



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Incomplete-Information Models of Guilt Aversion in the Trust Game

Giuseppe Attanasi, Pierpaolo Battigalli, Elena Manzoni

To cite this article:

Giuseppe Attanasi, Pierpaolo Battigalli, Elena Manzoni (2016) Incomplete-Information Models of Guilt Aversion in the Trust Game. *Management Science* 62(3):648-667. <http://dx.doi.org/10.1287/mnsc.2015.2154>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Incomplete-Information Models of Guilt Aversion in the Trust Game

Giuseppe Attanasi

Bureau of Economic Theory and Applications, CNRS, and University of Strasbourg, 67000 Strasbourg, France,
attanasi@unistra.fr

Pierpaolo Battigalli

Department of Decision Sciences and Innocenzo Gasparini Institute for Economic Research, Bocconi University,
20136 Milan, Italy, pierpaolo.battigalli@unibocconi.it

Elena Manzoni

Department of Economics, Management and Statistics, University of Milano-Bicocca, 20126 Milano, Italy,
elena.manzoni@unimib.it

In the theory of psychological games it is assumed that players' preferences on material consequences depend on endogenous beliefs. Most of the applications of this theoretical framework assume that the psychological utility functions representing such preferences are common knowledge. However, this is often unrealistic. In particular, it cannot be true in experimental games where players are subjects drawn at random from a population. Therefore, an incomplete-information methodology is needed. We take a first step in this direction, focusing on guilt aversion in the Trust Game. In our models, agents have heterogeneous belief hierarchies. We characterize equilibria where trust occurs with positive probability. Our analysis illustrates the incomplete-information approach to psychological games and can help to organize experimental results in the Trust Game.

Keywords: psychological games; Trust Game; guilt; incomplete information

History: Received July 4, 2013; accepted December 16, 2014, by Uri Gneezy, behavioral economics. Published online in *Articles in Advance* August 17, 2015.

1. Introduction

The Trust Game is a stylized social dilemma whereby player A takes a costly action (investment) that generates a social return, and player B decides how to distribute the proceeds between himself and player A . Experimental work on the Trust Game has shown systematic and significant departures from the standard equilibrium prediction implied by the assumption of common knowledge of selfish preferences (see Berg et al. 1995, Buskens and Raub 2013, §III.A of the survey by Cooper and Kagel 2013, and the references therein). Given the simplicity of this game, such deviations are hard to explain as the result of bounded rationality. Charness and Dufwenberg (2006) provide support for the hypothesis that the behavior of most subjects in the second-mover role (player B) is affected by aversion to letting down the first mover (player A) relative to A 's expectations, as in Dufwenberg's (2002) model of marital investment. This is an instance of the "simple guilt" model of belief-dependent preferences of Battigalli and Dufwenberg (2007). Recent experimental work corroborates this hypothesis (e.g.,

Reuben et al. 2009, Bellemare et al. 2011, Chang et al. 2011, Attanasi et al. 2013).¹

Of course, when subjects' preferences differ from the simple benchmark of selfish expected payoff maximization, the assumption that such preferences are common knowledge is farfetched. Therefore, it should be assumed that the game played in the laboratory is one with incomplete information, although the rules of the game (e.g., who plays when, information about previous moves, and material payoffs at terminal nodes) are made common knowledge in the experiment. This is consistent with the high heterogeneity of behavior and beliefs found in most experiments on other-regarding preferences (see Cooper and Kagel 2013). Our goal is to understand how such a game is played with incomplete information about guilt sensitivity.

¹ See also Dufwenberg and Gneezy (2000) and Guerra and Zizzo (2004). Charness and Dufwenberg (2011) find support for the "guilt-from-blame" model of Battigalli and Dufwenberg (2007), which assumes that i anticipates guilt if she thinks that j is going to blame i for letting him down. Vanberg (2008) and Ellingsen et al. (2010) question the guilt-aversion interpretation of prosocial behavior in the Trust Game.

We analyze the Bayesian equilibria (Harsanyi 1967, 1968a, b) of two incomplete-information models of guilt aversion in the Trust Minigame, a binary-choice version of the Trust Game similar to the one analyzed by Charness and Dufwenberg (2006).² A key feature of both models is that agents playing in a given role hold heterogeneous beliefs about the type of the coplayer, which implies heterogeneous first- and second-order beliefs about actions. In the first model it is common knowledge that player *A*, the “truster,” is selfish and that only player *B*, the “trustee,” can feel guilt. In the second model, instead, guilt sensitivity and beliefs about it do not depend on the role played in the game (player *A* or player *B*). The first model is more tractable and it may be appropriate for situations where the players come from different populations, e.g., when player *A* is a firm and player *B* is a worker. The second model may be more appropriate for situations where the subjects playing in roles *A* and *B* are drawn from the same population, as in most experiments.

