A theory of reciprocity with incomplete information

Citation for published version (APA):

Vostroknutov, A. (2012). *A theory of reciprocity with incomplete information*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 015
<https://doi.org/10.26481/umamet.2012015>

Document status and date:

Published: 01/01/2012

DOI:

[10.26481/umamet.2012015](https://doi.org/10.26481/umamet.2012015)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Alexander Vostroknutov

**A Theory of Reciprocity with
Incomplete Information**

RM/12/015

METEOR

Maastricht University School of Business and Economics
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands

A Theory of Reciprocity with Incomplete Information

Alexander Vostroknutov*

Maastricht University, The Netherlands

January 2012

Abstract

A model of belief dependent preferences in finite multi-stage games with observable actions is proposed. It combines two dissimilar approaches: incomplete information (Levine, 1998) and intentionality (Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006). Incomplete information is important because social preferences are not directly observable; intentions are found to be indispensable in explaining behavior in games (Falk, Fehr, and Fischbacher, 2008). In the model it is assumed that the players have social attitudes that define their social preferences. In addition, players care differently about the payoffs of other players depending on their beliefs about their social attitude and possibly on the beliefs of higher orders. As the game unfolds players update their beliefs about the types of other players. An action of a player shows intention when she chooses it anticipating future belief updating by others. A reasoning procedure is proposed that allows players to understand how to update beliefs by constructing a sequence of logical implications.

JEL classification: C72, C73.

Keywords: reciprocity, dynamic games of incomplete information, belief dependent preferences.

*I would like to thank Christian Bach, Pierpaolo Battigalli, Friederike Mengel, Joshua Miller, Andrés Perea, Elias Tsakas and all the participants of the seminars at Maastricht University, Università Bocconi and International ESA Meeting in Copenhagen for invaluable comments and suggestions. All mistakes are mine.

Correspondence: Department of Economics, P.O. Box 616, Maastricht 6200 MD, The Netherlands
e-mail: a.vostroknutov@maastrichtuniversity.nl

1 Introduction

Many experimental studies have shown that strategic behavior is influenced by reciprocity and/or various forms of social preferences (e.g. Fehr and Gächter, 2000; Falk, Fehr, and Fischbacher, 2008). This lead to the attempts to model such behavioral phenomena in games. Two broad classes of models have emerged: 1) *incomplete information* models, in which players in the game have private information about their own social preferences, but not about social preferences of others (Levine, 1998) and 2) *intentionality* models, where players infer good or bad intentions of the opponents from the observed actions and act upon their inferences (Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006; Cox, Friedman, and Gjerstad, 2007).

Both approaches are important in order to understand the effects of behavioral “biases” on strategic interactions.¹ For example, Fischbacher and Gächter (2010) provide strong evidence that people are highly *heterogeneous* in their propensity to conditionally cooperate in public goods game. This means that incomplete information is crucial for the modeling of social preferences and reciprocity since these traits cannot be directly observed *before* the game starts. In a very elegant experiment McCabe, Rigdon, and Smith (2003) show that the behavior in a simple sequential game depends on the availability of additional, rarely taken, action. This demonstrates that people in strategic situations evaluate other players’ *intentions* and react to counter-factuals or to what *could have happened* in the game.² Therefore, a good theory of reciprocal behavior has to include some form of evaluation of others’ intentions.

Even though the models mentioned above capture important behavioral features they have certain drawbacks. In the standard incomplete information models (like Levine, 1998) it is assumed that each player has *fixed* social utility which is unobserved by others. This implies that players in these models are unable to react to any intentionality as it requires a *change* in attitude towards other player depending on her action.³ Moreover, it is not inconceivable that people are able to “use” different social preferences when interacting with different people. The simplest example of this is food sharing practices among kin versus non kin. Fiske (1990) reports that in-

¹By “bias” I mean deviations from pure self-interest.

²Falk, Fehr, and Fischbacher (2008) also show importance of intentions.

³In the incomplete information game behavior *can* depend on the observed action of other player if the observer Bayesian-updates her beliefs regarding the type of the other player. However, this change is fundamentally different from reaction to intentions. In the case of Bayesian updating player just changes her best response because she has better information about the type of the other player and thus about other player’s future choices.

digenous people of the Moose Culture willingly divide food among kin, but rather reluctantly share it with other tribesmen. More economically relevant example is status seeking preferences in urban ethnically similar populations. Members of these groups invest into goods that signal their social status. However, they do not take into account the status of the outsiders (Charles, Hurst, and Roussanov, 2009). In addition, Cox, Friedman, and Gjerstad (2007) provide an extensive list of experimental results that can only be reconciled under the assumption that subjects' social preferences depend on the roles of others and context.

Intentionality models like DK and FF (correspondingly: Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006) incorporate intentions by assuming that players assess relative *kindness* of the action of the opponent by comparing payoffs that *could have been achieved* should the opponent choose differently with the payoffs that *can be achieved* given the actual choice. The action is deemed kind if the opponent chose so that achievable payoffs are better than now non-achievable (and unkind if the opposite is true). After the kindness of the action is calculated, player responds to kindness with kindness and vice versa.⁴ In both papers an equilibrium concept is developed to deduce the optimal behavior. It is also assumed that players have personal preferences parameter that regulates the strength of their reciprocal response to kind actions. Therefore, in equilibrium players should know the parameters of others *with certainty*. However, as was mentioned above, it is not clear how players can obtain this information before the game begins. Thus, the realism of the intentionality approach suffers from the lack of incomplete information.

The goal of this paper is to develop a framework in which incomplete information about social preferences of others coexists with inferences about intentionality (and, therefore, with changes of utility depending on these inferences). In order to achieve this a somewhat different perspective should be taken on *what it means to infer information about intentions*. In particular, the questions are: 1) How should such inferences be made if players have incomplete information about social preferences of others? and 2) How should reciprocity be defined in this situation? The approach taken in this paper is the following. There is a set of social utilities that both players can possibly exhibit and the set of "labels" called *social attitudes*.⁵ These objects are *game independent* similarly to standard incomplete information case where utilities of the players are defined only on the payoffs and not on any other features of the game. Before the

⁴The utility function changes in order to reflect this desire of the player.

⁵In what follows the case of two players will be considered for the ease of exposition. Social utility, thus, means utility defined over payoffs of a player herself and payoffs of the other player.

game each player knows her social attitude but is uncertain (holds probabilistic belief) about social attitudes of other players; their beliefs about her social attitude; and so on. The new element in the model is the connection between social attitudes and social utilities. In particular, it is assumed that social utility of a player can depend on her beliefs about social attitudes of the other player and beliefs of higher orders. Thus, each player's hierarchy of beliefs about social attitudes gives rise to the hierarchy of beliefs about social utilities. As the game unfolds players can update their beliefs about social attitudes depending on the game structure and the actions taken by the opponent. But, since social utilities of the players depend on their beliefs about social attitudes of others, this belief updating will also change players' social utilities.

