}{1-\delta_{1i}} \text{ (for t -> } \infty)\end{aligned} \tag{39}$$

Consequentle, for $\Delta\alpha$ it holds in this situation:

$$\begin{aligned}\Delta\alpha_{it+1} = \alpha_{it+1} - \alpha_{it} \quad = \quad \delta_{1i}\alpha_{it} + \delta_{2i}C_{jt} - \alpha_{it} \quad => \\ \Delta\alpha_{it+1} = \delta_{2i}C_{jt} - (1-\delta_{1i})\alpha_{it} \approx \delta_{2i}C_{jt} - (1-\delta_{1i})\frac{\delta_{2i}C^{e}q}{1-\delta_{1i}} = \delta_{2i}(C_{jt} - C^{eq})\end{aligned} \tag{40}$$

In the symmetric equilibria (the most commonly observed equilibria in our experiment) we have that, when the reference point with respect to α is fixed, the change in α approximates a linear function of the change in the other player's contribution.

C Tie Model Parameters for different Repeated PD Strategies

C.1 Introduction

In this section we look at a number of strategies for the repeated prisoner's dilemma game and determine which combinations of parameters in the tie model are compatible with those strategies. This analysis provides the basis for the hypothetical distribution of strategies presented in section 5, which we derived based on our parameter estimates from the fragile public good game. We take a prisoner's dilemma game of the following form, which corresponds to Fudenberg et al. (2012)'s b/c=4 case (only the required α-values will change in the analysis below if one uses other b/c ratios):

Table 7: Prisoner's Dilemma with b/c=4

	C	D
C	3,3	-1, 4
D	4,-1	0,0

We start from the basic Ties model in which agents have the following extended utility function:

$$V_{it} = U_{it} + \alpha_{ijt} U_{jt} \tag{41}$$

Where V_{it} denotes the extended utility of player i at time t and U_{it} and U_{jt} stand for the direct own utility (payoff) of i and j, respectively, at time t, while α_{ijt} represents the weight i attaches to the utility of j (i's tie with j) in period t, which is updated as follows:

$$\alpha_{ijt} = \delta_{1i}\alpha_{ijt-1} + \delta_{2i} I_{t-1} \tag{42}$$

With α_{ij1} denoting the initial tie. We define the impulse I_{t-1} to be the scaled amount by which the other deviated in the previous period from the standard (one-shot) Nash equilibrium choice. If the other player cooperated in the previous period the impulse equals 1, if the other defected it equals 0. In this model the choice between cooperating and defecting is fully dependent on the level of α. If α is larger than 1/4 cooperating is a dominant choice while for α smaller than 1/4 defecting is dominant, as in the standard models.

Both Dal Bó and Fréchette (2011) and Fudenberg et al. (2012) find experimental evidence that many subjects in their experiments use either tit-for-tat (TFT) or tit-for-two-tats (TF2T) as a strategy. In this exercise we will see if the simple and neurological underpinned Ties model (Bault et al., 2015, 2016) could explain the behavior described by these strategies.

C.2 Tit-for-tat

The TFT strategy is simple: A player begins by playing C and simply imitates the other player's action in future periods. We can derive conditions on the ranges of parameters δ_1, δ_2, and α_{ij1} for which this behavior is sustained. First, we investigate what levels of α can be reached after continuous play of either C or D. Since $\delta_1 < 1$, being exposed to infinitely repeated defection leads to α going to zero. Now, what about continous play of C? Noting that the initial value of α vanishes as t goes to infinity, we start from the expression for α in period 2 and iterate:

$$\begin{gathered}\alpha_{ij2} = \delta_{1i}\alpha_{ij1} + \delta_{2i}I_1 \quad => \quad \alpha_{ij3} = \delta_{1i}^2\alpha_{ij1} + \delta_{1i}\delta_{2i}I_1 + \delta_{2i}I_2 \ => \\ \alpha_{ijt} = \delta_{1i}^{t-1}\alpha_{ij1} + \delta_{1i}^{t-2}\delta_{2i}I_1 + ... + \delta_{1i}\delta_{2i}I_{t-2} + \delta_{2i}I_{t-1}\end{gathered} \tag{43}$$

