On the Dynamic Development of Social Ties: Theory and Application*

Benjamin Pelloux[1†], Nadège Bault[1,2,3], Johannes J. Fahrenfort[3,4,5], K. Richard Ridderinkhof[3,5] and Frans van Winden[2,3]

Abstract

We implement, expand with limited foresight, and test the performance of the theoretical social ties model of van Dijk and van Winden (1997). The model is estimated on various experimental data sets of repeated public good games involving different numbers of players. Our estimation results provide direct support for the proposed social ties mechanism, showing that the history of social interaction is an essential determinant of preferences. About a quarter of the subjects appear to be forward-looking. Both the within-sample and the out-of-sample predictive performance of the model turn out to be remarkably good, as it is able to track the often complex dynamic contribution patterns. This is true both for games played in pairs and in groups of four players, where the decision-making processes are much more complex. Moreover, we show that this model performs better than some other social preferences models that allow for a dynamic implementation. Our conclusion is that this simple, tractable, and psychologically grounded model of social preferences can account very well for dynamic behavioral patterns in repeated public good games.

Key-words: social preferences, social ties, public goods, experimental economics.
JEL Codes: H41, C92

[1] Groupe d'Analyse et de Théorie Economique (GATE), CNRS and University of Lumiere Lyon 2, 69130 Ecully, France
[2] Center for Research in Experimental Economics and political Decision-making (CREED), Amsterdam School of Economics, University of Amsterdam, 1018WB Amsterdam, the Netherlands
[3] Cognitive Science Center Amsterdam (CSCA), University of Amsterdam, 1018WS Amsterdam, the Netherlands
[4] Brain and Cognition, Department of Psychology, University of Amsterdam, 1018WS Amsterdam, the Netherlands
[5]Amsterdam center for the study of adaptive control in brain and behavior (Acacia), Dept. of Psychology, University of Amsterdam, 1018XA Amsterdam, the Netherlands
[†] Corresponding author (pelloux@gate.cnrs.fr)

* Financial support by the Research Priority Program Brain & Cognition of the University of Amsterdam is gratefully acknowledged. At the time of writing the first draft Benjamin Pelloux was affiliated with CREED. We thank Jean-Louis Rullière, Ben Loerakker and Marcelo Tyszler for their comments, and are further grateful for the remarks made by participants at the 5th Singapore Economic Review Conference and at ABEE 2013.

1. Introduction

Economists started to formalize human economic behavior in a way that is now referred as the *homo œconomicus* theory, which attributes to agents, among other characteristics, purely selfish preferences and infinite computing abilities allowing them to perfectly maximize their utility given their beliefs about their environment. However, repeatedly observing behaviors that seem to conflict with the predictions of this model, many researchers in various fields built alternative theories in order to explain these 'anomalies'. Most of these theories retained the rationality assumption and focused on the development of new types of preferences. Gathered under the broad label of "social preferences", these models assume interdependent utility functions of various forms, meaning that one's utility is also depending on the utility (or at least payoff) of the individuals one is interacting with. Some of these models are based on purely distributional preferences meaning that the choice of an action is made simply based on the different distributions of income that are available to the decision maker, independent of how these options were made available to the agent (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Kirchsteiger, 1994; Levine, 1998). Other models are implementing preferences directly built on reciprocity, that is answering an (un)kind act by a costly (un)kind action (already suggested by Kreps et al., 1982; through the existence of tit-for-tat types of players). Sobel (2005) argues that preferences might contain what he calls "intrinsic reciprocity", a genuine taste for reciprocity where preferences depend both on outcomes as well as on strategies. For example, Rabin (1993) lets utility be context-dependent, thus being a function of both the distributional outcomes available and the strategies used to arrive at the decision node. Other examples of reciprocity models can be found in Cox et al. (2007), Falk and Fishbacher (2006) and Dufwenberg and Kirchsteiger (2004).

Even though these approaches clearly brought valuable insights, they also raised several problems. First of all, evidence appears to be mixed and results to depend a lot on the type of game or environment being scrutinized. Second and as far as reciprocity models are concerned, their complexity caused by their reliance on psychological game theory as well as the very often observed presence of multiple (and sometimes unreasonable) equilibria have been pointed by some researchers (for a survey, see Fehr and Schmidt, 2006). Another more general problem has been identified by Sobel (2005): "The most basic problem is the specification of α [the weight put on other's payoff]. Charness and Rabin (2002), Martin Dufwenberg and Georg Kirchsteiger (2004), Falk and Fishbacher (2005), and Rabin (1993) present explicit functional forms for α, all motivated by plausible intuitive arguments and appeals to selected experimental evidence, but no one has described observable behavioral assumptions on α that best describe behavior." (op. cit.; p409) As we will see later, this paper is trying to tackle this problem.

Moreover, models describing purely distributional preferences typically assume that these preferences are stable. In an environment of repeated games, no time or history dependence is included. As a consequence, most of the empirical tests of these models are performed either on static distributional choices (*i.e.* not including an interaction), on one-shot games or on the last period of a repeated game, in order to evacuate learning or reputation building issues. For example, Levine (1998) focused on the last period of play when confronting his model to data (among which are public good games), "after the players have had time to learn an equilibrium" (op. cit., p599). As acknowledged by Fehr and Schmidt (1999, p 851): "an obvious limitation of our model is that it cannot explain the evolution of play over time in the experiments discussed. Instead, our examination aims at the explanation of the stable behavioral patterns that emerge in these experiments after several periods." Aside from the immediate question about the existence of such "stable behavioral patterns" after several periods[1], the question of how do players reach a cooperative or a competitive equilibrium is interesting and far from trivial.

This paper extends an alternative model originally developed by van Dijk and van Winden (1997) and investigates its performance in explaining and predicting data. They proposed a simple and tractable formalization of the development of affective ties between players facing the problem of private provision of public goods, one of the economists' favorite frameworks for the analysis of social preferences. In this model, social ties are simply modeled as the weight attached to another individual's well-being in one's own utility function, a weight that depends on the dynamics of provision from the players. More precisely, interaction history modulates the development of social ties in an automatic way, through individual parameters reflecting affective processes happening at the autonomic level.[2] Keeping the utility maximization hypothesis, this model makes social preferences endogenous and specific to each interaction partner and allow for studying the dynamics of play. In a nutshell, our model assumes that agents will experience feelings based on their

[1] Even if this may depend on games and contexts, a quick glimpse at the individual patterns of contribution from the three public good experiments presented later (see section 3) makes clear that such stable patterns are far from being always reached.

[2] Economists already took into account the impact of emotions in individual decision making, especially when modeling decision-making under uncertainty (see for example models on regret (Loomes and Sugden, 1982) or loss aversion (Tversky and Kahneman, 1992). However, it has been absent of most behavioral models of social decision making until very recently (some of the exceptions are Battigalli and Dufwenberg, 2007; Loewenstein and O'Donoghue, 2007; van Dijk and van Winden, 1997; van Winden, 2001).

emotional appraisal of the quality of the interaction. These emotions will then influence their attitude and attachment towards the other, making these concerns for others intrinsically dynamic. All these components give to the resulting social preferences an “individualized, historical, and context dependent content” (van Winden et al., 2008; p126).

The social ties model already received indirect support through several behavioral and neurological experiments. Behaviorally, van Dijk et al. (2002) showed that their discrete measure of social ties was significantly influenced by the success of the interaction (measured by contribution levels or earnings) in a public good game played in pairs. Sonnemans et al. (2006) replicated this finding with four-player groups. They showed that a subject developed different social ties with different counterparts and that these differences depended on individual contribution behaviors. Brandts et al. (2009) gathered evidence from a prisoner’s dilemma game that emotions were mediating the effect of the interaction on social ties. At the neural level, Fahrenfort et al. (2012) showed that interaction success as well as post-experimental liking ratings correlate with activity in the posterior superior temporal sulcus (pSTS; an area associated with social significance of stimuli and more broadly intention inference) during an allocation task.

In this paper, we will present a direct behavioral evaluation of the social ties model. To that purpose, we will first implement and extend the theoretical model of van Dijk and van Winden (1997), to allow for (limited) forward-looking behavior and stochasticity such that it can embed several canonical models of preferences.[3] This renders a specification of α and an explicit dynamic model to study the evolution of play. Subsequently, we’ll estimate and investigate the performance of the model through within-sample and out-of-sample tests, using three different datasets concerning (two-player as well as four-player) public good games. We show that both long-term history of the interaction and immediate reciprocity have a significant impact on behavior and that parameters estimates are stable when comparing different experiments with the same number of players. Moreover, the model is able to predict very well individual dynamic behavior, even out of sample.

In section 2, we will present the formal dual-process model of the development of social ties. Section 3 will present the different tests of this theoretical model against experimental datasets using public good games and section 4 will add some extended analyses. Finally, section 5 will conclude and consider perspectives for future research.

[3] We do not tackle here the analysis of dynamic equilibria since the focus of this article is empirical rather than theoretical. For the interested reader, equilibrium analyses of the original model are provided in van Dijk and van Winden (1997).

2. A Theoretical Model of the Dynamic Development of Social Ties

A lot of evidence from psychology (for a survey, see Baumeister and Leary, 1995) showed that the search for interpersonal attachment is one of the most important elements of human motivation. Evolutionarily sensible (e.g. as a way to fight external threats or to enforce trusting behavior), the question on how people develop affective ties have receive little scrutiny, especially in terms of economic modeling. The first version of the formalization of dynamic social ties in the framework of public good contributions appeared in van Dijk and van Winden (1997). We present here a linearized version of this theoretical model and extend it in two directions by including a probabilistic choice mechanism and (limited) foresight. Indeed, indeterminacy of behavior is getting more support from researchers of different fields (for a survey see Glimcher, 2005) and we think that this is a valuable approach. Concerning the latter, this assumption rests on several experimental results. It has been shown that most individuals were not using (full) backward induction (Johnson et al., 2002), as it is assumed by traditional game theory. Without assuming that agents do not plan at all (results from Bone et al. (2009) show that this is the case for more than half of their subjects) it seems as improbable that all individuals are able to reason backward from the last stage of the game. Cognitive hierarchy models (like Camerer et al., 2004) are examples of models assuming limited cognitive abilities for agents.

In this presentation of the model, we are implicitly considering the framework of analysis to be the private provision of public goods but are trying to stay as general as possible. We assume that the game is played repeatedly for *T* periods in groups of *N* players. The utility of an agent in one period is defined as follows[4]:

$$V_{it} = U_{it} + \lambda_i . U_{it+1} \qquad (1)$$

where V_{it} is the agent *i* expected utility in round *t*, which depends both on U_{it} and U_{it+1} , the utility he expects to get from period *t* and *t+1*, respectively. The parameter $\lambda_i \in [0; 1]$ is the discount factor applied to the utility expected from next period by agent *i*. Thus, agents have limited foresight since they will maximize their utility only over the two coming periods of play. Notice that an agent with $\lambda_i = 0$ is myopic since he does not consider the utility he will get from the next period. This limited foresight approach was already suggested in the past (Isaac et al., 1994) but to our knowledge its formalization in a model is novel.