The construction just described allows players to have incomplete information about other player's social utility and at the same time makes it possible for them to change their "attitude" towards other player depending on what the other did. To illustrate, suppose there are two social attitudes: *Selfish* and *Nice*. If a player has social attitude *Selfish*, then she has self-interested utility u_S regardless of the beliefs about other player's social attitude.⁶ However, if player is *Nice*, then her social utility depends on her beliefs. If she thinks that other player is also *Nice*, then her social utility u_N is altruistic (increases in the payoff of other player); but if she believes that the other person is *Selfish* and thus does not care about her payoff, then she does not feel like she has to be any different towards the other player and thus her utility is also self-interest (u_S). Given these assumptions, *Nice* player can now make inferences about social attitudes of the other player by observing actions taken in the game. If some action is only consistent with the other player being *Nice* (and believing that she is *Nice*), then her social utility becomes u_N which corresponds to reciprocal response.

The model is rather flexible and can be used to test variety of hypotheses regarding reciprocal behavior. For example, in the previous paragraph (positive) reciprocity was defined as altruism of *Nice* players towards other *Nice* players. However, there exists another view on the nature of reciprocal relationships. It is, in particular, possible that people who positively reciprocate do it not because they become *genuinely* altruistic towards others, but because they only want others to *believe* that they are positive reciprocators. Such belief of others can be very important for profitable future interactions. In our framework this can be captured by making utility depend on the second order beliefs (beliefs of the other player). We can assume that if player is *Selfish* then her utility is u_S if she believes that other player thinks she is *Selfish*;

⁶ u_S depends only on the payoffs of the player herself.

and $u_S + b$, where b is a positive constant, if she believes the other player thinks she is Nice. The same can be true about Nice player. The games then can be found in which two models of reciprocity make differential predictions.

The paper is structured as follows. Section 2.1 gives main definitions. In section 2.2 it is shown how the players reason about what is expected to happen in the game. Section 2.3 deals with the belief updating mechanism. Section 2.4 describes general reasoning procedure. Section 2.5 relates the model to psychological games. And section 3 provides some examples.

2 The Model

The model consists of two main ingredients: the game and the description of the behavioral types. Consider any finite extensive form game of perfect information with two players and the finite set Θ of *attitudes*. For example, Θ can be the set $\{S, N\}$ where S stands for Selfish and N stands for Nice. These attitudes reflect the social preferences of one player over the payoffs of the other player. In addition, players' social preferences can depend on what they *believe* the other player's attitude is and what other player's beliefs are. Moreover, as the game unfolds the beliefs of the players might change depending on what behavior they observe, thus influencing their social preferences.

To be more precise let $\mathbb{R}$ be the real line and $\mathbb{R}^2$ the space of all possible payoff pairs of both players. Let $\mathcal{P}$ be some finite set of utility functions $u : \mathbb{R}^2 \rightarrow \mathbb{R}$. In our example we can define $\mathcal{P} = \{u_S, u_N\}$ where $u_S(x_1, x_2) = x_1$ is self-interested utility function and $u_N(x_1, x_2) = x_1 + \alpha x_2$ is altruistic utility function ($0 < \alpha < 1$). Fix some integer $T \geq 1$ and define the *belief dependent utilities* to be a mapping $\mathcal{B} : \Theta^T \rightarrow \mathcal{P}$. To illustrate suppose $T = 2$ and consider $\mathcal{B}$ defined by the list:

$$\begin{aligned}\mathcal{B}(S, S) &= u_S \\ \mathcal{B}(S, N) &= u_S \\ \mathcal{B}(N, S) &= u_S \\ \mathcal{B}(N, N) &= u_N.\end{aligned}$$

The interpretation is the following. Selfish player who *believes* that other is also selfish (S, S) has utility u_S . Selfish player who believes that other is nice (S, N) has also utility u_S (in the end, this is what selfishness is all about). Nice player who believes that other is nice (N, N) acts altruistically (u_N). But if nice player believes other is

selfish he stops being altruistic and becomes selfish (u_S). This example shows how reciprocity can be modeled in the current framework: the utility of the player with attitude N changes depending on her beliefs about the attitude of the other player. Thus, if other player during the game gives her a reason to think that he is selfish she will adjust her utility accordingly. It is important to notice that this possibility of utility updating is crucial for modeling reciprocity like behaviors. If we define utilities as independent of beliefs, then we will be in the standard incomplete information setup where one cannot model situations with preferences dependent on characteristics of others.

For higher T the construction above can represent more complicated social attitudes of the players. For example, if $T = 3$, then $\mathcal{B}$ would assign different utility to each triplet $(\theta_1, \theta_2, \theta_3) \in \Theta^3$. The meaning is that player's utility depends on what he believes other player is and what other player believes he is. For example, it is possible to assume that nice player who thinks the other is nice and thinks other thinks he is selfish (N, N, S) might have some disutility in comparison to the case (N, N, N). This is plausible because if others think that the player is selfish they won't cooperate in social dilemma environments as much as those who think that the player is nice. Thus the model can incorporate preferences over beliefs of other player (and beliefs about beliefs about beliefs...).

2.1 Behavioral Types

In this section we construct *behavioral types* as infinite hierarchies of beliefs about attitudes Θ . Instead of standard recursive formulation we use a definition of belief hierarchy on an arborescence.⁷ This is done in order to simplify the future exposition. However, there is an obvious way to translate the arborescence representation into a recursive one (see below).

Consider a countably infinite set $\mathcal{K}$ of nodes and a partial order P on it with the following properties:

- K1** $\exists_{\kappa \in \mathcal{K}} \forall_{k \in \mathcal{K}} \kappa P k$
 There exists an upper bound $\kappa \in \mathcal{K}$
- K2** $\forall_{k, n, \ell \in \mathcal{K}} k P \ell \wedge n P \ell \Rightarrow k P n \vee n P k$
 The partial order is "tree like"

⁷Here and further, the term arborescence is used in graph theoretic sense.

K3 $\forall_{k \in \mathcal{K}}$ the set $C_k := \{n \in \mathcal{K} \setminus k : kPn \wedge (\nexists_{\ell \neq n, \ell \neq k} kP\ell Pn)\}$ of immediate successors of k is finite and non-empty

These assumptions assure that the pair $(\mathcal{K}, P)$ is an *arborescence* with finite number of branches coming from each node.

The interpretation of P is the following. Each node $k \in \mathcal{K}$ is associated with some player and his attitude in Θ . The immediate successors of k represent the support of the belief of this player about possible attitudes of the other player. The upper bound κ shows the original attitude of the player.