However, even when players are drawn from the same population, it is not implausible to assume that sensitivity to guilt is triggered only when playing in role *B*. This assumption resonates with (i) the evolutionary psychology of emotions (e.g., Haselton and Ketelaar 2006), which suggests that, when a single emotion (guilt) operates in a variety of different domains, its effects are moderated by contextual cues; and with (ii) the conceptual act theory of emotion (e.g., Barrett 2006), which posits that people experience an (in our case, anticipated) emotion by categorizing an instance of affective feeling (anticipated disappointment of the other); with this, it is plausible that the role played in the interaction is part of the categorization process. Finally, the different responses to oxytocin of trusters and trustees (e.g., Kosfeld et al. 2005, Zak et al. 2005), and findings in the animal literature on the reactivity of oxytocin to social cues (e.g., Carter and Keverne 2002) provide some indirect evidence supporting the role-dependent model of guilt aversion.³

² We coined the name “Trust Minigame” after the “Ultimatum Minigame” of Binmore et al. (1995), a binary-choice version of the Ultimatum Game.

³ On one hand, higher levels of oxytocin in trustees are correlated (between subjects) with a higher investment by trusters and higher sharing (Zak et al. 2005, Zak 2008). According to the guilt-aversion model, oxytocin can thus be interpreted as the transmitter from second-order beliefs (expected disappointment from not giving) to the prosocial action. However, in the case of trusters, higher levels of oxytocin are not correlated (between subjects) with a higher investment. Yet, Kosfeld et al. (2005) show that an exogenously induced increase in the oxytocin level increases investment, which suggests that it is the difference between baseline and actual levels of oxytocin that affects prosocial behavior. We submit that the role of truster is a social cue that shuts down the link between

Our approach finds its intellectual home in the theory of psychological games, i.e., the analysis of games with belief-dependent preferences (Geanakoplos et al. 1989, Battigalli and Dufwenberg 2009; see also the introductory surveys by Dufwenberg 2008 and Attanasi and Nagel 2008). To our knowledge, this is the first paper offering a fully-fledged Bayesian equilibrium analysis of guilt aversion.⁴

Our paper is related to Attanasi et al. (2013), who analyze experimentally the belief-dependent preferences, behavior, and beliefs of subjects in the Trust Minigame. They show that making the elicited belief-dependent preferences common knowledge between the subjects of each matched pair significantly affects behavior and beliefs. This can be interpreted as a comparison between a psychological game with incomplete information (control) and a psychological game with complete information (treatment). The theoretical comparison between treatment and control draws on the analysis of our paper, which therefore helps to organize the data of their experiment.

Also, under mild assumptions about the empirical distribution of types, our analysis implies the positive correlation between the second-order beliefs and prosocial action of *B* subjects found by Charness and Dufwenberg (2006). In §6, we comment extensively on such a correlation and, more generally, on the relevance of our models for experiments.


As argued above, the assumption that interacting individuals have belief-dependent social preferences naturally leads to an incomplete-information analysis. Therefore, we hope that our paper may have a pedagogical value for applied theorists and experimental economists who are interested in using psychological game theory to analyze social dilemmas. Indeed, we present the more abstract methodological material on Bayesian equilibrium and psychological games in §3 so that it can be easily extended and applied to different games.

The rest of the paper is structured as follows. Section 2 introduces the Trust Minigame with guilt aversion. Section 3 provides the methodology to analyze psychological Bayesian games, with a focus on the Trust Minigame with unknown guilt aversion.

second-order beliefs and oxytocin levels (see Carter and Keverne 2002); thus, oxytocin differences between trusters mainly reflect differences in baseline levels, not in second-order beliefs.

⁴ Some papers analyze incomplete-information models of games with belief-dependent preferences; see, for example, Caplin and Leahy (2004), Ong (2011) and Tadelis (2011). Unlike ours, none of these models features heterogeneous beliefs, which are instead allowed for by Battigalli et al. (2013). The latter paper analyzes the cheap-talk game of Gneezy’s (2005) experiment under the assumption that the sender is affected by an unknown sensitivity to guilt. This analysis is based on two rounds of elimination of nonbest replies under mild assumptions about heterogeneous beliefs.

Figure 1 The Trust Minigame Form with Material Payoffs



Section 4 puts forward and analyzes a model with role-dependent guilt, where player A is known to be selfish. Section 5 puts forward and analyzes a model with role-independent guilt. Section 6 concludes with a discussion of the empirical implications of our analysis and of the situations where it can be usefully applied. Formal proofs are collected in the appendix.

2. Guilt Aversion in the Trust Minigame

We analyze models of the Trust Minigame where players have different sensitivities to guilt feelings and incomplete information about the guilt sensitivity of the coplayer. All of the models we consider are based on the game form with material payoffs depicted in Figure 1.

In the analysis of this game form, we denote players' strategies as follows.