Furthermore, if C is played, I is always equal to one, so that:

$$\begin{gathered}\alpha_{ijt} = \delta_{1i}^{t-1}\alpha_{ij1} + \delta_{1i}^{t-2}\delta_{2i} + ... + \delta_{1i}\delta_{2i} + \delta_{2i} :=> \\ \alpha_{ijt} \approx \frac{\delta_{2i}}{1-\delta_{1i}} \ \text{ (as t -> } \infty\text{)}\end{gathered} \tag{44}$$

This is an important result as it gives an upper limit to the alpha level that can be reached in an infinitely repeated game. Now that we have established both the upper and the lower limit for α ,we can check the conditions for which the behavior according to the ties model is identical to the TFT strategy. The first condition is simply that the first action must be to play C. This leads to the simple condition (omitting the i and j subscripts, for convenience):

$$\alpha_1 \geq 1/4 \tag{45}$$

We also need that, no matter how high the current level of α is, after only one period of D played by the other, a TFT player weakly prefers D over C. Using eqs. (42) and (45), it follows that the tie value of a player who experienced infinitely repeated C and one period of D is equal to $\frac{\delta_1\delta_2}{1-\delta_1}$. Since D can only be weakly preferred if $\alpha \leq 1/4$, we get the condition that:

$$\frac{\delta_1\delta_2}{1-\delta_1} \leq 1/4 \tag{46}$$

On the other hand, we also need that, no matter how low α is, after only one period of C played by the counterpart C is weakly prefered over D, which requires (as $\alpha \geq 0$):

$$\delta_2 \geq 1/4 \tag{47}$$

By combining (46) and (47) we find that $\delta_1 \leq 1/2$ should hold. The intuitive explanation for these values is that a player has to be sufficiently sensitive to a changed impulse and must not have too strong of a 'memory'.

C.3 Tit-for-2-tats

After having defined the Tie-model parameters for which TFT is the resulting strategy, we now repeat the same exercise for the TF2T behavior. As this strategy also starts with playing C, we need (45) to hold. Furthermore, even after continuous D play, after only one period of C by the counterpart, C should be weakly dominating D, so (47) should hold as well.

Before we proceed with imposing further restrictions, we have to decide how strict we want to be in our interpretation of the TF2T strategy. If we take it at face value we have to assume that even after a history like $DDDDDDCD$ a player will still be patient and play C. It also implies that this player would constantly play C against a counterpart that keeps alternating between C and D. In order to account for such phenomena we evaluate two different versions of TF2T, one that takes the strategy literally and one that requires multiple periods of C before the trust in the other is restored. For convenience, the latter version will be called "qualified tit-for-two-tats" (QTF2T). For the standard TF2T we need that, even if we start out with $\alpha = 0$ and the other player cooperates in one round, only to defect immediately thereafter, a player would still reply with C to that D choice. This requires that:

$$\delta_1\delta_2 \geq 1/4 \tag{48}$$

In addition, we need that, even at the highest possible level of α, after two periods of D a player wants to choose D, which requires (using eqs. (45)):

$$\frac{\delta_1^2\delta_2}{1-\delta_1} \leq 1/4 \tag{49}$$

Combining (48) and (49) gives:

$$\frac{\delta_1}{1-\delta_1} \leq 1 \quad \text{or} \quad \delta_1 \leq 1/2 \tag{50}$$