To be more specific consider a mapping $A : \mathcal{K} \rightarrow \Theta$ that assigns an attitude to each node in $\mathcal{K}$. Consider some function $p_k : C_k \rightarrow (0, 1]$ such that $\sum_{n \in C_k} p_k(n) = 1$. p_k represents the probabilities that a player assigns to the attitudes in the order of beliefs following k . In addition consider a mapping $\iota : \mathcal{K} \rightarrow \{1, 2\}$ that assigns one of the two players to each node in $\mathcal{K}$ with the following property: for any $k \in \mathcal{K}$ if $\iota(k) = i$ then $\iota(n) = -i$ for all $n \in C_k$.⁸

Before we proceed, an additional assumption on the structure of the arborescence is in order. The reason for this assumption is two-fold: 1) it plays an important role in the next section; 2) it is sufficient to represent the arborescence of beliefs in a standard recursive way with *finitely* many types (see below). Let

$$H_{\mathcal{K}} := \{(\kappa, k_1, k_2, \dots) \subset \mathcal{K} : k_1 \in C_{\kappa} \wedge k_t \in C_{k_{t-1}}, t = 2, 3, \dots\}$$

be the set of all infinite *linear paths* on the arborescence.⁹ Say that a path $(\kappa, k_1, k_2, \dots)$ has *common attitude belief of level l* if, for all $t \geq 1$, $A(k_{l+2t}) = A(k_l)$ and $A(k_{l+2t+1}) = A(k_{l+1})$.

Assumption 1. *There is a number L such that all paths in H have common attitude belief of level no more than L .*

Call the tuple $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$ that satisfies Assumption 1 a *behavioral type of player $i = \iota(\kappa)$* .

To illustrate consider the example on Figure 1:

Here $\Theta = \{S, N\}$ and the branches represent the arborescence P . k_1 is the upper bound of P that represents the attitude of the player himself ($A(k_1) = N$). The next level $(k_2, k_3, k_4$: immediate successors of k_1) represents the beliefs of the player

⁸Here $i = 1, 2$ and $-i$ denotes the other player.

⁹All paths on the arborescence are necessarily infinite by the property **K3**.

Figure 1: Arborescence of beliefs.

about the attitude of the other player. For example, with probability $p_{k_1}(k_2) = 0.2$ player thinks that the other is nice ($A(k_2) = N$); with probability $p_{k_1}(k_3) = 0.3$ player also thinks that the other is nice ($A(k_3) = N$); with probability $p_{k_1}(k_4) = 0.5$ player thinks that the other is selfish ($A(k_4) = S$). The important distinction between k_2 and k_3 comes from the difference in the third order beliefs: beliefs of the other about the attitudes of the player. In k_2 player believes that the other believes he is nice with probability 0.1 and in k_3 player believes that the other believes he is nice with probability 0.9. Depending on the actions of the other player it is possible that the player will update his beliefs by throwing away one of the two N -branches.

To see how the arborescence can be related to the recursive formulation of types let T_1 and T_2 be some finite sets of *attitude types* for players 1 and 2. Associate with each T_i a mapping $\theta_i : T_i \rightarrow \Theta$. This mapping represents the attitude of each attitude type. Let $b_i : T_i \rightarrow \Delta(T_j)$ represent the belief of attitude types of player i about attitudes of player j . An arborescence is represented by the tuple $\langle (T_i, \theta_i, b_i)_{i=1,2} \rangle$ if T_1 and T_2 are the sets of nodes for each player; θ_i are derived from function A ; and b_i are derived from functions p_k . Notice that Assumption 1 guarantees that T_1 and T_2 are finite. To see how this is achieved consider a partial representation of the arborescence on Figure 1:

This construction is given as example. Section 2.5 deals with the explicit relation of behavioral types to dynamic psychological games (Battigalli and Dufwenberg, 2009).

2.2 Reasoning about Expected Play

Consider a finite two-player extensive form game of perfect information. Let $\{1, 2\}$ be the set of players. G denotes the set of nodes with G^i being the set of nodes in which player i moves. $Z \subset G$ is the set of terminal nodes. Let A_g be the set of actions in node

Figure 2: Recursive representation of the arborescence of beliefs.

$g \in G \setminus Z$ and $\pi_i : Z \rightarrow \mathbb{R}$ the material payoff function of player i . Let $\Delta(A_g)$ be the set of probability distributions over actions A_g in node g . Denote by $S_i = \times_{g \in G^i} A_g$ the set of pure strategies of player i and by $S = S_1 \times S_2$ the set of all pure strategies. Let $S_i^\delta = \times_{g \in G^i} \Delta(A_g)$ be the set of behavioral strategies of player i and $S^\delta = S_1^\delta \times S_2^\delta$ the set of all behavioral strategies.

We define the reasoning about the game as it happens in the mind of a single player. Given some behavioral type any player should be capable of calculating how the game would unfold. There is no need to invoke other player in any way since all possible data about the beliefs of the other player are contained in the type tree under consideration.

The first step in understanding how a player reasons about the game is to figure out what he thinks other player *expects*. Consider the simplest behavioral type of player $i = \iota(\kappa)$ with $\mathcal{K} = \{1, 2, \dots\}$; P is a linear order inherited from $\mathbb{N}$; and $A : \mathcal{K} \rightarrow \Theta$ is some function that generates the sequence $\varepsilon = (\theta_1, \theta_2, \dots) \in \Theta^\infty$. Suppose that $\mathcal{K}$ has common attitude belief of level 3. This means that

$$\varepsilon = (\theta_1, \theta_2, \theta_3, \theta_4, \theta_3, \theta_4, \dots).$$

Given this behavioral type player i 's payoffs in the game are different from the original payoffs defined by the function π_i . Player i "perceives" the payoffs through the social utility $u_{1,T} := \mathcal{B}(\theta_1, \theta_2, \dots, \theta_T) = \mathcal{B}(\varepsilon_{1,T})$.¹⁰

Definition 1. For any players $i, j \in \{1, 2\}$ say that j **thinks i 's payoffs are determined by $u : \mathbb{R}^2 \rightarrow \mathbb{R}$** if player j considers i 's payoff in any terminal node $z \in Z$ to be $u(\pi_i(z), \pi_{-i}(z))$.

¹⁰Notation: $\varepsilon_{m,n} := (\theta_m, \theta_{m+1}, \theta_{m+2}, \dots, \theta_{m+n-1})$. $u_{m,n} := \mathcal{B}(\varepsilon_{m,n})$.

Say that *i*'s payoffs are determined by $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ if *i* thinks *i*'s payoffs are determined by $u : \mathbb{R}^2 \rightarrow \mathbb{R}$.

Thus, player *i*'s payoffs are determined by $u_{1,T}$. Analogously, *i* thinks that other player's payoffs (player *j*) are determined by $u_{2,T} = \mathcal{B}(\varepsilon_{2,T})$. However, this is not everything we can deduce. *i* also knows that

1. *j* thinks that *i*'s payoffs are determined by $u_{3,T}$
2. *j* thinks that *i* thinks that *j*'s payoffs are determined by $u_{4,T}$
3. *j* thinks that *i* thinks that *j* thinks that *i*'s payoffs are determined by $u_{5,T}$...