Strategy	Notation
<i>In</i>	I
<i>Out</i>	O
<i>Share if In</i>	S
<i>Keep if In</i>	K

To investigate the effects of guilt feelings on behavior, we need to consider the players' first- and second-order beliefs about strategies. We denote with α_i player i 's first-order beliefs, and with β_i the second-order beliefs. Specifically, we use the notation described in the following table.⁵

Belief	Notation	Definition
Ann's initial <i>first-order belief</i>	α_A	$\mathbb{P}_A[S]$
Bob's initial <i>first-order belief</i>	α_B	$\mathbb{P}_B[I]$
A feature of Bob's initial <i>second-order belief</i>	β_B^\emptyset	$\mathbb{E}_B[\alpha_A]$
A feature of Bob's <i>conditional second-order belief</i>	β_B^I	$\mathbb{E}_B[\alpha_A I]$

⁵ We use bold symbols to denote random variables. Because B does not know α_A , this number is a random variable from B 's point of view, and its expectation is $\mathbb{E}_B[\alpha_A]$. Similarly, we write $\mathbb{E}_A[\mathbf{m}_A]$ for the expected material payoff of A .

Figure 2 The Trust Minigame with Psychological Utilities



Note that we distinguish between initial and conditional second-order beliefs of Bob, and we refer to the features of such beliefs that are relevant in our analysis. Indeed, we assume below that Bob's choice depends on his expectation of Ann's disappointment if he Keeps, which can be written as a function of the expected value of Ann's first-order belief. The second-order beliefs of Ann will be introduced later as needed.

According to the model of simple guilt (Battigalli and Dufwenberg 2007), player i suffers from guilt to the extent that she believes that she is letting the coplayer $-i$ down. In particular, player i has belief-dependent preferences over material payoff distributions represented by the following psychological utility function:

$$u_i = m_i - \theta_i \max\{0, \mathbb{E}_{-i}[\mathbf{m}_{-i}] - m_{-i}\}, \quad (1)$$

where m_i is the material payoff of i , $\theta_i \geq 0$ is the guilt sensitivity of player i , and $\max\{0, \mathbb{E}_{-i}[\mathbf{m}_{-i}] - m_{-i}\}$ measures the extent of the coplayer's disappointment given the coplayer's subjective beliefs.

We first assume that guilt sensitivity is role-dependent: only the second mover can be affected by guilt ($\theta_A = 0, \theta_B \geq 0$), and this is common knowledge. Ignoring players' beliefs about parameters, we can represent the strategic situation by the parametrized psychological game depicted in Figure 2.

Indeed, Ann can only be disappointed after terminal history (I, K) , in which case the extent of her disappointment is

$$\begin{aligned} & \max\{0, \mathbb{E}_A[\mathbf{m}_A] - m_A(I, K)\} \\ & = 2 \cdot \alpha_A + 0 \cdot (1 - \alpha_A) - 0 = 2\alpha_A, \end{aligned}$$

where $m_i(z)$ denotes the material payoff of i at terminal history $z \in \{O, (I, K), (I, S)\}$. Thus, the psychological utility at $z = (I, K)$ of Bob (expressed as a function of Ann's first-order belief α_A) is

$$\begin{aligned} u_B(I, K, \alpha_A) & = m_B(I, K) - \theta_B \max\{0, \mathbb{E}_A[\mathbf{m}_A] - m_A(I, K)\} \\ & = 4 - 2\theta_B\alpha_A. \end{aligned}$$

Of course, when Bob evaluates his alternatives and chooses his optimal strategy, he compares the utility from choosing S with the expected psychological utility from choosing K , which depends on his second-order beliefs. As long as Bob initially assigns a strictly

positive probability to I ($\alpha_B = \mathbb{P}_B[I] > 0$), the comparison between strategies S and K can equivalently be made either ex ante, or conditional on I , because the difference between the ex ante expected utilities of S and K is proportional to the difference between the conditional expected utilities of S and K

$$\begin{aligned} \mathbb{E}_B^S[\mathbf{u}_B] - \mathbb{E}_B^K[\mathbf{u}_B] &= \mathbb{E}_B[u_B(I, S, \alpha_A) \\ &\quad - u_B(I, K, \alpha_A) \mid I] \cdot \mathbb{P}_B[I] \\ &= 2(\theta_B \beta_B^I - 1)\alpha_B. \end{aligned} \quad (2)$$

By definition, $0 \leq \beta_B^I \leq 1$; thus, in an equilibrium with $\alpha_B > 0$, Bob Shares if $\alpha_B > 0$, $\theta_B > 1$, and $\beta_B^I > 1/\theta_B$.

The assumption that only player B may be sensitive to guilt is removed in §5, where we analyze a model with role-independent guilt.

3. Methodology: Bayesian Psychological Games

We are going to model incomplete information about θ using the methodology first proposed by Harsanyi (1967, 1968a, b), suitably extended to psychological games (see also Battigalli and Dufwenberg 2009, §6.2). We define type structures that implicitly determine the possible hierarchies of the subjective beliefs of the players.