[4] A period refers here to a repetition of a simultaneous game, as in the public good framework. It can also be thought, of course, as a decision node in a sequential game.

Before going further, we need to specify what we mean by social tie. We use here the definition given by van Winden et al. (2008; p 128): "a social tie refers to a caring about the interests of a specific other person, based on feelings experienced while interacting with that other person." First, the fact that a tie is individual-specific means that an agent will develop an individualized relationship with every other person he interacts with. Thus, he may develop different attitudes towards different agents depending on how the interaction unfolds. Second, we restrict our definition of social ties only to the environments where an interaction takes place. This differentiates this concept from others like sympathy and empathy that do not require the actual occurrence of an interaction.[5] These features are represented formally in the following linear definition of the expected utility for each period:

$$U_{it} = P_{it} + \sum_{j=1}^{N} \alpha_{ijt} . P_{jt} \tag{2}$$

where P_{it} is agent *i* payoff from the game in period *t* and α_{ijt} is the social tie that agent *i* has developed with agent *j* in period *t*.[6] We model this social tie simply as the weight in one's utility to the payoff received by the considered counterpart. This weight can be positive, null or negative. One might consider the limitation of its absolute value to 1 (which is the weight of one's own payoff) a reasonable assumption: occurrence of individuals valuing more the well-being of another person than their own may arise only in very specific situations (e.g. heroism or kin relationships. However we do not put any limitation to it in our model, leaving this door open.

As we said, the "kindness" or "spitefulness" of an agent in our model is not a fixed trait as the weight he puts on the well-being of the others will fluctuate through the interaction. In this way, we depart from the traditional assumption of stability of preferences and shift this stability to a deeper level by anchoring our parameters to psychological processes. As shown in the following equation, the parameters of our model are thus reflecting history retention (the number of periods of play influencing the social tie) and

[5] Whereas one may feel empathy just by observing or thinking about another person's situation, he would have to interact in some way with this person in order to develop a specific affective attachment. Following this definition, the social tie one develops with another person can only arise if an interaction takes place, i.e. if actions are perceived and if the other person's behavior is identifiable.

[6] It should be noticed that this individualization of the social tie is possible only if the necessary information about the behavior of the counterparts is provided to the subjects. Otherwise, they will simply develop a unique attitude towards the group, without being able to discriminate between players.

emotionality (the extent to which an emotional impulse will influence the social tie), traits that we believe to be more basic than social preferences[7]:

$$\alpha_{ijt} = \delta_i^1.\alpha_{ijt-1} + \delta_i^2.I_{ijt-1} \tag{3}$$

where δ_i^1 is the "tie persistence" parameter (related to memory), δ_i^2 the "tie proneness" parameter (related to emotionality), and $I_{it-1}$the emotional impulse for player *i* caused by player *j* in period *t-1*. The history of play is introduced through $\delta_i^1 > 0$, a memory related parameter, which captures the persistence of the previous tie and determines the number of periods of play that will influence behavior significantly. The parameter $\delta_i^2 > 0$, related to emotionality, captures the emotional impact of the impulse generated by counterparts' behavior. This impulse, I_{ijt-1}, is defined as the last action taken by the other player compared to some reference point. In our public good game setting, the impulse function is simply:

$$I_{ijt} = g_{jt} - g_{jt}^{ref} \tag{4}$$

where g_{jt} is the contribution made by the player *j* in round *t* and g_{jt}^{ref} is the reference level about player *j*'s contribution for round *t*. At this stage, we do not constraint the type of reference and many options are worth considering, such as a fixed norm, the standard Nash equilibrium of the static game, the expectation about other's behavior, or one's own behavior. It can thus be exogenous as well as endogenous. We want to concentrate here on the first emotional appraisal of a situation, contrary to social preferences models based on beliefs about the intentions or the type of the opponents, where more emphasis is put on the cognitive side of decision making[8] (e.g. Levine, 1998). This view was stimulated by the works of psychologists and neuroscientists like Zajonc (1984) and Ledoux (1997), arguing in favor of the primacy of affect. It is also in line with work being in done in psychology on the difference between description versus experience based choices (for a survey, see Rakow and Newell, 2010).[9] In our model, while preferences will be shaped emotionally and out of

[7] Even though the nature of certain social preferences appears to be, at least to some extent, innate. See for example Matthews et al. (1981) or Rushton (1991).

[8] However, we could imagine that perceived intentions are entering the model through their influence on the reference point that agents use to gauge the kindness of the other. See below.

[9] Rakow and Newell (2010) argue that experience-based and description-based choices may trigger very different processes in terms of information treatment (acquisition, representation, weighting and integration). It has also been shown that people's choice are sometimes at odds with their judgments and representations (Barron and Yechiam, 2009). More directly linked to our social decision-making context, Denrell (2005) showed that the formation of social impressions (and thus the desire to continue or not an interaction) was greatly impacted by the sequential sampling of experiences. All these results point towards the importance, in terms of preferences and decision making, of both emotional responses generated by experience and dynamic paths.

cognitive control by what is experienced, cognition will then kick in to make the best possible decision according to these preferences.

Notice also that we need to model the way expectations concerning the next period ($t+1$) contribution(s) of the other player(s) are formed by forward-looking agents. This is assumed to follow the following simple process:

$$g_{ijt+1}^{exp} = \beta . g_{it} + (1-\beta) . g_{ijt}^{exp} \tag{5}$$

Thus, we assume that the expected contribution of another player in the next period is a convex combination of the expectation for the current period and the player's own contribution for the current period. The parameter β is thus measuring the expected reciprocity from the interaction partners. Interestingly, it should be noticed that such a mechanism is identical to the implementation of a simple reinforcement learning model where the adaptation of the expectation from one period to the next (i.e. $g_{ijt+1}^{exp} - g_{ijt}^{exp}$) is a function of the player's belief about his counterpart's prediction error in the current period (i.e. $g_{it} - g_{jit}^{exp}$; with g_{jit}^{exp} standing for *i*'s belief about *j*'s expectation about *i*'s contribution in *t*). The only hypothesis we need to make in order for this equivalence between the two models to hold is that a player assumes for the others a simpler expectation formation mechanism than the one he uses. In line with cognitive hierarchy models, player *i* assumes that player *j* will take his own contribution as his expectation about *i*'s behavior. Under this hypothesis, our specification is similar to a reinforcement learning model.

Before continuing, we can illustrate how this model would work by considering the simple case of a two-player game. An agent will consider some action as a reference. He will then observe the action chosen by the other agent and will compare it to his reference. This comparison is what we call the emotional impulse. An action perceived by the agent as kinder than his reference will generate a positive impulse and will increase the social tie. On the other hand, an action perceived as less kind than the reference will trigger negative emotions and will lower the social tie (eventually to negative values). Thus, the social tie an agent will develop towards another can be seen as a stock variable of the emotional impulses triggered by the interaction. The strength of the impact of an impulse on the tie is determined by δ_i^2 (the tie-proneness or emotionality parameter) and the persistence of this impact on the tie is captured by δ_i^1 (the tie-persistence or memory parameter). We can also point out that reciprocity is introduced in the model through the tie-proneness parameter. Indeed, if $\delta_i^2 = 0$, the counterpart's behavior does not enter the tie and the utility function. On the other hand, a strictly positive value of the parameter will increase the tendency of the agent to reciprocate: actions from the other perceived as unkind will trigger negative ties that will push the agent to try to lower the other's payoff while actions perceived as kind will generate a positive tie and will entail more prosocial decisions from the agent. The tie mechanism described in equation 3 constitutes the affective part of our dual-process model. Direct cognitive control can be exercised neither on the emotional impulse nor on the social tie. These variables cannot simply be chosen on purpose by the agent to achieve a certain goal. Preferences are formed through emotional reactions triggered in an 'autonomic' way by the interaction.

Moving to the cognitive part of the model, we implement a probabilistic choice mechanism using the traditional logit form. The choice probability of each action is thus determined by:

$$\pi_{ikt} = \frac{\exp(\theta_i . V_{ikt})}{\sum_k \exp(\theta_i . V_{ikt})} \tag{6}$$

where π_{ikt} is the probability that agent *i* chooses action k ($k \in [1; K]$) in round t ($t \in [1; T]$). The parameter θ_i captures the distance between the agent's behavior and perfect utility maximization. When $\theta_i \to 0$, every action is equiprobable and behavior is random. On the other hand, when $\theta_i \to +\infty$, the agent is choosing with certainty the action maximizing is utility. The log-likelihood of choosing a certain stream of decisions over the whole game for a given individual can thus be written:

$$\mathrm{LogL}_i = \sum_{t=1}^{T} \sum_{k=1}^{K} d_{ikt} . \ln(\pi_{ikt}) \tag{7}$$

where $d_{ikt} = 1$ when action *k* is chosen in round *t* and is null otherwise.

We can already see that this model captures various important features of potential behaviors. First, we saw that, by setting λ_i to zero, agents will not be forward looking to the next period and will act like myopic individuals. Moreover, it has to be noticed that this limited foresight gives a more strategic aspect to the model: indeed, a forward-looking agent would be able to contribute to the public good for purely strategic reasons (i.e. without developing any social tie with other agents) as he expect his own contribution to have a positive impact on the other players next contributions. This feature appears very important as it can potentially explain the decay of contributions towards the end of a finitely repeated game (i.e. the often called "end effect"). Second, it obviously embeds the standard selfish preferences model, simply by allowing the values of α_{ij0} (the initial social tie) and of δ_i^2 both being null. Third, stable 'standard' social preferences can be obtained by having $\alpha_{ij0} \neq 0$ (for a specific or for all j), $\delta_i^1 = 1$ (so that there is no decaying trend in the tie in between periods of play) and $\delta_i^2 = 0$ (so that there is no update of the value of α_{ij0}). The direction of such stable preferences will then depend on the value of α_{ij0}: spitefulness will occur if $\alpha_{ij0} < 0$ while altruism arises if $\alpha_{ij0} > 0$. Interestingly, more complex distributional social preferences like inequality aversion can also be approached by our model. Consider the case of a symmetric public good game where contributions are costly and assume that agents are taking their own contribution as a reference. A higher contribution of the agent compared to his counterpart in one period not only implies disadvantageous inequality, but also a negative emotional impulse in our social tie framework. In both cases, the agent will supposedly decrease his contribution for the next period. In the same fashion, a lower contribution than the other will create advantageous inequality and a positive impulse and will thus push the agent towards a higher contribution in the next round. More generally, the social tie model can capture phenomena related to direct and indirect reciprocity and the in- and out-group distinction (see van Winden, 2012).