Given all this information we would like to construct the reasoning of player *i*. First, player *i* has to understand how player *j* would reason about the game. Thus, player *i* can put himself "in the shoes" of player *j* and consider the information that (*i* believes) player *j* has. Player *j* wants to predict how player *i* would choose in the game given that *i* has utility $u_{3,T}$ and thinks that *j* has utility $u_{4,T}$. Notice now that by assumption $\mathcal{K}$ has common attitude belief of level 3. This means that at any level of beliefs after 3 player *j* thinks that player *i* has utility $u_{3,T}$ and player *i* thinks that player *j* has utility $u_{4,T}$. Thus, both players have *common belief* about their utilities.

To summarize, player *i* thinks that player *j* thinks that players have common belief that their utilities are $u_{3,T}$ and $u_{4,T}$. Thus, player *i* can use Backward Induction on the game tree with the payoffs modified by $u_{3,T}$ and $u_{4,T}$ to arrive at the strategy $s_3 \in S_i$ that player *j* expects *i* to play.

The next thing to notice is that strategy $s_3 \in S_i$ was obtained by backward induction procedure together with $s_4 \in S_j$ under the belief that player *j* has utility $u_{4,T}$. However, player *i* believes that player *j*'s actual utility is given by $u_{2,T}$. Therefore, player *i* can expect that player *j* might not want to stick to the strategy s_4 , but will rather best respond to s_3 in a game where *j*'s payoffs are transformed by $u_{2,T}$. So, *i* should expect player *j* to play $s_2 := BR(s_3)$.¹¹

Overall, player *i* thinks that player *j* expects him to play according to s_3 and thinks that player *j* will play according to s_2 . Using the same logic we can now conclude that most likely it will not be rational for player *i* to follow s_3 as his real utility is $u_{1,T}$. However, since we have reached the beginning of the belief arborescence the deviations from s_3 will be considered *unexpected* by player *j* (at least in the mind of player *i*) and so *j* might then update his beliefs. Such unexpected deviations and

¹¹Strategy s_2 is formed by player *j* best responding in all nodes in which he moves given s_3 .

updates will be considered in the next section. For now though we need to formulate an important principle of expected play.

Principle 1. *Players do not update their beliefs if the game unfolds according to their expectations. All players assume this to reason about others and to reason about how others reason about them.*

Remember that the equilibrium path that j expects was reasoned about in the mind of player i . Using Principle 1 player i can now conclude that j will not update his beliefs if i sticks to s_3 . In the same vein, if it is actually in the interest of i to follow s_3 , then i will not update his beliefs as long as j follows s_2 .

Let us look at the slightly more complicated case. Behavioral type of player $i = \iota(\kappa)$ is the same as above but $\mathcal{K}$ has common attitude belief of level $M > 3$. We use the same logic as above, only now we should start from the beliefs of level M . In player i 's mind, at the belief of level M the utilities $u_{M,T}$ and $u_{M+1,T}$ are commonly believed in by both players. Thus, the player with (believed) attitude θ_{M-1} would expect that the game unfolds according to the strategies $(s_M, s_{M+1}) \in S$ determined by the Backward Induction procedure on the game with payoffs transformed by $u_{M,T}$ and $u_{M+1,T}$. Now however, this player, whose believed utility at this level of beliefs is $u_{M-1,T}$, will not stick to the strategy s_{M+1} that he thinks other player expects him to play, but will rather try to maximize his believed utility $u_{M-1,T}$ by playing best response to s_M . Notice that this behavior is *expected by both players* since we are talking about higher levels of beliefs inside both players' minds (from the perspective of player i). Thus we can define $s_{M-1} = BR(s_M)$. In similar vein we can go recursively closer to the beginning of the arborescence and define $s_k = BR(s_{k+1})$. This procedure stops when we reach node 2 in $\mathcal{K}$: the first occurrence of belief about player j . Thus player i in whose mind this reasoning takes place thinks that the other player is going to play the strategy s_2 and expects him to play according to s_3 .

Now we can generalize this reasoning procedure to any behavioral type $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$. For any linear path $\varepsilon = (\kappa, k_1, k_2, \dots) \in H_{\mathcal{K}}$ let $cb(\varepsilon) \in \varepsilon$ refer to a node after which the players have common attitude beliefs. For a linear path with common attitude belief level l $cb(\varepsilon)$ points at the $(l + 1)$ th node on the path.

Let $I_1 := \{n = cb(\varepsilon) : \varepsilon \in H_{\mathcal{K}}\}$ be the set of all nodes after which players have common attitude beliefs. Fix any node $v \in I_1$ and let $i = \iota(v)$ and $j = -i$. Player i can find a pair of commonly believed utilities $U_v^i = \mathcal{B}(A(v), A(\ell), A(v), A(\ell), \dots)$ and $U_v^j = \mathcal{B}(A(\ell), A(v), A(\ell), A(v), \dots)$ where ℓ is immediate successor of v . Thus

player j in the node which is an immediate predecessor of v will expect that i will use Backward Induction to reason about the game transformed by the utilities U_v^i and U_v^j . This gives rise to the strategy s_v^i of player i *expected in node* v . By following this procedure we can determine strategies s_v^i expected in node v for all $v \in I_1$.¹²

Now, there is a finite number of predecessors of nodes in I_1 . This guarantees that there is a non-empty set of nodes $I_2 \subseteq \mathcal{K}$ such that any $\mu \in I_2$ has immediate successors only from the set I_1 . For any such μ player $j = \iota(\mu)$ first finds his utility at μ . Let $B_{k,\ell} := \{n \in \mathcal{K} : kPnP\ell\}$ with $|B_{k,\ell}| = T$ be a linear path between k and ℓ of length T . Let $\rho_{\mu,\ell} = \prod_{n \in B_{\mu,\ell} \setminus \ell} p_n(C_n \cap B_{\mu,\ell})$ be the probability of the occurrence of $B_{\mu,\ell}$ according to j 's beliefs. Then the *expected social utility* of player j is $U_\mu^j := \sum_{B_{\mu,\ell}} \rho_{\mu,\ell} \mathcal{B}(A(\mu), \dots, A(\ell))$.

Now, player j recognizes that player i has expected strategies s_v^i in all $v \in C_\mu$. Player j forms an overall expected strategy $s_{C_\mu}^i = \bigcirc_{v \in C_\mu} (s_v^i, p_\mu(v))$ and finds pure best response $s_\mu^j = BR(s_{C_\mu}^i)$ to it using his expected social utility U_μ^j .¹³

Going backwards we can now define the set $I_3 \subseteq \mathcal{K}$ of nodes whose immediate successors are all in $I_1 \cup I_2$ and find best replies in these nodes using the exactly same procedure. Eventually after r steps we will arrive at the nodes C_κ whose immediate predecessor is the upper bound κ . Here the process stops. The procedure generates the strategies $s_k^{\iota(k)}$ for all $k \in I_1 \cup I_2 \cup I_3 \cup \dots \cup I_r$ that define what players expect to happen depending on the resolution of uncertainty.