Although our methodology is fully standard from the abstract theory perspective, it is not widely used in applied theory. Therefore, it is useful to describe carefully the building blocks of our approach. The main concepts are illustrated by a leading example.

Here we present a note on terminology. We call “exogenous” a belief about an exogenous variable or a parameter: a belief about θ is an exogenous first-order belief, a joint belief about θ and exogenous first-order beliefs of the coplayer is an exogenous second-order belief, and so on. We call “endogenous” a belief about a variable that we try to explain, or predict, with the strategic analysis of the game. In particular, a belief about strategies is an endogenous first-order belief, a joint belief about strategies and endogenous first-order beliefs is an endogenous second-order belief, and so on. We also call “endogenous” a joint belief about exogenous and endogenous variables.

3.1. Type Structures

We consider situations where the psychological utility functions of players A and B are determined by parameters $\theta_A \in \Theta_A$, $\theta_B \in \Theta_B$ known to players A and B , respectively. Formally, the psychological utility of i is a parametrized function

$$u_i: \Theta_i \times Z \times \mathcal{H}_i \times \mathcal{H}_{-i} \rightarrow \mathbb{R},$$

where Z is the set of terminal histories (play paths) of the game, and \mathcal{H}_i (\mathcal{H}_{-i}) is a space of endogenous hierarchical beliefs of player i ($-i$).⁶ Because in our applications θ_i is the guilt sensitivity parameter of player i , we call θ_i the “guilt type.” When the parameter set Θ_i is a singleton, the guilt type of i is common knowledge. In models with role-dependent guilt sensitivity, we have $\Theta_A \neq \Theta_B$; in particular, we assume that Θ_A is a singleton, because player A is commonly known to be a self-ish expected material-payoff maximizer; i.e., $\Theta_A = \{0\}$.

The subjective exogenous beliefs of players A and B about each other’s private information and exogenous beliefs are implicitly represented by a type structure, i.e., a tuple

$$\mathcal{T} = \langle N = \{A, B\}, (\Theta_i, T_i, \mathfrak{D}_i: T_i \rightarrow \Theta_i, \tau_i: T_i \rightarrow \Delta(T_{-i}))_{i \in N} \rangle.$$

Elements of T_i are called “Harsanyi types” or simply “types.” A Harsanyi type specifies both the guilt type (more generally, the utility function) and the exogenous beliefs of player i .⁷ Note that we use bold symbols to denote functions interpreted as random variables, i.e., functions that depend on the state of the world (t_A, t_B) . Function \mathfrak{D}_i specifies the psychological utility (guilt sensitivity) of type t_i , and function τ_i determines the exogenous beliefs of t_i about the utility and beliefs of the coplayer $-i$. In particular, we explain below how each type t_i in the structure determines a whole hierarchy of exogenous beliefs for player i . Given a random variable $x_i: T_i \rightarrow X_i$, we denote events about x_i either directly as subsets of T_i , or according to a convention that is common in statistics. For example, both $\mathfrak{D}_i^{-1}([0, x])$ and $\mathfrak{D}_i \leq x$ denote the set $\{t_i: \mathfrak{D}_i(t_i) \leq x\}$, i.e., the event that the guilt type of i is at most x . We use whatever notation is more convenient and transparent in the given context.

It may be assumed without essential loss of generality that the set of types is a Cartesian product $T_i = \Theta_i \times \mathcal{E}_i$, so that a type is a pair (θ_i, e_i) , and that beliefs about the coplayer’s type are determined only by the second element e_i , also called “epistemic type.” In this case, function \mathfrak{D}_i is the projection map $(\theta_i, e_i) \mapsto \mathfrak{D}_i(\theta_i, e_i) = \theta_i$, and function τ_i depends only on e_i ; hence, it makes sense to write $\tau_i(e_i)$.

Once we append a type structure to the profile of parametrized utility functions, we obtain a Bayesian psychological game

$$\begin{aligned} \Gamma &= \langle N = \{A, B\}, (\Theta_i, u_i: \Theta_i \times Z \times \mathcal{H}_i \times \mathcal{H}_{-i} \rightarrow \mathbb{R}, T_i, \\ &\quad \mathfrak{D}_i: T_i \rightarrow \Theta_i, \tau_i: T_i \rightarrow \Delta(T_{-i}))_{i \in N} \rangle. \end{aligned}$$

⁶ See Geanakoplos et al. (1989) and Battigalli and Dufwenberg (2009). In the latter, \mathcal{H}_i is a space of hierarchical conditional beliefs.

⁷ All of our models satisfy the following technical assumptions: for each player i , T_i is a compact metric space, the set of Borel probability measures $\Delta(T_{-i})$ is endowed with the topology of weak convergence (hence, it is compact and metrizable), and the functions \mathfrak{D}_i, τ_i are continuous. This implies that the sets and functions we consider satisfy the necessary measurability requirements.