3. Testing the Model on Public Good Game Data

An important goal of this paper is to apply our model to various experimental datasets. We will restrict ourselves to three experiments involving a repeated public good game using a partner matching protocol (sections 3.1 and 3.2). The first two datasets (Bault et al., 2013; van Dijk et al., 2002) concern two-player public good games while the third one (Sonnemans et al., 2006) involves a four-player game. These datasets will allow us to estimate the model at the group-level (section 3.3). Based on these estimations, dynamic within-sample as well as out-of-sample predictions about the development of the game across periods are generated (section 3.4).

3.1. Experimental Designs

All three experiments use the same general structure with a repeated public good game interleaved between two monetary allocation tasks. In this section, we will simply discuss the design features of the public good game. For an extended description of all experimental procedures used during these experiments, please see Bault et al. (2013), van Dijk et al. (2002) and Sonnemans et al. (2006).

3.1.1. Bault et al. (2013)

In Bault et al. (2013), a subject in a Magnetic Resonance Imaging (MRI) scanner was playing simultaneously with a subject outside the scanner a non-linear public good game for 29 rounds. In each round, subjects were endowed with 12 tokens to split between a public and a private account. A token in the public account yielded 14 monetary units (MU) while the value of i markers in the private account was $32*i-i^2$ MU, so the non-linearity of the game comes from the private account. On top of this, a fix cost of 160 was subtracted each period. In this configuration, both the standard Nash Equilibrium (a contribution of 3 tokens to the public account) and the Pareto Optimum (a contribution of 10 tokens to the public account) are interior in the action space. This non-linear version, which leaves room for "reward" and "punishment"[10], was chosen for prediction matters. Indeed, a linear public good game would have led to "all-or-nothing" predictions depending on the relative values of the social tie and of the return of tokens to the public account. Also, in a linear public good game, any null or negative social tie would lead to a null contribution whereas in our game, subjects with anti-social and purely selfish preferences would not select the same contribution level. The same holds for the difference between agents with a social tie equal to one and agents for whom it is greater than one.

Once both subjects made their decisions, they were asked about their expectation about the contribution of the other during the current round (without incentives to avoid more complexity). After that both subjects indicated their expectations, they received feedback about the round. The feedback screen showed the contribution of the other, their payoff in MU for the past round and their cumulated payoff in MU over the whole public good game. Important is the fact that the experiment was totally anonymous and that we made sure that the two subjects never met at any point.

3.1.2. van Dijk et al. (2002)

The experimental procedures of van Dijk et al. (2002) and Sonnemans et al. (2006) have the same general design as the one of Bault et al. (2013). These two experiments were only behavioral and were run in the CREED laboratory, in Amsterdam. van Dijk et al. (2002) used a non-linear public good game repeated for 25 periods, in which participants were matched in pairs (fixed for the whole duration of the game). In this game, each subject has an endowment of 10 tokens to allocate to a public account, which yields 14 MU to both members of the pair,

[10] Contributing more than the Pareto-Optimum increases the other's payoff at a cost for oneself and also decreases the total payoff of the pair. On the other hand, contributing less than the Nash equilibrium decreases both payoffs, but the other's payoff more strongly.

or a private account, where the value of i tokens was $28i$-i^2 MU. A fixed cost of 110 MU was subtracted each period. The standard Nash equilibrium of the one shot game is again to contribute 3 tokens to the public account. The Pareto optimal choice is the full contribution (10 tokens).

3.1.3. Sonnemans et al. (2006)

Sonnemans et al. (2006) follow the same general timeline. Subjects were matched in fixed groups of four subjects and played a nonlinear public good game for 32 periods. The endowment of 10 tokens could be invested in a public account, which yields 7 MU to each of the four members of the group, or in a private account, where the value of i tokens was $21i$-i^2 MU. Fixed costs of 60 MU were subtracted each round, with the one-shot game Nash equilibrium being a contribution of 3 tokens to the public good and the Pareto optimum a contribution of the whole endowment. At the end of a round, information about the individual contributions was made available to all members of the group. Subjects got to know also their payoff for the round.

As can be seen, the three experimental designs are very similar. These three different datasets provide a good opportunity to check the robustness of the behavioral findings obtained in different frameworks.

[INSERT FIG 1 ABOUT HERE]

3.2. Aggregate Behavioral Results

We will first present aggregated behavioral results of the three experiments. In Bault et al. (2013), the average contribution over the whole game is 6.30 tokens (over 12 tokens available), which corresponds to 52.5% of the endowment. Interestingly, the average contribution to the public good is very stable over the 29 rounds, ranging from 5.26 tokens in the last round to 7.00 tokens in round 18 and 22 (see Figure 1). The very often observed end-effect is very slight here since average contribution drops by only 1.07 tokens between round 28 and round 29. This quite high level of contribution and the absence of any decreasing trend might be due to the pair setting.

The data from van Dijk et al. (2002) exhibit a little bit more cooperation as the mean contribution is at 6.01 tokens (over 10 tokens available; 60% of the endowment). Starting from 5.00 tokens in the first round, cooperation is slowly increasing to stabilize slightly above 6.50 tokens until the last rounds where a more noticeable end-effect took place. Interestingly, the data from Sonnemans et al. (2006), concerning four-player groups, reveal an even higher level of cooperation with an average contribution of 6.52 tokens (65% of the endowment). Once again, we have the pattern of a progressive increase towards stabilization around 7 tokens and an important decrease over the last periods where contributions fall at 3.39 tokens. However, as we will see, these aggregated results are hiding a great heterogeneity at the individual level that we will try to understand.

3.3. Estimation Procedures at the Group Level

As we have seen in section 1, there is evidence that many people do not plan ahead and uses backward induction in order to make a decision, even in very simple games. This was a hint for the use of a mixture-model approach that would allow subjects to be from different types, either myopic or forward looking. However, in order to estimate the forward-looking model, we need data about expected contributions from the other in t in order to build expectations for the next period $t+1$ (see equation 5). Since these data are only available for Bault et al. (2013), we will use the mixture approach on this dataset only. Estimations on van Dijk et al. (2002) and Sonnemans et al. (2006) will only consider the myopic model (i.e. with λ constrained to be equal to 0), which will turn out to be not so unreasonable, however.

In our case, the myopic and forward-looking models are nested. Thus, we cannot apply a mixture approach at the choice level (see for example Harrison and Rutström, 2009; Wang and Fischbeck, 2004) where a decision is considered to have a certain probability of being generated by each model, the likelihood of a choice being the combination of the two models weighted by a "mixture probability". In that case, because of the linearity of our utility function (see equation 1), this probability and the discount parameter would be perfectly confounded and thus impossible to estimate. More conceptually, because of the dynamic nature of the model and of the path dependence of each contribution decision, reasoning at the choice level appears neither relevant nor correct. As a consequence, we opted for a mixture model at the individual level, where each individual can be categorized as being of one of the two alternative types. Once again, it was impossible for us to mimic a procedure previously used by El-Gamal and Grether (1995) in their article on Bayesian updating (for greater details, especially on their method and algorithm, see op. cit., p1141). Indeed, while they attribute subjects' type according to the best-fitting model too, their models are fixed decision rules that do not require any optimization. However this is not the case for us: we do have to classify subjects in two categories (myopic or forward-looking) but we also have to find, for each group, the optimal parameters values. This main difference constrains us to categorize subjects based on estimations at the individual level in order to classify subjects based on optimal parameters' values (i.e. to make our classification at the optimum). We chose to categorize subjects based on the results of a Likelihood Ration Test (LRT; at the 5% level).[11] Based on the categorization, we then estimate the corresponding model on the whole sample using maximum likelihood. Each individual contributes to the total likelihood based on the likelihood corresponding to the model selected by the classification. In the end, we get one set of parameters for each group (myopic and forward-looking) and a categorization of individuals between each group.[12]

Different static and dynamic specifications of the reference point for the emotional impulse have been considered, with the standard Nash contribution (which is 3 tokens) showing the best performance. [13] Higher contributions will generate positive emotional impulses that will increase the social tie while lower contributions will generate negative ones and thus decrease the current social tie so that the development of both positive and negative ties is possible. Another assumption concerning our data is that players start with a null tie: there are no antecedents in the relationship so we consider that subjects start with selfish preferences that may evolve throughout the interaction.[14]

Concerning forward-looking individuals and the way they form expectations for $t+1$, we have to mention that β, the weight representing expected reciprocity, is interacting with λ, the weight put on future utility (see equations 1 and 5). Indeed, the more weight is put on future utility, the more profitable it becomes profitable to contribute in the current period due to the reciprocity expected in the next round. In the same way, for a given λ, the utility for the contribution in t increases with the expected reciprocity. This is why we cannot estimate both parameters together. Thus, we chose to fix β to an arbitrary fixed value of 0.5 and to estimate λ.[15] The last remark concerns the agent's own contribution in $t+1$. Since the agent will try to choose the optimal contribution in t (g_{it}) and since g_{it} and $g_{it+1}$are independent (due to the

[11] On top of results based on a LRT at the 5% level shown in table 1, results based on the categorization obtained from the BIC, AIC and the LRT at the 10% level are shown in Appendix A.

[12] For comparison with the other two datasets, we also estimate the myopic model over the whole sample. Results of this estimation are consistent and can be found in Table 2.

[13] As alternative reference points the following have been considered (see Appendix B): the expected contribution by the other, one's own contribution, no (zero) contribution, the Pareto-optimal contribution, and the standard Nash contribution. The standard Nash contribution level may have been a focal point during the experiments because it is easy to retrieve in the payoff table available to the subjects, but most probably also because it represents the demarcation between competitive and prosocial actions.

[14] This hypothesis can of course be disputed since it has been shown in many cases that some people exhibit to some extent unselfish preferences even in one-shot games using double-blind procedure. In order to address this concern, we computed a value for the starting tie based on choices made during the monetary allocation task run before the public good game. Results of estimations with this value of the initial tie are shown in Appendix C. Quite expectedly, estimates are not changing that much since this value is only affecting significantly the first periods of the game.

[15] Results are similar if β is estimated instead of λ (see Appendix D).

additivity of the utility function; see equation 4), g_{it+1} is actually factored out and thus does not require any supplementary assumption.