2.3 Belief Updating

There is no requirement on the behavioral types that asks for the consistency of the beliefs. This implies that player i 's actual utility can be different from what player j expects it to be. Thus it might be the case that i wants to deviate from the equilibrium path that j expects to take place. In order to deviate optimally i has to predict what j will think after a move that is unexpected in j 's mind. In addition, player i naturally takes into account that player j can perform Bayesian updating of his beliefs as the game unfolds.

As before we consider the behavioral type $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$ of player $i = \iota(\kappa)$. We can use the procedure from the previous section to construct the strategies that constitute the expected play in the game depending on the resolution of uncertainty.

¹²For simplicity it is assumed that Backward Induction generates unique strategies for all players.

¹³Symbol $\bigcirc$ denotes the mixing of the collection of strategies s_v^i into the behavioral strategy with specified probabilities $p_\mu(v)$.

For player i relevant strategies of player $j = -i$ that i expects j to play are s_k^j for all $k \in C_\kappa$. Player j at some node $k \in C_\kappa$ expects i to play strategies s_n^i attached to the nodes $n \in C_k$ on the arborescence. Both players have uncertainty over which of these strategies will be actually played by the opponent. Consider the example on Figure 3:

Figure 3: Expected strategies on the belief arborescence.

Here $C_\kappa = \{k_1, k_2\}$, $C_{k_1} = \{k_3, k_4\}$ and $C_{k_2} = \{k_5, k_6\}$. Depending on the resolution of the uncertainty player i expects j to play according to $s_{k_1}^j$ or $s_{k_2}^j$ and player j at the node k_1 expects player i to play according to either $s_{k_3}^i$ or $s_{k_4}^i$ etc.

Now suppose that player i moves first in the game in the node $g \in G^i$. Then i should consider how player j would update the arborescence depending on i 's choice of the first action. Since at this point player i knows only probabilistically which behavioral type (starting at different nodes in C_κ) of player j he is facing, player i should consider updates for all different beliefs of player j . Fix any node $k \in C_\kappa$ (in our example above, say, $k = k_1$) and suppose i decides to choose some action $a \in A_g$. Then three things can happen:

1. For player j in k all strategies s_n^i with $n \in C_k$ prescribe that i will take action a (in the example: both $s_{k_3}^i$ and $s_{k_4}^i$ say that i chooses a);
2. For player j in k only some strategies s_n^i with $n \in C_k$ prescribe that i will take action a (in the example, say, only $s_{k_3}^i$ says that i chooses a);
3. For player j in k none of the strategies s_n^i with $n \in C_k$ prescribe that i will take action a (in the example neither $s_{k_3}^i$ nor $s_{k_4}^i$ say that i chooses a).

In the first case since nothing unexpected has happened and no new information was obtained player j will not update any beliefs (according to the Principle 1). So the branch of the arborescence following k will stay as it is.

In the second case i 's choosing action a delivers some information. Player j in node k can do Bayesian updating on his arborescence by cutting all the branches that have strategies s_n^i that prescribe something else from action a (in the example, behavioral types starting from k_4). Of course, player j has to redefine the probability function p_k according to the Bayes rule. This procedure gives new updated arborescence which is consistent with observed choice of a of player i .

The third case is the most interesting of all. Here the choice of action a is inconsistent with anything that player j expects to happen (in the mind of player i). However, as was mentioned above, player i might still be interested in following such unexpected course of action because his true social utility might be different from what player j thinks it is. In this case player i should predict what the beliefs of j will be after such an unexpected choice. Or in other words, how would player j rationalize i 's action a ?

In our model unexpected actions can be considered inside the “world” of belief dependent utilities that exist outside the game. When player j sees some action a which is not consistent with any of his beliefs he can try to “rationalize” action a by looking at sequences of reasoning about the behavioral types that could have generated it. Player j constructs additional behavioral (sub-)types; follows the strategic reasoning about the expected play given these types; chooses those that prescribe the choice of action a and update beliefs accordingly.

In general this procedure can be very complicated since the beliefs might get updated later in the game; player j might take into consideration doing something unexpected himself etc. That is why we make simplifying assumptions about how player j rationalizes the unexpected move of player i .

Assumption 2. *When player j observes an unexpected move by player i he only tries to rationalize it by considering linear paths with common attitude belief of level no more than some fixed number $R \geq 1$.*

Assumption 2 says that player j is using only the simplest linear arborescences to rationalize the behavior of player i . In particular, j does not consider the possibility that him or i can be uncertain about the attitudes somewhere along the arborescence. This assumption is made mostly to keep the sequence of reasoning *finite* as the goal

of this paper is to come up with the *finite algorithm* that allows to calculate the optimal play in the game with belief dependent utilities. If we assumed that player j rationalizes the unexpected move of i by any general arborescence with uncertainty the algorithm of finding the optimal play can become infinite.

Now, suppose player i takes an unexpected action a such that in node $k \in C_\kappa$ of the arborescence player j has to update beliefs according to the case 3. We assume player j in this situation goes through the following process of updating. He takes the set of all possible paths with common attitude belief of level d where $1 \leq d \leq R$ for all d . Then he finds all paths in this set that generate the expected play where player i chooses a (as described in Section 2.2). And finally he updates his arborescence by appending all the paths found in the previous step to it.

We make an additional assumption. Player j puts equal probabilities on all the new discovered paths that rationalize the move of player i . Moreover, to make the model more flexible we assume that player j does not discard his prior beliefs about i (those inconsistent with action a) but rather multiplies their probabilities of occurrence by some number $\alpha \in [0, 1]$. Thus the probability put on all new paths sums up to $1 - \alpha$.

There might be several reasons to introduce the parameter α into the model. When α is small or zero we are in the case when player j really starts believing something very different about player i than what he thought before. However, this creates the situation when the choices of player i become heavily influenced by the exact mechanism of the rationalization assumed above which might be undesirable. When α is big player j does not “trust” that much to his new found rationalizations, but still puts considerable weight on his previous beliefs. In this situation player i ’s decisions might be not much dependent on the assumption of rationalization mechanism.

Now we are finally ready to specify how the arborescence changes depending on any choice of player i . Suppose at *any* node $g \in G^i$ in the game player i considers taking action $a \in A_g$ and the current behavioral type is $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$. Player i first calculates all the strategies s_k^j for all $k \in C_\kappa$ and all the strategies s_n^i for all $n \in C_k$ for all k following the procedure of Section 2.2 which is applied to the subgame starting at g . These strategies point to the expected plays of player j and i in node g (depending on the resolution of uncertainty). Then player i constructs the new *updated arborescence after a* according to the following steps:

1. Take any node $k \in C_\kappa$;
2. Update the branch of the arborescence following k using the procedure of the

one of the three cases above. The procedure depends on s_k^j and s_n^i for $n \in C_k$;

3. Repeat the two steps above until all nodes in C_κ are updated.

This generates the updated arborescence after action a in node g .