In this paper, we focus on Bayesian psychological games based on the Trust Minigame with guilt aversion, i.e., the game form of Figure 1 with parametrized utility functions given by Equation (1).

3.2. Higher-Order Exogenous Beliefs

The exogenous first-order belief of a type t_i is determined by the equation

$$\mathbf{p}_i^1(t_i)[E_{-i}^0] = \tau_i(t_i)[(\Theta_{-i})^{-1}(E_{-i}^0)] \\ (E_{-i}^0 \subseteq \Theta_{-i} \text{ Borel measurable}).$$

For example, $\mathbf{p}_A^1(t_A)[\{t_B: \Theta_B(t_B) > 2\}]$ is the subjective probability assigned by type t_A to the event that the guilt sensitivity of player B is more than 2. We can write this probability more compactly as $\mathbf{p}_A^1(t_A)[\Theta_B > 2]$.

With this, we obtain a map $(\Theta_i, \mathbf{p}_i^1): T_i \rightarrow \Theta_i \times \Delta(\Theta_{-i})$ for each $i \in \{A, B\}$. Then the exogenous second-order belief of a type t_i is determined by the equation

$$\mathbf{p}_i^2(t_i)[E_{-i}^1] = \tau_i(t_i)[(\Theta_{-i}, \mathbf{p}_{-i}^1)^{-1}(E_{-i}^1)] \\ (E_{-i}^1 \subseteq \Theta_{-i} \times \Delta(\Theta_i) \text{ Borel measurable}).$$

For example, $\mathbf{p}_B^2(t_B)[\{t_A: \mathbf{p}_A^1(t_A)[\Theta_B > 2] \geq \frac{1}{2}\}]$ is the subjective probability assigned by type t_B to the event that player A believes that $\Theta_B > 2$ is at least as likely as $\Theta_B \leq 2$.

Proceeding this way, we can associate a hierarchy of exogenous beliefs with each type. However, beliefs beyond the second order will not be used in the analysis below.

3.3. Leading Example: Exogenous Beliefs

We illustrate these abstract concepts with an example, which is (essentially) a special case of our model with role-dependent guilt. Suppose that it is common knowledge that player A is selfish, whereas player B can either have a low-guilt type $\theta^L < 1$, or a high-guilt type $\theta^H > 2$; therefore, $\Theta_A = \{0\}$ and $\Theta_B = \{\theta^L, \theta^H\}$. The exogenous beliefs of each player i are determined by the epistemic type $e_i \in \mathcal{E}_i$. Because player A has only one possible guilt type, the epistemic and Harsanyi types of A coincide, and we can ease notation, writing $e_A = t_A$. There is a continuum of epistemic types on both sides. Specifically, we let $e_A = t_A \in T_A$ parametrize the subjective probability assigned by A to the high-guilt type of B : $t_A = \mathbb{P}_{t_A}[\Theta_B = \theta^H]$.⁸ Therefore, we let $T_A = \mathcal{E}_A = [0, 1]$. Furthermore, we assume that the set of possible epistemic types of player B is $[0, 1]$ as well. This is just a convenient parametrization. Thus $T_B = \Theta_B \times \mathcal{E}_B = \{\theta^L, \theta^H\} \times [0, 1]$. Although the set of

epistemic types of player A is easily seen to be isomorphic to the set of exogenous first-order beliefs of player A , the meaning of the epistemic types of player B can only be understood by considering the belief maps and unraveling the higher-order beliefs corresponding to each type.

The belief maps have the following features. All types of player A believe that the guilt and epistemic types of player B are statistically independent; furthermore, they hold the same marginal belief about the epistemic type of player B given by a strictly positive density function $f: [0, 1] \rightarrow \mathbb{R}$, e.g., the uniform distribution $f(e_B) = 1$. On the other hand, different epistemic types of player B may hold different beliefs about the type of player A : the belief of each e_B is given by a strictly positive density function $f_{e_B}: [0, 1] \rightarrow \mathbb{R}$.⁹

To sum up, the belief maps $\tau_i: T_i \rightarrow \Delta(T_{-i})$ ($i \in \{A, B\}$) satisfy

$$\tau_A(t_A)[\Theta_B = \theta^H \cap \mathbf{e}_B \leq y] \\ = \tau_A(t_A)[\Theta_B = \theta^H] \cdot \tau_A(t_A)[\mathbf{e}_B \leq y] = t_A \int_0^y f(e_B) de_B,$$

and

$$\tau_B(\theta_B, e_B)[\mathbf{t}_A \leq x] = \int_0^x f_{e_B}(t_A) dt_A,$$

for all $t_A, e_B, x, y \in [0, 1]$, $\theta_B \in \{\theta^L, \theta^H\}$.