[INSERT TABLE 1 ABOUT HERE]

Results of this estimation are presented in Table 1. We get 15 subjects classified as forward-looking (the corresponding set of parameters will be indexed by FL in the text) while the other are considered as myopic (the corresponding set of parameters will be indexed by M in the text). As expected, the more conservative classification is restricting the forward-looking sub-group such that the proportion of forward-looking subjects is slightly above one fourth of our sample.[16] We can remark a few important things concerning parameters estimates. First of all, the estimates for θ_M and θ_{FL} are in line with results obtained from previous experiments testing social preferences models in public good games environments (Corrazini and Tyszler, 2010). Second, all parameters are significant at the 1% level, meaning that both immediate reciprocity and long-term history are significantly impacting decisions. As we expected, deviations from the selfish Nash contribution are generating significant impulses and thus influencing the social tie. This influence last for a significant number of periods showing that dynamics matter and that actions from past rounds still impact current choices. Third, about a quarter of the subjects attributes a significant weight to future utility. Finally, we do not see big differences between the two sub-groups in terms of parameters values for $\delta^1$and δ^2. Both groups have values of δ^1 slightly below 0.5 and values of δ^2 around 0.075.

[INSERT TABLE 2 ABOUT HERE]

Table 2 presents the results of group-level estimations concerning the three datasets using the myopic model. It is interesting to notice that results from van Dijk et al. (2002) look very close to what we obtain with the Bault et al. (2013) data (when using the same model). This suggests a certain robustness of our results since two very similar experimental settings (an almost identical game played in pairs) yield very similar estimates. Results are slightly different for the data from Sonnemans et al. (2006) where the game was played in groups of four. All parameters are still significant but we can notice that δ^1 (tie persistence) takes a lower value, meaning that the history of play that impacts a decision is shorter. On the other hand, δ^2 (the tie proneness parameter) takes the highest value of our three samples so that

[16] In their article, Bone et al. (2009) found that over half of the population was acting as myopic. This proportion corresponds to the one we obtained with the AIC and BIC criterions. The LRT classification reduces it a bit further.

impulses have a stronger immediate impact on tie formation. This simultaneous decrease in δ^1 and increase in δ^2 is hinting at the fact that subjects' immediate reciprocity is more intense in this game. This may be due to a harder information retrieval because keeping track of all the individual streams of contributions among the group is more difficult in larger groups. As a consequence, subjects may base their behavior much more on the immediate history of play.

3.4. Behavioral Predictions

Our next task consists in formulating predictions about the dynamics of individual contributions to the public good. We will use parameters estimated at the group-level, in line with what is mostly done in economics. In order to check the quality of the fit, we will first present results of predictions made on each datasets using the parameters estimated on the same dataset. Even though this may not be the most challenging task, since estimation and predictions are made on exactly the same data, but the interest of this first exercise is to be informative: even in this 'comfortable' environment, is the model able to track round-to-round changes in contributions?

For the Bault et al. (2013) data predictions using group-level estimates, we attributed to every individual the set of parameters corresponding to the sub-group (myopic or forward-looking) into which he was classified and predictions were formulated according to the corresponding model.[17] For the two other datasets, we used the estimates of the myopic model. Using the parameters values, we are able to predict individual contribution behavior round after round but some assumptions are required. First, we take as given the first round contribution, as we have no special assumption concerning the first round of play[18]. Second, we also take as given the behavior of the other player and then generate predictions based on this "pre-determined" behavior from counterparts. Indeed, in this interactive choice framing, letting the two players' behaviors free would have led to very poor results since the main factor driving behavior in a repeated game is the behavior of the others. Such a strategy would have more to do with a simulation approach than with predictions. Also in line with what is

[17] Given the slight differences between parameters estimates between categorizations (see Appendix B), no big changes should be expected if we would have used a different categorization.

[18] Notice that as $\alpha_{ij0} = 0$, all myopic players should be selfish and thus play the Nash contribution in the first round. The picture is more complex for forward-looking individuals since it would depend on several factors (α_{ij0}, $\lambda_i$and the effect of contributions in t on the expectation in t+1). As a consequence, we let the first round predicted contribution as the real one and let the social tie start building (if there is any) from round two onward. Thus, predictions start from round two.

usually done when predicting behavior based on model estimations, the predicted contribution is the one that is the most probable according to the model.

[INSERT FIG 2 ABOUT HERE]

The predictions using group-level estimates for each of the three datasets appear in Appendix E while Figure 2 presents predictions for a subset of representative interactions. In these figures, P (G) represents the subject's pair (group) number and S the subject's number within the pair (group). The first thing to notice about the Bault et al. (2013) data is that for a huge majority of pairs the model seems to do a good job in terms of quality of predictions. Contributions of pairs exhibiting a quite stable behavior for a significant number of periods (like P3, P7, P11, P18, P19 or P23) are unsurprisingly well-captured by predictions. For pairs exhibiting more complex contributions dynamics, two different cases can be observed: either the model does a very good job in predicting round-to-round changes (like for P1, P2, P8, P12 or P22 for example) or it is only able to track the general trend of behavior, as if the predictions' line was a smoothed version of the real data line (like for example P4, P5, P10, P12, P13 or P21). Except maybe for P9 where the two players chose opposite strategies without adapting to the other, there is no pair for which predictions are completely off the mark, which is very encouraging.

Similar conclusions can be drawn for the van Dijk et al. (2002) data. Results are even better with these data since predictions are very close to real behavior for a huge majority of pairs. Even though the overall fit of predictions looks also good for the Sonnemans et al. (2006) data, we can notice a slight deterioration due to some groups where behavior was very chaotic and where coordination failed (like G3 or G14). However, it should be noticed that our model appears to perform well in predicting individual contributions to the public good when individuals make their decisions in groups of four, a task rarely undertaken in economics.

In order to go further into testing the model, we will present results of out-of-sample predictions. To achieve this, we used the parameters estimates of the myopic model of one of our three datasets (see Table 2) in order to predict behavior in the two remaining ones. We repeated this exercise three times, in order to cover all potential combinations. Good predictions in this more demanding exercise will be a good argument in favor of the model explanatory power over individual behavior in voluntary contribution environments.

To assess statistically the quality of fits, Table 3 presents the average absolute prediction error as well as the Root Mean Square Errors (RMSEs) for all the predictions. The best performance appear to be for the van Dijk et al. (2002), maybe because quite a few pairs reach a stable state quite quickly. It is to be noticed that even though the performance in the four-player case is a bit worse, it remains quite close to in terms of the error measures. Overall, the decrease in performance is of low magnitude when shifting from within- to out-of-sample predictions. The highest increase in RMSEs concerns the Bault et al. (2013) data for which the measure rises by 15% or 26% when predicted out-of-sample. The raise is limited to 5% or less for the two other datasets.

4. Extended Analyses

This section presents complementary analyses on the Bault et al. (2013) data in order to give a more complete picture of the performance of the model as well as of individual heterogeneity. We will first compare the performance of our model to several fixed social preferences models. We will then present results about the heterogeneity of individual parameter estimates.

4.1. Model Comparison

Our previous results from the estimation of the model on the three datasets show that several behavioral models can be rejected. First of all, because δ^2 is significantly different from zero, we can reject the standard model of selfish preferences.[19] More interestingly, because δ^1 is also significantly different from zero, we can rule out the hypothesis that subject would adapt their behavior only based on the last period of play, like it is implicitly assumed by most reciprocity models and would hold for tit-for-tat behavior. Longer-term past history matters and counterparts' actions continue to significantly influence contribution behavior for a substantial number of periods.

We next compare the performance of the model to that of a fixed social preferences model where the utility can be described as follows:

$$U_{it} = P_{it} + \sum_{j=1}^{N} \alpha_i . P_{jt} \tag{8}$$

[19] Estimation of the myopic version of the model does not enable to distinguish strategic motives behind the contribution behavior. Nevertheless, if all subjects were contributing only for selfish reasons, they would choose the Nash contribution in the last period, since there were no more gains to be expected from cooperation at that point. However, this is not the case: the mean contribution in the last period is 5.268 tokens and is significantly different from 3 (Wilcoxon signed-rank test : z = 0.346 ; *p-value* = 0.001).

The only change compared to equation (2) resides in the subscript next to α: here we do not have neither a *j* nor a *t*, meaning that the social attitude is generalized towards all other counterparts (i.e. not individualized) and time invariant. Table 4 presents the parameter estimates of this model, and the AIC and BIC scores of these estimations as well as the same scores for the corresponding estimations of the (myopic) social tie model (see Table 2). The two information criteria are clearly in favor of the endogenous and dynamic social ties model.

[INSERT TABLE 4 ABOUT HERE]

As we already mentioned in section 2, the social ties model generates behavior that is similar to inequality aversion if players' own contribution is the reference contribution in the impulse. Our results in section 3, however, show that this specification can be rejected (see also Appendix B). Because subjects' expectations regarding their partner's contribution in each round is available in the Bault et al. (2013) dataset, we can more directly specify an inequality aversion model of the Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) type. A caveat is immediately in order, though, because the authors of these models do not aim to study dynamic problems and the evolution of play. Nevertheless, it seems still interesting to investigate how this type of model performs, especially since the (non-strategic) myopic model appeared to perform well on this data set. To that purpose, we estimated a model with the following expected utility:

$$U_{ikt} = P_{ikt} - \alpha.\max\left(P_{jkt} - P_{ikt}; 0\right) - \beta.\max\left(P_{ikt} - P_{jkt}; 0\right) \tag{9}$$

where P_{ikt} is the expected payoff of player *i* in period *t* for a contribution level *k*, given the expectation *i* has about *j*'s contribution in *t*. Thus, α represents disadvantageous inequality aversion while β represents advantageous inequality aversion. Concerning the estimation, we only kept the constraints of positivity of the parameters as the question of the relative values of α and β as well as the limitation of the value of β to one appear to us as an empirical question. Estimation of this model renders the following estimates for these parameters: $\theta = 0.017$, $\alpha = 0.944$ and $\beta = 1.586$. We can first notice the very low value of θ, reflecting the lack of grasp of this model on the data. Moreover, these results show that the weight of disadvantageous inequality aversion is lower than the one of advantageous inequality aversion. This suggests that for a given payoff, one would rather earn less than the other, compared to more than the other. On top of this strange result, it appears that individuals would be willing to pay more than one dollar to reduce advantageous inequality by one dollar, which is also against the original formalization of Fehr and Schmidt (1999). All in all, this suggests that this model is maybe not fit as it is strictly formalized now to capture this kind of problems, as it was already suggested by the authors.