2.4 Reasoning Procedure

In this section we describe the reasoning procedure that player i with behavioral type $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$ uses to find optimal strategy in the game when all belief updates are taken into account. Let the other player be called j and consider any first move of player i in node g_0^i .¹⁴ In case player i is not a first mover in the game, he calculates the expected play given $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$ and Bayesian updates his beliefs to only those behavioral types of player j that expectedly take the action that leads to g_0^i . Call this updated beliefs $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle_{g_0^i} =: \langle g_0^i \rangle$ and assign $\langle g_0^i \rangle$ to node g_0^i . In case player i is the first mover in the game assign the original not updated behavioral type to node g_0^i .

Now the procedure can be described in the following steps.

1. At g_0^i consider any action $a \in A_{g_0^i}$ that leads to node g_0^j of player j .
2. At g_0^j update beliefs of player j in $\langle g_0^i \rangle$. Obtain $\langle g_0^j \rangle$ and assign it to node g_0^j . $\langle g_0^j \rangle$ generates some expected play by player j which in general depends on the resolution of uncertainty.¹⁵ Let $\sigma(g_0^j)$ denote the mixed action of j in g_0^j generated by expected play.
3. Choose any action a_0 in support of $\sigma(g_0^j)$ that leads to some node g_1^i of player i . Bayesian update $\langle g_0^j \rangle$ to leave only behavioral types of player j that expectedly choose a_0 . Obtain $\langle g_1^i \rangle$ and assign it to g_1^i .
4. Repeat steps 1 to 3 with node g_1^i and then with nodes $g_2^i, g_3^i, \dots, g_r^i$ where g_r^i is the last move of player i . Consider any action a_r in $A_{g_r^i}$. If a_r leads to the terminal node z then assign $\langle z \rangle := \langle g_r^i \rangle$ to node z and calculate the expected social utility at z :

$$U_z^i[\langle z \rangle] := \sum_{(k_\ell)_{\ell=1..T-1}} p((k_\ell)_{\ell=1..T-1}) \mathcal{B}(\kappa, (k_\ell)_{\ell=1..T-1}) (\pi_i(z), \pi_j(z))$$

¹⁴In case player i is not a first mover in the game he would have several “first” moves.

¹⁵Keep in mind that throughout the reasoning procedure the whole game is considered. So, expected play of player j at g_0^j is found by the expected play procedure on the whole game.

where the sum goes over all truncated paths $(\kappa, k_1, \dots, k_{T-1})$ of length T on the belief arborescence; $p((k_\ell)_{\ell=1..T-1}) = p_\kappa(k_1) \prod_{t=1..T-2} p_{k_t}(k_{t+1})$ is the probability of occurrence of the path; and $\mathcal{B}(\kappa, (k_\ell)_{\ell=1..T-1})(\pi_i(z), \pi_j(z))$ is the payoff transformed with social utility $\mathcal{B}(\kappa, (k_\ell)_{\ell=1..T-1})$.

If a_r leads to some node g_r^j of player j that leads to some terminal node z' then repeat steps 1 to 3 and obtain $\langle g_r^j \rangle$. In step 3 it is important to Bayesian update player i 's beliefs as z' is reached. Assign this updated belief $\langle z' \rangle$ to z' . After that calculate the expected social utility $U_{z'}^i[\langle z' \rangle]$ using the formula above.

5. Repeat steps 1 to 4 until all possibilities of considering various actions of player i and j are exhausted.

The reasoning procedure gives the set $Z_1 \subseteq Z$ of terminal nodes for which the social utilities were calculated together with the sets $G_1^i \subseteq G^i$ and $G_1^j \subseteq G^j$ of nodes that lead to terminal nodes in Z_1 . In addition, for all nodes $g \in G_1^j$ we know the belief of player i about the choices of player j given by $\sigma(g)$. The optimal strategy s_0 of player i can then be found as a best response to his belief about the strategy of player j which consists of $\sigma(g)$ for all $g \in G_1^j$ with the payoffs of player i given by expected social utilities calculated in nodes Z_1 .

2.5 Relation to Psychological Games

In this section we build an epistemic model which describes the game, the beliefs about the attitudes and the reasoning procedure in the framework of psychological games of Battigalli and Dufwenberg (2009). First we need to understand how to relate any belief arborescence $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle$ to the conditional belief hierarchies defined over the strategies in the game.¹⁶ Consider a subset $\overline{\mathcal{K}}$ of the arborescence $\mathcal{K}$ that for each linear path with common attitude belief of level l contains only first $l + 1$ levels of beliefs. By the definition of common attitude belief the rest of the attitudes on the linear path are just the repetitions of the last l th and $(l + 1)$ th elements. By Assumption 1 $\overline{\mathcal{K}}$ is finite. Suppose as well that $\overline{\mathcal{K}}$ does not contain the upper bound of the arborescence κ . The function ι partitions $\overline{\mathcal{K}}$ into subsets of nodes belonging to each player in the game. Let T_1 and T_2 be the sets of nodes for players 1 and 2 so that $\overline{\mathcal{K}} = T_1 \cup T_2$. These sets represent the types of the two players.

¹⁶For the intuition see the example in the end of Section 2.1.

Let H denote the set of all histories in the game.¹⁷ Consider conditional beliefs for all types in T_i described by the function $b_i : T_i \times H \rightarrow \Delta(T_j)$. Somewhat abusing notation we can define $b_i(t_i, h) := p_{t_i}$ where p_{t_i} is the function that defines probabilities on the nodes of the arborescence following t_i (thus, distribution over T_j). For any type $t_{i,l+1} \in T_i$ without arborescence successors in $\bar{\mathcal{K}}$ (for each linear path with common attitude belief level l these are the $(l+1)$ th nodes) let $b_i(t_{i,l+1}, h)$ point at the type $t_{j,l}$ which is its immediate predecessor. Thus, after any history, types $t_{j,l}$ and $t_{i,l+1}$ have common certain belief in each other. Notice, in addition, that each type in T_i has the same beliefs after all possible histories in the game.¹⁸

Consider some functions $f_i : T_i \rightarrow S_i$ that assign a strategy to each type. The utility for each player can then be captured by $V_i : Z \times T_i \rightarrow \mathbb{R}$. Since utilities in our model are defined by T levels of beliefs we can set $V_i(z, t_i) := U_{t_i}^i(\pi_i(z), \pi_j(z))$ where $U_{t_i}^i$ are as determined in the reasoning about expected play in Section 2.2. Notice that V_i can be viewed as the psychological utility defined in Battigalli and Dufwenberg (2009). For each type t_i the epistemic model recursively generates an infinite hierarchy of conditional beliefs about strategies through functions b_i and f_i . Lastly, given the utilities V_i we can assign $f_i(t_i) = s_{t_i}^i$ as defined in Section 2.2. The construction of $s_{t_i}^i$ guarantees that each type t_i best responds to his beliefs given the utility $V_i(\cdot, t_i)$ after any history in H .