All of the above (including the belief map $e \mapsto f_e$) is assumed to be common knowledge. This gives the type structure \mathcal{T} and a Bayesian psychological game Γ based on the Trust Minigame with guilt aversion.

Exogenous hierarchies of beliefs are relatively simple. In particular, beliefs about the guilt type of player A and beliefs about such beliefs are trivial because player A is commonly known to be selfish, whereas the first- and second-order beliefs of, respectively, players A and B about the guilt type of player B are

$$\mathbf{p}_A^{1,H}(t_A) = \mathbf{p}_A^1(t_A)[\Theta_B = \theta^H] = t_A,$$

and

$$\mathbf{p}_B^2(\theta_B, e_B)[\mathbf{p}_A^{1,H} \leq x] = \tau_B(\theta_B, e_B)[\mathbf{t}_A \leq x] = \int_0^x f_{e_B}(t_A) dt_A,$$

for all $t_A \in T_A$, $(\theta_B, e_B) \in T_B$, and $x \in [0, 1]$.

3.4. Equilibrium

A Bayesian equilibrium of the Trust Minigame with incomplete information is given by a pair of measurable decision functions $(\sigma_A: T_A \rightarrow \{I, O\}, \sigma_B: T_B \rightarrow \{S, K\})$ such that, for each player $i \in \{A, B\}$ and type $t_i \in T_i$, choice $\sigma_i(t_i)$ maximizes player i 's expected utility, given the endogenous beliefs of type t_i about

⁸ We often write $\mathbb{P}_{t_A}[\cdot]$ instead of $\tau_A(t_A)[\cdot]$ to ease notation in the context of examples and models.

⁹ For the solution of this example, we do not have to specify a particular parametric family of density functions.

the coplayer’s choice and beliefs.¹⁰ Note that, in general, this is a subjective notion of equilibrium, because players’ exogenous beliefs are not necessarily derived from an objective distribution of types. For more on this, see §3.8.

In a perfect Bayesian equilibrium, player B maximizes his conditional expected utility on observing I , with conditional beliefs computed by Bayes’ rule, if possible. Of course, if $\sigma_A(t_A) = O$ for every t_A , the conditional second-order belief $\beta_B^I(t_B) = \mathbb{E}_{t_B}[\alpha_A \mid \sigma_A = I]$ cannot be determined by Bayes’ formula, and we cannot rule out the possibility that $\beta_B^I(t_B) = 0$, hence $\sigma_B(t_B) = K$, for every t_B . This in turn implies that each type of player A is certain of K ($\alpha_A(t_A) = 0$), which justifies $\sigma_A(t_A) = O$ for every t_A . This explains the following remark.¹¹

REMARK 1. Every Bayesian psychological game based on the Trust Minigame with guilt aversion has a perfect Bayesian equilibrium with $\sigma_A(t_A) = O$ and $\sigma_B(t_B) = K$ for all types t_A and t_B .

Having established this once and for all, the rest of the analysis is focused on nondegenerate equilibria where a positive fraction of A ’s types choose I . Under our assumptions on exogenous beliefs, this implies that every type of B assigns positive probability to I . As we noted in §2, in this case ex ante maximization of psychological utility is equivalent to conditional maximization (see Equation (2)); therefore, nondegenerate Bayesian equilibria are also perfect. When A is commonly known to be selfish, such equilibria have another interesting feature: since the positive fraction of A ’s types who choose I in equilibrium do this to maximize expected payoff, then, on observing I , player B must conclude that A chose I

¹⁰ Such decision functions are often called “strategies.” We avoid this terminology for two reasons. First, we are not studying a situation where player i decides how to play the game before being informed about her type; rather, we study decisions of different agents playing in role i , where each agent is characterized by some type t_i . Second, we want to avoid confusion with the strategies of the Trust Minigame, such as “Share if In.”

¹¹ Remark 1 is an instance of a more general observation. Fix a game form with material payoffs and no chance moves. Let G denote the corresponding complete-information game obtained when the game form and the fact that players are selfish are common knowledge. Let $\Gamma(G)$ denote any psychological game obtained from G by adding to each player’s material payoff a (possibly null) guilt-aversion term, and possibly allowing for incomplete information about guilt parameters. Then, every pure-strategy sequential equilibrium of the material-payoff game G is also a perfect Bayesian equilibrium of the psychological game $\Gamma(G)$. The intuition is that off the equilibrium path players may believe that deviations occurred by mistake and hence do not signal expectations of high material payoffs. Thus, the best reply at off-path information sets is to maximize one’s own expected material payoff. See Battigalli and Dufwenberg (2007, Observation 2).

rationally, hence that A assigned at least 50% probability to S . In other words, in a nondegenerate equilibrium, each type t_B is certain, conditional on I , that $\alpha_A \geq 1/2$: $\beta_B^I(t_B)[\alpha_A \geq 1/2] = 1$. This is exactly the same inference imposed by forward-induction reasoning (see Dufwenberg 2002, and Battigalli and Dufwenberg 2009, §5). Hence, our focus on nondegenerate equilibria can also be motivated as a forward-induction refinement. Indeed, some of our insights are solely based on a kind of step-by-step forward-induction reasoning, whereas others need the fully fledged Bayesian equilibrium analysis.