4.2. Estimations at the Individual Level

To give an idea of the heterogeneity of the parameters at the individual level, we estimated our model on each individual separately. Even though the task of estimating three or four parameters on a sample with a limited number of observations (between 25 and 32, depending on the dataset) may be detrimental to the estimation precision, it can still give an impression of the variability of individual parameters values that our group estimates are hiding. Figure 3 shows the distribution of individual parameters values for 127 subjects[20]. All parameters exhibit a wide range of individual values. The range of estimates of δ_i^1 is quite broad and the distribution appears uniform except for the mode at 0. This suggests that, even though most individuals take a longer-term history into account, a non-negligible minority of our sample seems to exhibit only immediate reciprocity. For δ_i^2 we can notice that the distribution is quite flat between 0 and 0.2 (90% of the values are below 0.218, the median being 0.092). Only about 10% of the subjects seem not to develop any social ties at all ($\delta_i^2 = 0$). Similar observations can be made concerning λ_i, with a quite uniform distribution between 0.2 and 0.95 and a mode at 1.[21]

A question arising from the observation of this heterogeneity is to what extent the taking into account of these individual differences – by using individual instead of group-level estimates – can improve the predictive performance of the social ties model. This will be informative to gauge the potentialities of the model in terms of predictive power. Table 5 displays the average RMSEs and average absolute errors of such predictions. Even though the performance is improved, the average RMSEs decreases are not huge, spreading from 11% for the van Dijk et al. (2002) data to 18% for the Sonnemans et al. (2006) data. Finally, for exploratory reasons, we compute Pearson correlations coefficients between individual parameters to see if there are specific relationships between them. We find a negative and significant correlation between δ_i^1 and δ_i^2($\boldsymbol{z} = -0.204$; $p - value = 0.014$) and regarding λ_i only a weakly significant positive correlation with δ_i^2 ($\boldsymbol{z}$ = -0.489**;** *p-value* = 0.064; where the relatively few observations regarding λ_i may play a role). The negative correlation

[20] We had to remove the parameters of 23 subjects from this figure. Nineteen of them were removed because parameters took very extreme values due to a lack of variability of contribution behavior. The algorithm did not converge for the four remaining subjects.

[21] It concerns only the fifteen subjects classified as forward-looking from the Bault et al. (2013) data set.

between $\delta_i^1$and δ_i^2 suggests that a longer lasting impact of impulses goes together with a smaller impact on the social tie.

5. Discussion

This paper proposes a model of the dynamic development of social ties where preferences are made endogenous and dynamic. We believe that this psychologically grounded model represents a natural way to understand repeated social interactions. When people have the possibility to track their partners' behavior, which is the case in most real life interactions, affective bonds will be created dependent on how each partner emotionally experiences the behavior of the other. This is how friendships or hate relationships are built, which may also be relevant for economic interactions such as in case of team work and business connections.[22] We showed that this model is quite general and embeds several other models of social preferences. Confrontation with data from repeated public good games - involving both pairs and groups of four players - showed that the model is able to explain the dynamic patterns of contributions as well as to perform well when predicting behavior out-of-sample.

Interestingly, several recent studies suggest the usefulness and relevance of making preferences endogenous and dynamic as our social tie model does. First, Goette et al. (2012) show that there is a qualitative difference in behavior in the interaction between groups formed under the minimal group paradigm and groups that have experienced real social interactions. More specifically, they find that in-group favoritism when making cooperative decision is stronger when social ties are present. They also find differences in punishment behavior in the sense that groups involving social ties do not punish more strongly out-group defectors than in-group defectors - whereas groups formed according to the minimal group paradigm do - but that they punish more strongly when the victim of defection is an in-group member. These important differences make them conclude that "both conceptually and empirically, economists should take into account that social ties are an important factor in group interactions, within organizations and societies." (op. cit.; p114). Closer to our modeling concerns, Malmendier and Schmidt (2012) are studying gift giving in a three-parties setting and show that gift giving strongly affects the recipient's decisions in favor of the gift giver even if this comes at the expense of a third party. In their view, "a gift creates a bond between the giver and the recipient of the gift. Before the gift is given the decision maker is equally concerned about the welfare of all other players. However, once he receives the gift

[22] "Colleagues in office, partners in trade, call one another brothers; and frequently feel towards one another as if they really were so. Their good agreement is an advantage to all." Adam Smith, The Theory of Moral Sentiments, Part VI, section II.

the welfare of the gift giver gets a higher weight in DM's utility function." (op. cit., p32). As they mention, this finding is in line with the social tie model proposed by van Dijk and van Winden (1997) while other models of social preferences fail to explain this result. They also develop a new model, quite similar to ours[23], where "by giving or withholding a gift the potential gift giver receives a larger or smaller weight in the utility function of the decision maker" (op. cit., p35), rendering this weight endogenous. Building on these results, Liang and Meng (2013) study the impact of social connections (through club membership) on indirect reciprocity. In an indirect investment game where the trustor and the beneficiary of the trustee's decision are not the same person, they find that a social connection between the trustor and the beneficiary increases the repayment of the trustee, but only when the trustor has been kind enough. They explain this result by the conjunction of two facts: first, a sufficiently trusting decision from the trustor creates a positive tie with the trustee; second, social connections are transitive ("friends of my friends are also my friends"). As a consequence, the trustee is more generous to the beneficiary as he anticipates that it will please the trustor, a person he is now positively caring about. Findings that are in line with our model.

This study constitutes a first step in the testing of the social ties model and it opens many doors for future research. Because of the nature of the datasets we chose to tackle, this article has mostly focused on positive social ties resulting from cooperative interactions. Indeed, even if antisocial behavior was an option that was to some extent available to the subjects in our non-linear public good games[24], the room for pro- and anti-social behavior was far from being symmetric. Interesting avenues for future work would be to focus on more symmetric environments where negative ties are potentially as possible to occur as positive ties. Related to that is the potential relevance of this model in competitive environments, which may very well depend on the market structure. Indeed, in a very atomistic market where agents are price-takers and cannot have a direct impact on their competitors, ties are not expected to develop. However, in case of restricted competition (e.g. an oligopoly), ties could very well matter. While positive ties could generate collusion, negative ones could yield very aggressive (cut-throat) competition. In the first case, this may even reduce the effectiveness of anti-trust actions since the collusion would take place because of ties, which are tacit, and not because of explicit agreements.

[23] Given that they want to explain their experimental data, which concerns a one-shot game, they do not consider any forward-looking behavior.

[24] Remember that by contributing less than the interior Nash equilibrium contribution, subjects were decreasing the other's payoff at their own cost, *ceteris paribus*.

Finally, an important question concerns the stability of the parameters making up the tie mechanism (related to tie proneness and tie persistence). First of all, it would be interesting to test the stability of the group-level estimates across different games to see whether we can retrieve comparable values or, if this is not the case, to study and try to explain the differences. The second direction concerns individual parameters. One of the features of this model is that its parameters do not directly represent preferences but are presumed to be linked to psychological traits, which suggest that they might be more stable across environments and interactions. A test of this conjecture would be to have the same participant play several games with different counterparts to see whether his parameters stay close.

References

Barron, G., and E. Yechiam (2009): "The coexistence of everestimation and underweighting of rare events and the contingent recency effect," *Judgment and Decision Making*, 4(6), 447-460.

Battigalli, P., and M. Dufwenberg (2007): "Guilt in Games," *American Economic Review*, 97(2), 170-176.

Baumeister, R. F., and M. R. Leary (1995): "The need to belong: Desire for interpersonal attachments as a fundamental human motivation," *Psychological Bulletin*, 117(3), 497-529.

Bault, N., B. Pelloux, J. J. F. Fahrenfort, K. R. Ridderinkhof, and F. van Winden (submitted): "Neural dynamics of social tie formation in economic decision-making".

Bolton, G. E., and A. Ockenfels (2000): "A theory of equity, reciprocity and competition," *American Economic Review*, 90(1), 166-193.

Bone, J., J. D. Hey, and J. Suckling (2009): "Do people plan?," *Experimental Economics*, 12(1), 12-25.

Brandts, J., A. Riedl, and F. van Winden (2009): "Competitive rivalry, social disposition, and subjective well-being: An experiment," *Journal of Public Economics*, 93(11-12), 1158-1167.

Camerer, C. F., T. H. Ho, and J. K. Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119(3), 861-898.

Charness, G., and M. Rabin (2002): "Understanding Social Preferences with Simple Tests," *The Quarterly Journal of Economics*, 117 (3), 817-869.

Conte, A., and J.D. Hey (2012): "Assessing multiple prior models of behavior under ambiguity," Discussion Papers 12/01, Department of Economics, University of York.

Corrazini, L., and M. Tyszler (2010): "Preference for Efficiency or Confusion? A Note on a Boundedly Rational Equilibrium Approach to Individual Contributions in a Public Good Game," ISLA Working Papers, 37.

Cox, J. C., D. Friedman, and S. Gjerstad (2007): "A tractable model of reciprocity and fairness," *Games and Economic Behavior*, 59(1), 17-45.

Denrell, J. (2005): "Why most people disapprove of me: Experience sampling in impression formation," *Psychological Review*, 112, 951−978.

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

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

Dufwenberg, M., and G. Kirchsteiger (2004): "A theory of sequential reciprocity," *Games and Economic Behavior*, 47(2), 268-298.

El-Gamal, M. A., and D. M. Grether (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90(432), 1137-1145.

Fahrenfort, J. J., F. van Winden, B. Pelloux, M. Stallen, and K. R. Ridderinkhof (2012): "Neural correlates of dynamically evolving interpersonal ties predict prosocial behavior," *Frontiers in Neuroscience*, 6:28.

Falk, A., and U. Fischbacher (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54(2), 293-315.

Fehr, E., and K. M. Schmidt (1999): "A Theory of Fairness, Competition, and Cooperation," *The Quarterly Journal of Economics*, 114(3), 817-868.

Fehr, E., and K. M. Schmidt (2006): "The Economics of Fairness, Reciprocity and Altruism – Experimental Evidence and New Theories," in *Handbook of the Economics of Giving, Altruism and Reciprocity*, Eds S.-C. Kolm and J. Mercier Ythier, 1, 615-691.

Glimcher, P. W. (2005): "Indeterminacy in Brain and Behavior," *Annual Review of Psychology*, 56, 25-56.

Goette, L., D. Huffman and S. Meier (2012): "The Impact of Social Ties on Group Interactions: Evidence from Minimal Groups and Randomly Assigned Real Groups," *American Economic Journal: Microeconomics,* 4(1), 101-115.

Harrison, G. W., and E. E. Rutstrom (2009): "Expected utility theory and prospect theory: one wedding and a decent funeral," *Experimental Economics*, 12(2), 133-158.

Isaac, M. R., J. M. Walker and A. W. Williams (1994): "Group size and the voluntary provision of public goods: experimental evidence utilizing large groups," Journal of Public Economics, 54, 1-36.