Now we are ready to define the epistemic model that relates the general reasoning procedure from Section 2.4 to a psychological game. Let $\langle \mathcal{K}, P, \kappa, A, (p_k)_{k \in \mathcal{K}}, \iota \rangle =: \langle g_0 \rangle$ be the beliefs of player i before the game starts. Assign $\langle g_0 \rangle$ to the first node g_0 in the game. Given this, the reasoning procedure from Section 2.4 generates belief arborescence $\langle g \rangle$ for each node g in the subset $G_1 := G_1^i \cup G_1^j \cup Z_1 \subseteq G$. For any $g \in G_1$ let $T_i^{\langle g \rangle}$ and $T_j^{\langle g \rangle}$ be the types corresponding to $\langle g \rangle$ together with the functions $(b_m^{\langle g \rangle}, f_m^{\langle g \rangle}, V_m^{\langle g \rangle})_{m=1,2}$ as defined above.

Let

$$\begin{aligned} \mathcal{T}_i &:= \bigcup_{g \in G_1 \cup \{g_0\}} T_i^{\langle g \rangle} \cup \{t_0\} \\ \mathcal{T}_j &:= \bigcup_{g \in G_1 \cup \{g_0\}} T_j^{\langle g \rangle} \end{aligned}$$

¹⁷Since we are only looking at the games of perfect information, H coincides with the set of nodes G .

¹⁸Type that corresponds to κ will not share this property.

be the types of players i and j . For all types excluding t_0 define functions $(b_m, f_m, v_m)_{m=1,2}$ as $(b_m^{(g)}, f_m^{(g)}, v_m^{(g)})_{m=1,2}$ from corresponding $g \in G_1 \cup \{g_0\}$. It is left to define the beliefs, the strategies and the utilities for the special type t_0 which represents player i at the upper bound κ of the arborescence. For each node $g \in G_1 \cup \{g_0\}$ there corresponds a history $h_g \in H$. Again slightly abusing notation let $b_i(t_0, h_g) := p_\kappa^{(g)}$. Here $p_\kappa^{(g)}$ is the function defining probabilities over the immediate successors of κ on the arborescence $\langle g \rangle$. Thus, at history h_g type t_0 holds beliefs only about player j 's types from arborescence $\langle g \rangle$. Notice that type t_0 holds different beliefs at any history corresponding to the nodes in $G_1 \cup \{g_0\}$. These beliefs reflect player i 's expectations about how beliefs after each move are updated. For the nodes $\bar{g} \in G \setminus G_1$ that are not allowed by any strategy $f_j^{(g)}$ for $g \in G_1 \cup \{g_0\}$ set $b_i(t_0, h_{\bar{g}}) := b_i(t_0, h_{g_0})$.¹⁹ For the nodes $Z_1 \subseteq G_1$ let $v_i(z, t_0) := U_z^i[\langle z \rangle]$ as defined in Section 2.4. For the rest of the terminal nodes $z' \in Z \setminus Z_1$ define $v_i(z', t_0) := U_\kappa^i(\pi_i(z'), \pi_j(z'))$ where U_κ^i is the expected social utility of player i at the upper bound κ of the arborescence $\langle g_0 \rangle$ (see Section 2.2). Finally, let $f_i(t_0) = s_0$ where s_0 is the best response of player i given all the belief updates as described in the end of Section 2.4.

To summarize, the epistemic model $(\mathcal{T}_m, b_m, f_m, v_m)_{m=1,2}$ represents the reasoning procedure that was described in the previous sections. Notice that by construction of the beliefs, strategies and utilities each type of either player best responds to his beliefs after any history. Thus, we have obtained an epistemic model that rationalizes the reasoning procedure.

3 Examples

In the examples below we use the following definitions of belief dependent utilities unless specified otherwise: $\Theta = \{S, N\}$; $\mathcal{P} = \{u_S, u_N\}$ where $u_S(x_1, x_2) = x_1$ is “selfish” utility function and $u_N(x_1, x_2) = x_1 - \alpha \max\{x_2 - x_1, 0\} - \beta \max\{x_1 - x_2, 0\}$

¹⁹Since these histories are not allowed by any strategy of player j that player i can possibly believe in, his beliefs at these nodes are irrelevant.

is “nice” utility function;²⁰ $T = 2$ and

$$\begin{aligned}\mathcal{B}(S, S) &= u_S \\ \mathcal{B}(S, N) &= u_S \\ \mathcal{B}(N, S) &= u_S \\ \mathcal{B}(N, N) &= u_N.\end{aligned}$$

3.1 Sequential Prisoner’s Dilemma

One of the main advantages of the Battigalli and Dufwenberg (2009) framework is the possibility to update beliefs as the game unfolds. We would like to provide some evidence that in the experimental lab subjects do indeed update their beliefs in a consistent way which influences their behavior. We use sequential Prisoner’s Dilemma and its slight modification in Figure 4. The game on the right is different from the usual PD in only one payoff of player 1 marked in red. The important difference between the two games though is that in the modified Prisoner’s Dilemma player 1 with utility u_S has dominant action L .

Figure 4: Sequential Prisoner’s Dilemma (left) and its modification (right).

We are interested in the inferences that player 2 does in the two games upon observing action R of player 1. Consider the belief arborescence of player 2 on Figure 5. Here player 2 has attitude S and believes with probabilities $\frac{1}{4}$ that player 1 can be one of the four shown behavioral types.

Assume that $u_N(4, 4) = 4 > u_N(5, 0)$ and $u_N(1, 1) = 1 > u_N(0, 5)$. Then upon observing action R of player 1 player 2 will update his beliefs. In the modified version

²⁰Here inequality aversion is not a crucial assumption. It is possible to assume u_N to be altruistic utility with the same results.

Figure 5: Belief arborescence of player 2 in sequential Prisoner's Dilemma.

of sequential PD action R of player 1 can mean only one thing: player 1 has behavioral type $N - N - N - N - \dots$ with corresponding utility u_N and expects player 2 to go $R2$ (Backward Induction outcome given common belief that both players have utility u_N). In other three cases player 1 has utility u_S and thus should prefer the dominant action L . Therefore, player 2 in modified PD updates his beliefs after R to $S - N - N - N - N - \dots$. Therefore, the only way for player 2 to rationalize action R in modified PD is to believe that player 1 holds belief $N - N - N - N - \dots$.

In the standard PD the situation is different. Here player 1 can have selfish preferences and still go R as long as he believes that player 2 will go $R2$. This happens for example if player 1 believes $S - N - N - N - N - \dots$, which means that he is himself selfish, but thinks that player 2 believes that both players commonly believe in each other having attitude N . Therefore, after seeing R player 2 updates his beliefs to the following:

Notice that now player 2 does not have a sure belief about what utility player 1 actually has.

In the experiment subjects played two games one after the other. In treatment 1 the first game was standard sequential PD followed by a Trust game. In treatment 2 the first game was modified PD followed by the same Trust game (see Figure 6).

Subjects who played as player 1 in PD then became Responders in the Trust game. Subjects who played as player 2 in PD became Proposers in the Trust game. In the beginning of the experiment subjects did not know that after PD they will face Trust game. They only knew that they will play several games. After the first game subjects were rematched. Before playing the Trust game the Proposers received the information about the choices of the Responders in the first game. For example, a Proposer in

Figure 6: Experimental design.