3.5. Higher-Order Endogenous Beliefs

It is important to understand how the type structure and decision functions σ_i generate the players’ endogenous beliefs. We analyze psychological games where the utility of i (determined by her guilt type θ_i) depends on what $-i$ plans to do ($-i$ ’s strategy) and on the endogenous first-order beliefs of $-i$. For example, the utility of each guilt type θ_B depends on B ’s material payoff (determined by the sequence of actions) and on the disappointment of A ; the latter is positive if A plans to choose I and carries out such plan, and then B replies with K ; in this case, A ’s disappointment is determined by the first-order belief of A about the choice of B , that is, the probability α_A assigned by A to strategy S .

This probability is an endogenous first-order belief determined by the type of A and the equilibrium decision function of B

$$\alpha_A(t_A) = \tau_A(t_A)[\sigma_B = S]. \quad (3)$$

For player B (and the analyst), $\alpha_A: T_A \rightarrow [0, 1]$ is a random variable. Type t_B of player B can compute his initial expectation of α_A as follows:¹²

$$\beta_B^{\emptyset}(t_B) = \mathbb{E}_{t_B}[\alpha_A] = \int \alpha_A(t_A) \tau_B(t_B) [dt_A]. \quad (4)$$

This initial second-order belief reflects how B reasons about the game before playing it.¹³ Player B , however, takes an action only if he observes I ; therefore, his choice depends on his second-order belief conditional on I

$$\begin{aligned} \beta_B^I(t_B) &= \mathbb{E}_{t_B}[\alpha_A \mid \sigma_A = I] \\ &= \frac{1}{\alpha_B(t_B)} \int_{t_A: \sigma_A(t_A)=I} \alpha_A(t_A) \tau_B(t_B) [dt_A], \\ &\quad \text{if } \alpha_B(t_B) > 0, \end{aligned} \quad (5)$$

¹² Given a real-valued random variable $x_{-i}: T_{-i} \rightarrow \mathbb{R}$ and a measure $\mu \in \Delta(T_{-i})$, $\mathbb{E}_\mu[x_{-i}]$ denotes the expectation of x_{-i} according to μ . To ease notation for the expectation of x_{-i} according to the belief of type t_i , we write $\mathbb{E}_{t_i}[x_{-i}]$ instead of $\mathbb{E}_{\tau_i(t_i)}[x_{-i}]$.

¹³ Hence it is an interesting feature of beliefs that is worth eliciting in experiments. See the discussion in §6.

where $\alpha_B(t_B) = \tau_B(t_B)[\sigma_A = I]$ is the initial endogenous first-order belief of t_B (see §2).

As the above equations illustrate, all of the endogenous beliefs are implicitly determined by the equilibrium decision functions $\sigma = (\sigma_A, \sigma_B)$ given the type structure \mathcal{T} . For the sake of clarity, we will make the key features of endogenous beliefs explicit.

Besides endogenous first- and second-order beliefs, the type structure and decision functions determine other random variables that will be used in our analysis (all in bold). For example, the random variable “material payoff of player i ” is¹⁴

$$m_i(t_A, t_B) = \begin{cases} m_i(O) & \text{if } \sigma_A(t_A) = O, \\ m_i(I, K) & \text{if } \sigma_A(t_A) = I, \sigma_B(t_B) = K, \\ m_i(I, S) & \text{if } \sigma_A(t_A) = I, \sigma_B(t_B) = S, \end{cases}$$

and the random variable “psychological utility of player i ” is

$$u_i(t_A, t_B) = m_i(t_A, t_B) - \mathfrak{D}_i(t_i) \max\{0, \mathbb{E}_{t_{-i}}[\mathbf{m}_{-i}] - m_{-i}(t_A, t_B)\},$$

where, of course, in the computation of $\mathbb{E}_{t_{-i}}[\mathbf{m}_{-i}]$, type t_{-i} assigns probability one to the choice $\sigma_{-i}(t_{-i})$.

As in the leading example, we use models where the type set of player i can be factorized as $T_i = \Theta_i \times \mathcal{E}_i$, with the convenient parametrization $\mathcal{E}_i = [0, 1]$. The second component of $t_i = (\theta_i, e_i)$ —the epistemic type of player i —is a random variable from the point of view of the coplayer $-i$. Formally, this random variable is just the projection from $T_i = \Theta_i \times \mathcal{E}_i$ onto $\mathcal{E}_i = [0, 1]$: $e_i(t_A, t_B) = e_i$ if and only if $t_i = (\theta_i, e_i)$ for some $\theta_i \in \Theta_i$. Thus, for example, $[e_i > x]$ denotes the event that the epistemic type of i is higher than threshold x . Given this, we can ease notation, writing the belief maps as a function of e_i only, as in $\tau_i(e_i)$, $\alpha_i(e_i)$, $\beta_i(e_i)$.