Johnson, E. J., C. F. Camerer, S. Sen, and T. Rymon (2002): "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *Journal of Economic Theory*, 104(1), 16-47.

Kreps, D. M., P. Milgrom, J. Roberts, and R. Wilson (1982): "Rational cooperation in the finitely repeated prisoners' dilemma", *Journal of Economic Theory*, 27(2), 245-252.

Kirchsteiger, G. (1994): "The role of envy in ultimatum games," *Journal of Economic Behavior & Organization*, 25(3), 373-389.

LeDoux, J. E. (1996): *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, New York: Simon & Schuster.

Ledyard, J. O. (1995): “Public Goods: A Survey of Experimental Research,” in *The Handbook of Experimental Economics,* eds. A. Roth & J. Kagel. (Princeton, NJ: Princeton University Press), 111–181.

Levine, D. K. (1998): “Modeling Altruism and Spitefulness in Experiments,” *Review of Economic Dynamics,* 1(3), 593-622.

Liebrand, W. B. G. (1984): “The Effect of Social Motives, Communication and Group-Size on Behavior in an N-Person Multi-Stage Mixed-Motive Game,” *European Journal of Social Psychology*, 14**,** 239-264.

Loewenstein, G., and T. O'Donoghue (2007): “The heat of the moment: Modeling interactions between affect and deliberation,” Working Paper.

Loomes, G., and R. Sudgen (1982): “Regret Theory: An Alternative Theory Of Rational Choice Under Uncertainty,” *Economic Journal*, 92(368), 805-824.

Malmendier, U., and K. M. Schmidt (2012): “You owe me,” Working Paper.

Matthews, K. A., C. D. Batson, J. Horn, and R. H. Rosenman (1981): “‘Principles in his nature which interest him in the fortune of others…’: The heritability of empathic concern for others,” *Journal of Personality,* 49, 237–247.

Mehrabian, A., and N. Epstein (1972): “A measure of emotional empathy,” *Journal of Personality*, 40**,** 525-543.

Rabin, M. (1993): “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, 83(5), 1281-1302.

Rakow, T., and B. R. Newell (2010): “Degrees of Uncertainty: An Overview and Framework for Future Research on Experience-Based Choice,” *Journal of Behavioral Decision Making*, 23**,** 1-14.

Rubinstein, A. (1979): “Equilibrium in supergames with the overtaking criterion,” *Journal of Economic Theory*, 21-1, 1-9.

Rushton, J. P. (1991): “Is Altruism Innate?” *Psychological Inquiry,* 2-2, 141-143.

Sobel, J. (2005): “Interdependent Preferences and Reciprocity,” *Journal of Economic Literature*, 43(2), 392-436.

Sonnemans, J., F. van Dijk, and F. van Winden (2006): “On the dynamics of social ties structures in groups,” *Journal of Economic Psychology*, 27(2), 187-204.

Tvesky, A., and D. Kahneman (1191): “Loss Aversion in Riskless Choices: A Reference-Dependent Model,” *Quarterly Journal of Economics*, 106(4), 1039-1061.

Wang, M., and P. S. Fischbeck (2004): “Incorporating Framing into Prospect Theory Modeling: A Mixture-Model Approach,” *Journal of Risk and Uncertainty*, 29(2), 181-197.

Wilcox, N. T. (2007): “Predicting Risky Choices Out-of-Context: A Monte Carlo Study,” University of Houston Working Paper.

van Winden, F. (2001): “Emotional Hazard Exemplified by Taxation-Induced Anger,” *Kyklos*, 54, 491-506.

van Winden, F., M. Stallen, and K. R. Ridderinkhof (2008): “On the nature, modeling, and neural bases of social ties,” *Adv Health Econ Health Serv Res*, 20**,** 125-159.

van Winden (2012): “Affective Social Ties – Missing Link in Governance Theory,” *Rationality, Markets and Moral*, 3, forthcoming.

Zajonc, R. B. (1984): “On The Primacy of Affect,” *American Psychologist*, 39(2), 117-123.

Figures

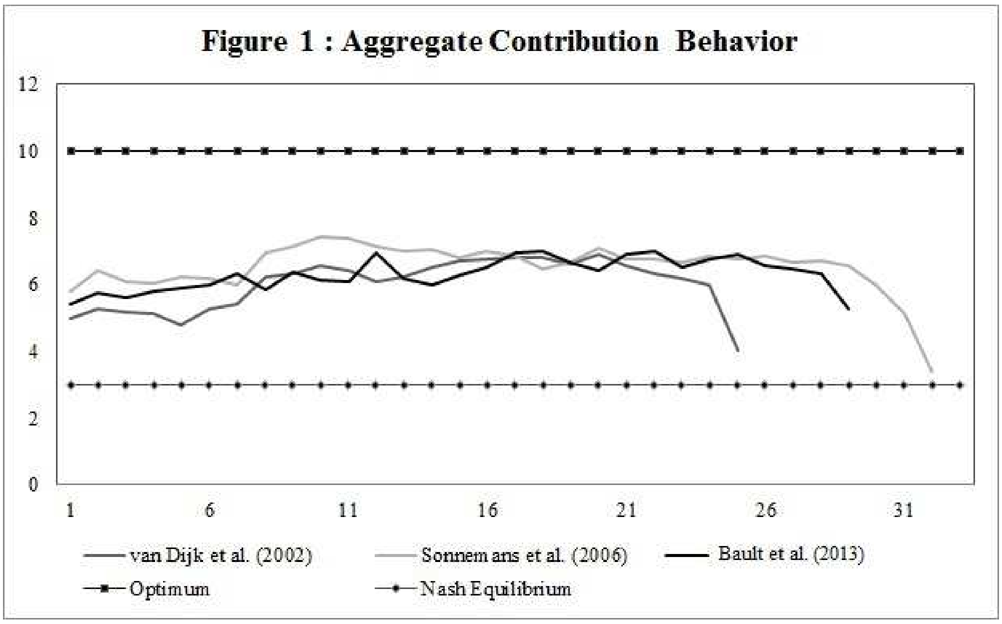

Figure 1 : Aggregate Contribution Behavior
12
10
8
6
4
2
0
1
6
11
16
21
26
31
van Dijk et al. (2002)
Sonnemans et al. (2006)
Bault et al. (2013)
Optimum
Nash Equilibrium

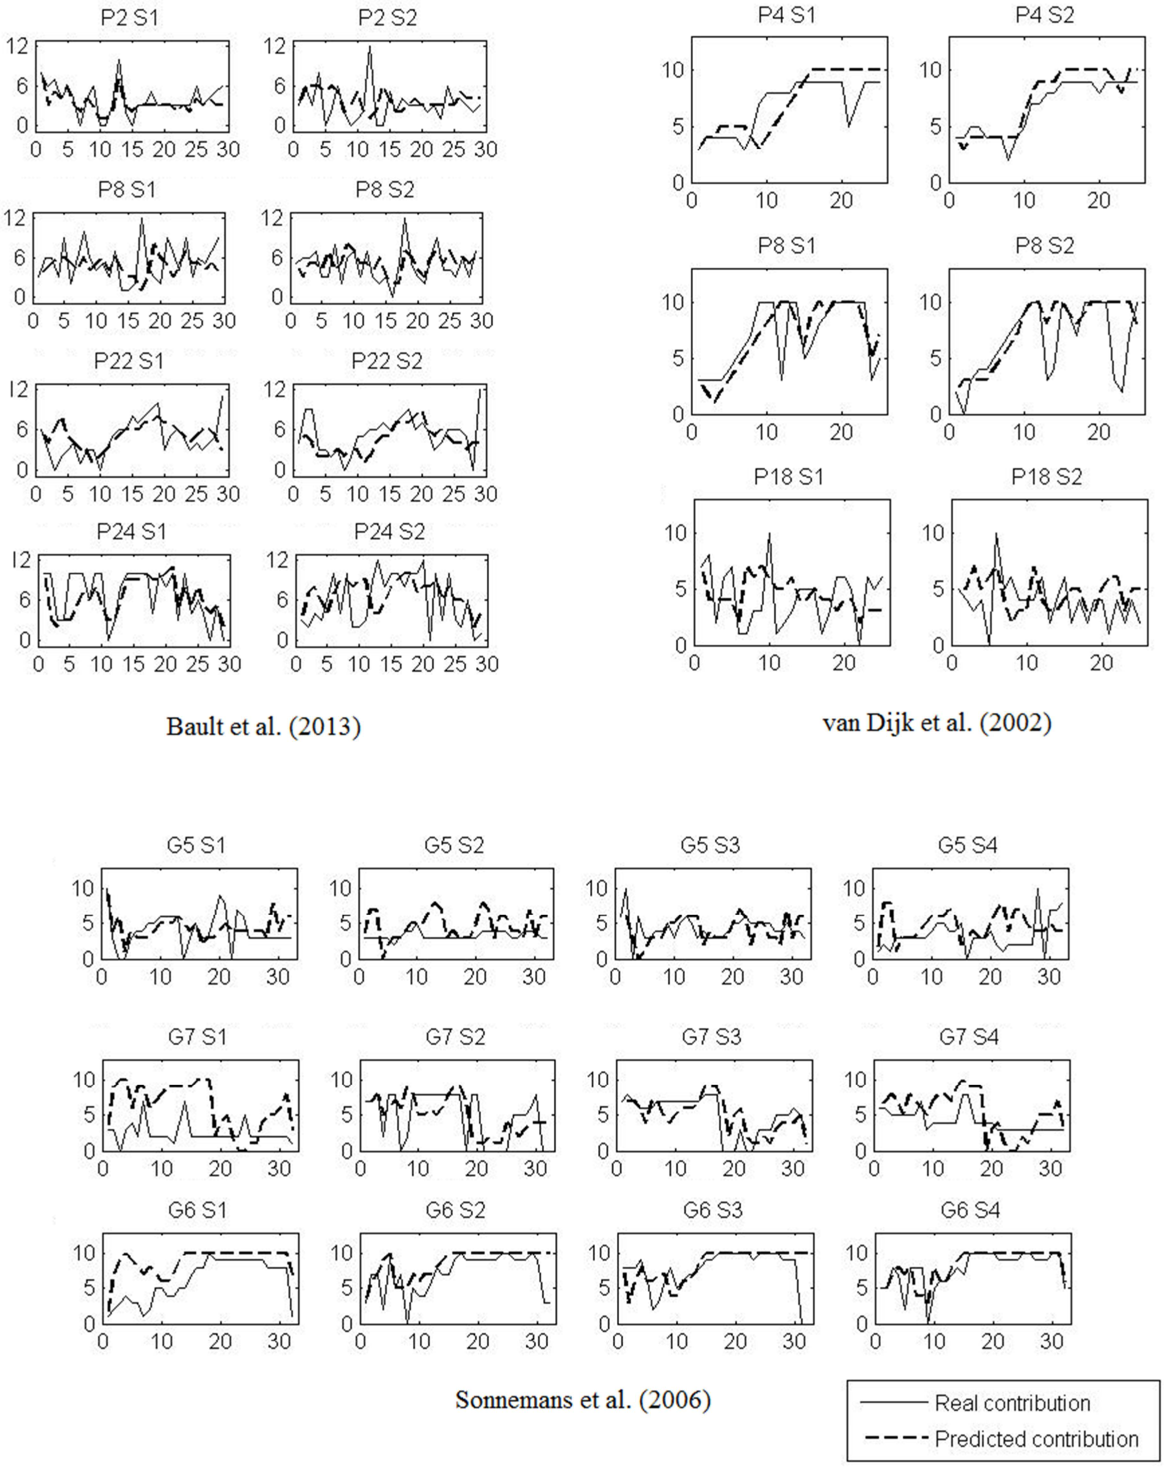

Figure 2: Within sample behavioral predictions (selected interactions)