So we have:

$$\delta_2 \geq 1/2 \tag{51}$$

A potential problem for the result above is that the upper bound of α, $\frac{\delta_{2i}}{1-\delta_{1i}}$, will be larger or equal to 1 (with equality only if δ_1=δ_2=1/2), since if we combine (48) and (44) :

$$\begin{aligned} \frac{\delta_2}{1-\delta_1} = \frac{\delta_2\delta_1}{\delta_1(1-\delta_1)} \geq \frac{1}{4(\delta_1(1-\delta_1))} => \\ \frac{\delta_2}{1-\delta_1} \geq 1 \end{aligned} \tag{52}$$

This 'problem' can be solved by using a QTF2T strategy where we assume here (for simplicity) that the value of α must be maximized for a player to play C after the other player chooses D. In this case, it is required that:

$$\frac{\delta_1\delta_2}{1-\delta_1} \geq 1/4 \tag{53}$$

D Payoff Matrix and SVO Example

Your possible choices		Other's possible choices: Take 7	Take 6	Take 5	Take 4	Take 3	Take 2	Take 1	0	Contribute 1	Contribute 2	Contribute 3	Contribute 4	Contribute 5	Contribute 6	Contribute 7
Take	7	84	94	104	114	124	134	144	154	164	174	184	194	204	214	224
Take	6	98	108	118	128	138	148	158	168	178	188	198	208	218	228	238
Take	5	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
Take	4	120	130	140	150	160	170	180	190	200	210	220	230	240	250	260
Take	3	128	138	148	158	168	178	188	198	208	218	228	238	248	258	268
Take	2	134	144	154	164	174	184	194	204	214	224	234	244	254	264	274
Take	1	138	148	158	168	178	188	198	208	218	228	238	248	258	268	278
	0	140	150	160	170	180	190	200	210	220	230	240	250	260	270	280
Contribute	1	138	148	158	168	178	188	198	208	218	228	238	248	258	268	278
Contribute	2	134	144	154	164	174	184	194	204	214	224	234	244	254	264	274
Contribute	3	128	138	148	158	168	178	188	198	208	218	228	238	248	258	268
Contribute	4	120	130	140	150	160	170	180	190	200	210	220	230	240	250	260
Contribute	5	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
Contribute	6	98	108	118	128	138	148	158	168	178	188	198	208	218	228	238
Contribute	7	84	94	104	114	124	134	144	154	164	174	184	194	204	214	224

Figure 5: Payoff Matrix

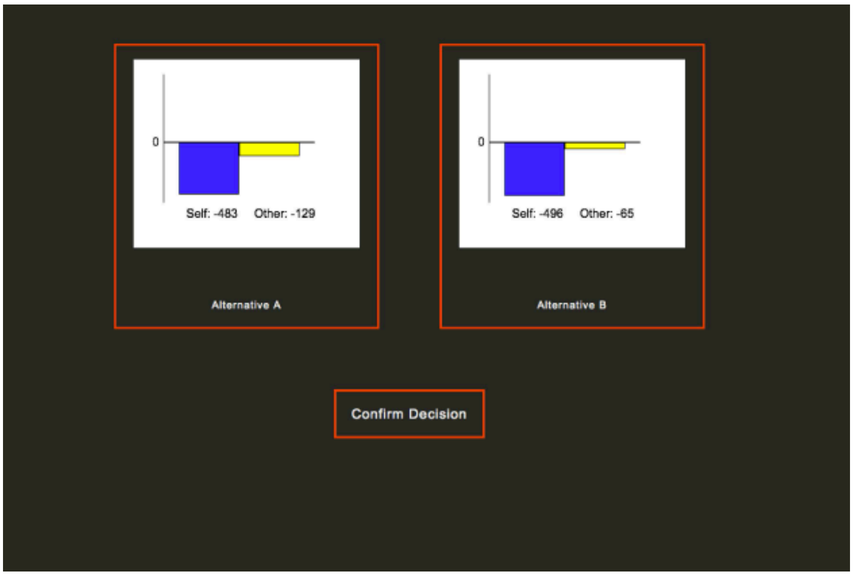

Figure 6: SVO example