Trust game could have been told: “Now you are playing Game 2 with a person who chose R in the previous game.”²¹

The graph on Figure 7 shows the amount proposed in the Trust game depending on the treatment and the information provided to the Proposers.

In a qualitative accordance with the prediction of the model proposals are twice higher after observing that the Responder played R in a modified PD than in other three conditions. In particular, observing the action that the Responder chose in the standard PD does not influence the amount proposed. In our framework the explanation is that action R in modified PD reveals the type and beliefs of the Responder whereas in standard PD it does not.

It is worthwhile to look at this result in the light of other models that propose explanations for cooperative behavior in games. Levine (1998) considers standard incomplete information framework in which players have private information about their social utility function (say, u_S or u_N) and share common prior about the distribution of the types in the population. According to this model in modified PD players with utility u_N would go R if there are sufficiently many people with utility u_N in the population (so that the probability of them going $R2$ is high enough). Assuming this, only separating equilibrium exists in modified PD. Following the same logic standard PD should have only pooling equilibrium in which all types go R . Given

²¹The rules of the Trust game (Game 2) were described in the instructions beforehand.

Figure 7: Proposals in the Trust game after seeing the action of the Responder in PD and modified PD.

these observations it is in principle conceivable that players update their prior after observing the action of the Responder before the Trust game and thus make higher proposals after seeing *R* in modified PD. However, these considerations are beyond the framework of the incomplete information games and are not explicitly modeled. In our setup players are able to reason about the meaning of information they receive and incorporate it into their future decisions.

Another two models that tackle the same issue are Dufwenberg and Kirchsteiger (2004) and Falk and Fischbacher (2006). Both consider reciprocal behavior in dynamic games using the approach that allows for some form of belief updating. In particular, after a player chooses some action in the game the opponent calculates how “kind” that action was towards him by comparing the payoffs that can still happen in the future and the payoffs that are no longer available after the current action. These models are capable of predicting cooperative outcomes in the PD, however they can say nothing about the effect that the revelation of information between the games has on the behavior in the Trust game. This is a consequence of the absence of uncertainty about the opponents type in these models.

3.2 Repeated Sequential Prisoner's Dilemma

Many experimental studies show that people cooperate in finitely repeated social dilemmas (e.g. Clark and Sefton (2001)).²² This results are usually hard to explain in standard game theoretic settings given simple recursive logic that, as long as selfish players are involved, obviously profitable defection in the last period of the game essentially makes finitely repeated social dilemma one period shorter thus asking for defection in the period one to last etc. We propose a model in which selfish players maintain cooperation in finitely repeated sequential Prisoner's Dilemma. The construction crucially depends on two assumptions: 1) players update their beliefs about the behavioral types of other player as the game unfolds; 2) players have *mismatched beliefs*. In particular, consider two players who both have the following belief arborescence before the game starts:

$$S - N - N - N - N - N - N - \dots$$

This means that both players have selfish utility u_S , but believe that the other player maintains common belief that both players have utility u_N .

Consider finitely repeated stage game on the left of Figure 4. We assume that $u_N(4, 4) = 4 > u_N(5, 0)$ and $u_N(1, 1) = 1 > u_N(0, 5)$. Under common belief that both players have utility u_N Backward Induction generates the unique strategies: player 1 goes R in all periods; player 2 goes $(L1, R2)$ in all periods. In finitely repeated game with D periods this gives the payoff of $4D$ to each player. As long as both players play the aforementioned strategies no belief updating is taking place. However, players might have the incentive to deviate. The important thing to notice is that after any deviation players will update their beliefs about the beliefs of the opponent.

If player 1 deviates to L he knows that player 2 should then think that player 1's utility is actually u_S as prescribed by the updated belief arborescence that now includes a variety of linear paths with combinations of N 's and S 's excluding $N - N - N - \dots$ ²³ Given this updated belief player 1 as a player whose actual utility is u_S will prefer to go L in all remaining periods expecting player 2 to choose $(L1, L2)$. This gives player 1 a payoff of 1 in each period after deviation which is strictly less than without deviation. For player 2 the situation is similar. After moving $L2$ instead

²²Cooperation normally decreases over time.

²³Here we assume a model with $\alpha = 0$ (see Section 2.3). This means that when unexpected action happens players rationalize it with some beliefs and forget their original belief. It is interesting to check whether high α can generate the "return to cooperation" behavior as observed in the lab.

of R_2 player 2 updates his beliefs about the beliefs of player 1 who now thinks that player 2's utility is u_S . Player 2 gets one time payoff of 5 and then receives 1 in all consecutive periods as after the deviation player 1 starts to optimally go L . Deviation is not profitable for player 2 if

$$5 + (D - 1) < 4D \Rightarrow D > \frac{4}{3}.$$

Thus for any repeated game with more than 1 period we should expect to observe strategies R in all periods for player 1 and (L_1, R_2) in all periods but the last and (L_1, L_2) in the last period for player 2. Player 2 does want to defect in the last period as he still has selfish utility u_S .

References

- BATTIGALLI, P., AND M. DUFWENBERG (2009): "Dynamic psychological games," *Journal of Economic Theory*, 144, 1–35.
- CHARLES, K. K., E. HURST, AND N. ROUSSANOV (2009): "Conspicuous consumption and race," *Quarterly Journal of Economics*, 124(2), 425–467.
- CLARK, K., AND M. SEFTON (2001): "The Sequential Prisoner's Dilemma: Evidence on Reciprocation," *The Economic Journal*, 111(468), 51–68.
- COX, J. C., D. FRIEDMAN, AND S. GJERSTAD (2007): "A tractable model of reciprocity and fairness," *Games and Economic Behavior*, 59, 17–45.
- DUFWENBERG, M., AND G. KIRCHSTEIGER (2004): "A theory of sequential reciprocity," *Games and Economic Behavior*, 47, 268–298.
- FALK, A., E. FEHR, AND U. FISCHBACHER (2008): "Testing theories of fairness – Intentions matter," *Games and Economic Behavior*, 62, 287–303.
- FALK, A., AND U. FISCHBACHER (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54, 293–315.
- FEHR, E., AND S. GÄCHTER (2000): "Fairness and Retaliation: The Economics of Reciprocity," *Journal of Economic Perspectives*, 14(3), 159–181.
- FISCHBACHER, U., AND S. GÄCHTER (2010): "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments," *American Economic Review*, 100(1), 541–556.
- FISKE, A. P. (1990): "Relativity within Moose ("Mossi") Culture: Four Incommensurable Models for Social Relationships," *Ethos*, 18(2), 180–204, Moral Relativism.
- LEVINE, D. K. (1998): "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics*, 1(3), 593–622.
- MCCABE, K. A., M. L. RIGDON, AND V. L. SMITH (2003): "Positive reciprocity and intentions in trust games," *Journal of Economic Behavior and Organization*, 52, 267–275.