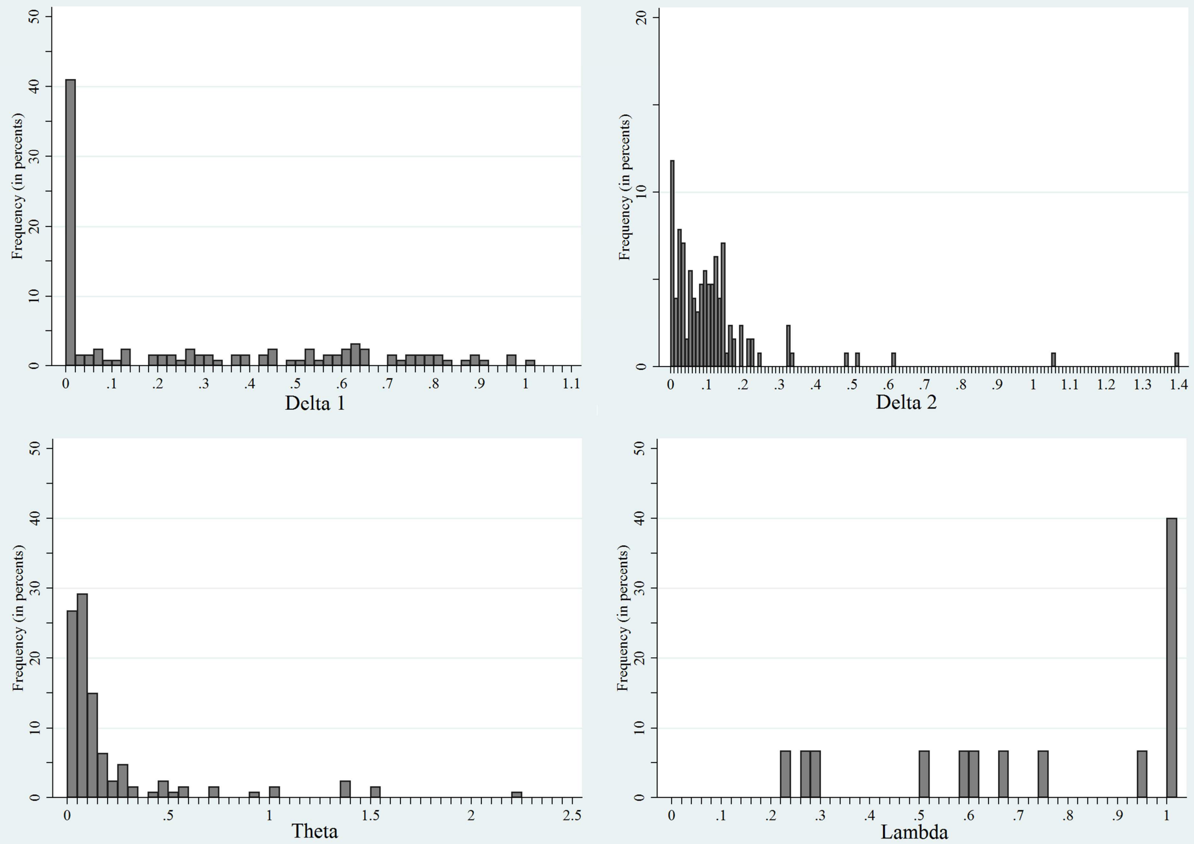

Figure 3: Distributions of individually estimated parameters

Tables

Table 1: Group Level Estimates (Bault et al., 2013)

Parameters	Myopic Group (Std Err)	Forward-Looking Group (Std Err)
θ	0.0418***	0.1053***
	(0.0027)	(0.0092)
δ^1	0.4756***	0.5302***
	(0.0529)	(0.0659)
δ^2	0.0742***	0.0743***
	(0.0081)	(0.0097)
λ	-	0.5049***
		(0.0482)
Log-Likelihood	-3671.20	
Forward-Looking Subjects	15 (out of 56)	

Note: *** indicates significance at the 1% level.

Table 2: Group Level Estimates

Parameters	Van Dijk et al. (2002) (Std Err)	Sonnemans et al. (2006) (Std Err)	Bault et al. (2013) (Std Err)
θ	0.0813***	0.0319***	0.0448***
	(0.0057)	(0.0030)	(0.0024)
δ^1	0.5489***	0.1840***	0.5094***
	(0.0370)	(0.0554)	(0.0384)
δ^2	0.0861***	0.1460***	0.0786***
	(0.0067)	(0.0133)	(0.0058)
Log-Likelihood	-1857.19	-3672.05	-3737.14
AIC	3720.38	7350.10	7480.28
BIC	3725.29	7356.18	7486.36

Note: In order to be able to make meaningful comparisons, the fourth column presents the results of the estimation of the “myopic” model (without the mixture approach) on the Bault et al. (2013) dataset. *** indicates significance at the 1% level.

Table 3: Average RMSEs from predictions based on group-level estimates

Predicted dataset	Van Dijk et al. (2002)	Sonnemans et al. (2006)	Bault et al. (2013)
		Within-sample predictions	
	2.006	2.794	2.475
	(1.333)	(2.002)	(1.7149)
		Out-of-sample predictions	
Set of parameters estimates used for prediction :			
van Dijk et al. (2002)	-	2.935	2.850
		(2.288)	(2.046)
Sonnemans et al. (2006)	2.103	-	3.142
	(1.373)		(2.374)
Bault et al. (2013)	1.908	2.643	-
	(1.257)	(1.874)	

Note: Average absolute error appears in parentheses.

Table 4: Fixed preferences model estimation

Parameters	Bault et al. (2013) (Std Err)
θ	0.0187***
	(0.0022)
α	0.5303***
	(0.0264)
Log-Likelihood	-4121.21
AIC	8246.42
BIC	8250.47

Table 5: Average RMSEs from predictions based on individual estimates

	Van Dijk et al. (2002)	Sonnemans et al. (2006)	Bault et al. (2013)
Average RMSE	1.783 (1.109)	2.332 (1.484)	2.166 (1.508)

Note: Average absolute error appears in parentheses.

Appendix A: Mixture-model estimation of Bault et al. (2013) with other classification criterions

Classification Criterion	BIC		AIC		LRT (p<0.10)	
Parameters	M	F-L	M	F-L	M	F-L
θ	0.043	0.070	0.042	0.071	0.040	0.082
δ^1	0.487	0.500	0.486	0.486	0.463	0.539
δ^2	0.070	0.081	0.069	0.082	0.076	0.074
λ	-	0.350	-	0.355	-	0.352
Log-Likelihood	-3692.49		-3686.07		-3689.05	
Forward-Looking Subjects	24		25		20	

Note: For exposition purposes standard errors are not displayed. All parameters are significant at the 1% level. The M (resp. F-L) column refers to the parameters of the myopic (resp. forward-looking) group.

Appendix B: Log-likelihoods from the myopic model estimations using different reference points

Reference Point	Fahrenfort et al. (2012)
Predicted Other's Contribution	-4020.62
Own Contribution	-4036.39
(Fixed) Null Contribution	-3752.07
(Fixed) Pareto-Optimal Contribution	-3864.29
(Fixed) Nash Equilibrium Contribution	-3671.20

Note: The reference point is subtracted to the contribution of the other to generate the emotional impulse. The reference point used in the rest of the paper (because of better performance) is the Nash equilibrium contribution of the one-shot game. The mixture model was used here.

Appendix C: Estimation with the starting value of the social tie based on the Social Value Orientation test

	Bault et al. (2013)		Van Dijk et al. (2002)	Sonnemans et al. (2006)
Parameters	M	F-L	M	M
θ	0.043	0.070	0.061	0.031
δ^1	0.487	0.500	0.448	0.190
δ^2	0.070	0.081	0.113	0.148
λ	-	0.350	-	-
Log-Likelihood	-3659.39		-1900.79	-3671.98

Appendix D: Estimation of the social tie model fixing λ and estimating β

Group Level Estimates (Bault et al., 2013)

Parameters	Myopic Group (Std Err)	Forward-Looking Group (Std Err)
θ	0.0418***	0.1042***
	(0.0026)	(0.0096)
δ^1	0.4756***	0.5320***
	(0.0526)	(0.0594)
δ^2	0.0742***	0.0740***
	(0.0069)	(0.0087)
β	-	0.2514***
		(0.0256)
Log-Likelihood	- 3675.81	
Forward-Looking Subjects	15 (out of 56)	

Note: We keep the same classification than when estimating the model with λ (since models with the β formalization are not directly nested anymore). We fixed λ to be equal to 1. Logically, when β is fixed to 0.5, we get λ=0.505 while when λ is fixed to nearly twice this value (i.e. 1), the value of β compensates and is nearly divided by 2 (i.e. is equal to 0.251).

Appendix E: Behavioral within sample predictions

Bault et al. (2013)

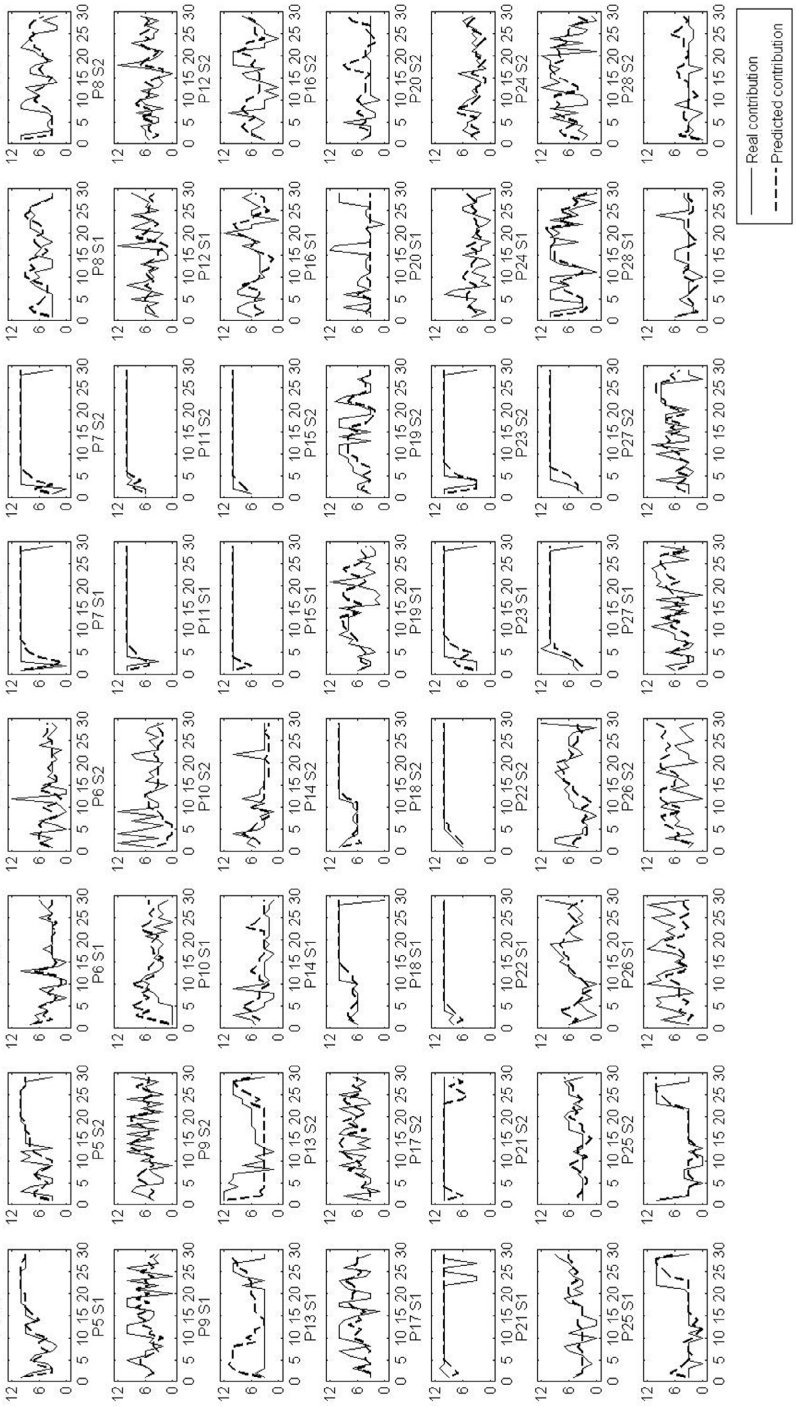

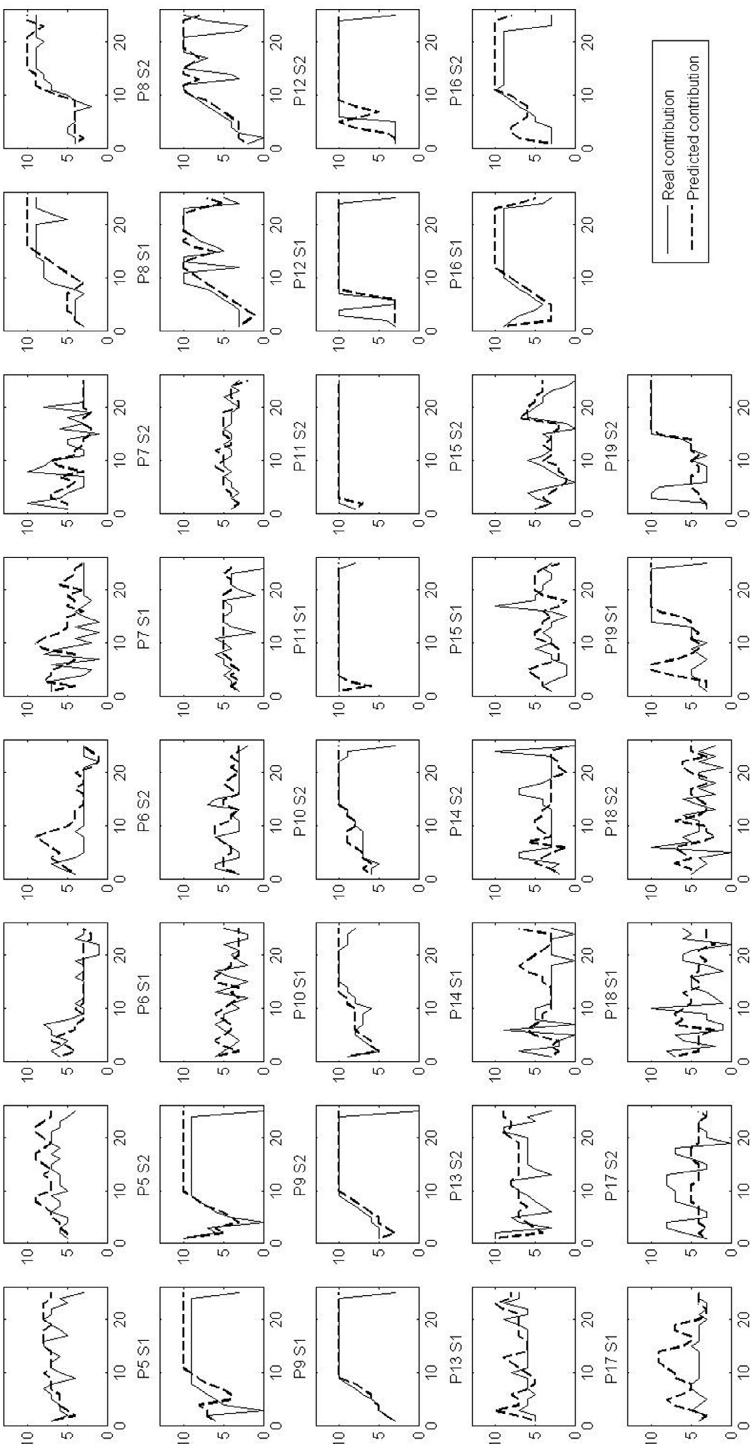
P1 S1
P1 S2
P2 S1
P2 S2
P3 S1
P3 S2
P4 S1
P4 S2
P5 S1
P5 S2
P6 S1
P6 S2
P7 S1
P7 S2
P8 S1
P8 S2
P9 S1
P9 S2
P10 S1
P10 S2
P11 S1
P11 S2
P12 S1
P12 S2
P13 S1
P13 S2
P14 S1
P14 S2
P15 S1
P15 S2
P16 S1
P16 S2
P17 S1
P17 S2
P18 S1
P18 S2
P19 S1
P19 S2
Real contribution
Predicted contribution

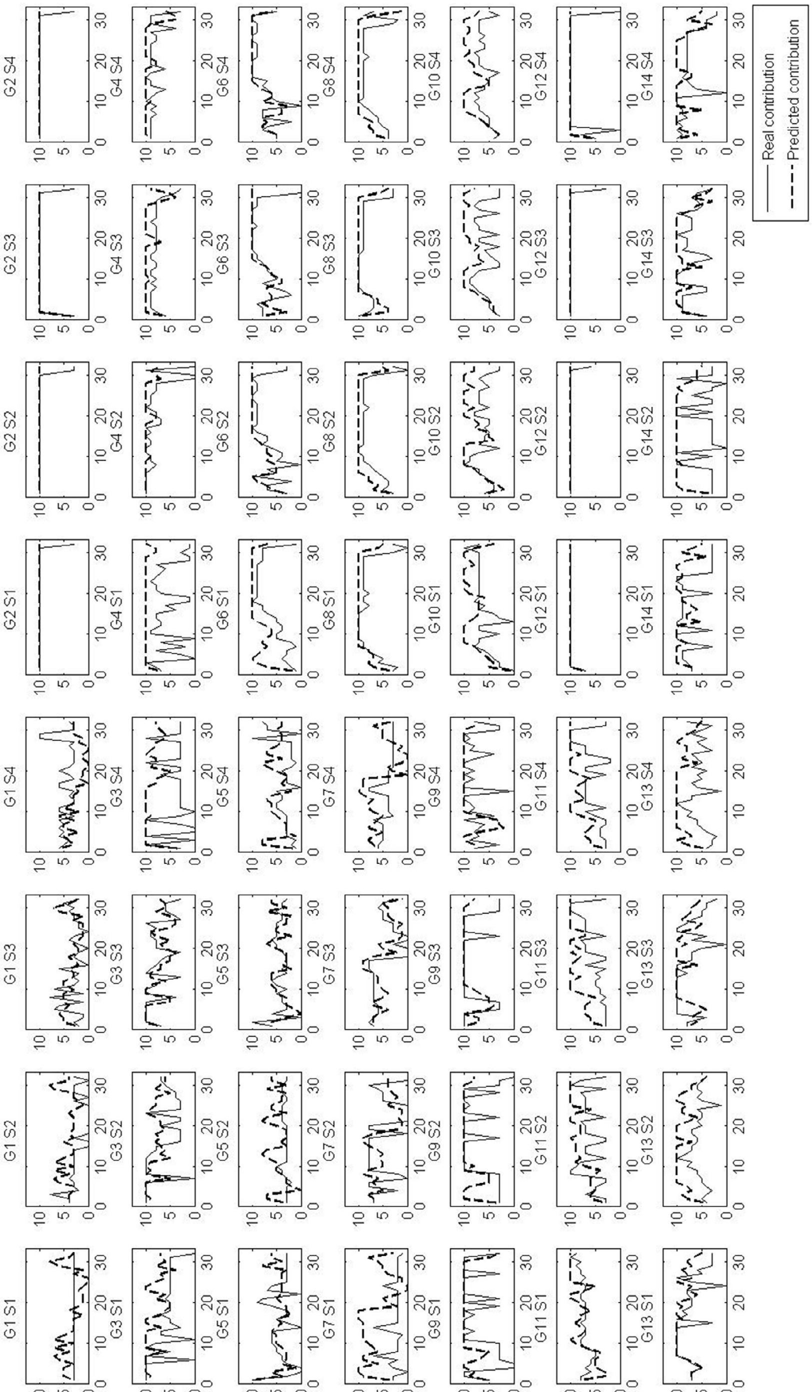
Real contribution
Predicted contribution