3.6. Leading Example: Equilibrium and Endogenous Beliefs

Under the simplifying assumptions of the example (see §3.3), there is a unique nondegenerate equilibrium pair of decision functions (σ_A, σ_B) that can be determined with the forward-induction argument mentioned in §3.4. Because in a nondegenerate equilibrium a positive fraction of A types choose I and each f_{e_B} has full support on $[0, 1]$, each epistemic type e_B assigns positive probability to I ; i.e.,

$$\alpha_B(e_B) = \int_{t_A: \sigma_A(t_A)=I} f_{e_B}(t_A) dt_A > 0.$$

¹⁴ Recall that $m_i(z)$ is the material payoff of player i at terminal history z .

Hence, $\beta_B^I(e_B)$ is determined by Bayes’ rule (see Equation (5))

$$\beta_B^I(e_B) = \frac{1}{\alpha_B(e_B)} \int_{t_A: \sigma_A(t_A)=I} \alpha_A(t_A) f_{e_B}(t_A) dt_A.$$

Player A ’s rationality implies¹⁵

$$\sigma_A(t_A) = \begin{cases} O & \text{if } \alpha_A(t_A) < \frac{1}{2}, \\ I & \text{if } \alpha_A(t_A) > \frac{1}{2}. \end{cases}$$

Therefore,

$$\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A \mid \alpha_A \geq \frac{1}{2}] \geq \frac{1}{2}.$$

Player B ’s rationality implies

$$\sigma_B(\theta_B, e_B) = \begin{cases} K & \text{if } \theta_B \beta_B^I(e_B) < 1, \\ S & \text{if } \theta_B \beta_B^I(e_B) > 1. \end{cases}$$

Because $\theta^L < 1$, $\theta^H > 2$, and $\beta_B^I(e_B) \geq 1/2$, the choice of B is independent of the epistemic type e_B : B Keeps if selfish and Shares if prone to guilt feelings; i.e.,

$$\sigma_B(\theta_B, e_B) = \begin{cases} K & \text{if } \theta_B = \theta^L, \\ S & \text{if } \theta_B = \theta^H. \end{cases}$$

Therefore, the endogenous first-order belief of a type t_A coincides with the exogenous one, i.e., the probability assigned to the high-guilt type of B

$$\alpha_A(t_A) = \tau_A(t_A)[\mathfrak{D}_B = \theta^H] = t_A,$$

and the decision function of A is

$$\sigma_A(t_A) = \begin{cases} O & \text{if } t_A < \frac{1}{2}, \\ I & \text{if } t_A > \frac{1}{2}. \end{cases}$$

Although in this example the fine details of higher-order beliefs do not matter, we derive the endogenous second-order beliefs of B for illustrative purposes

$$\beta_B^O(e_B) = \mathbb{E}_{e_B}[\mathbf{t}_A] = \int_0^1 t_A f_{e_B}(t_A) dt_A,$$

$$\beta_B^I(e_B) = \mathbb{E}_{e_B} \left[\mathbf{t}_A \mid \mathbf{t}_A > \frac{1}{2} \right] = \frac{\int_{1/2}^1 t_A f_{e_B}(t_A) dt_A}{\int_{1/2}^1 f_{e_B}(t_A) dt_A}.$$

3.7. Harsanyi’s Method and Psychological Games

There is a noteworthy difference between the definition of equilibrium in psychological games with complete information and in Bayesian psychological games: in the former, it is necessary to assume that endogenous beliefs of all orders are correct (see

¹⁵ We can ignore knife-edge cases because beliefs about the coplayer’s type are absolutely continuous.

Geanakoplos et al. 1989, and Battigalli and Dufwenberg 2009); in the latter, it is instead assumed that conjectures about the coplayers' decision functions are correct, but there is no explicit condition concerning belief hierarchies. In other words, the analysis of Bayesian psychological games just requires one to apply the equilibrium concept that was already on the shelves of standard game theory, whereas an extension of the traditional definition of equilibrium is needed for the analysis of complete-information psychological games. How can the two definitions be reconciled?

Following the method proposed by Harsanyi (1967, 1968a, b) in his three-part article, we posit a Bayesian game Γ and a profile of decision functions σ . The pair (Γ, σ) is an interactive beliefs structure generating a description of the possible hierarchies of beliefs about parameters (utility functions) and about choices, as illustrated in the leading example. These hierarchies satisfy by construction all of the necessary coherence conditions, given σ . In a psychological game, if the profile of decision functions σ is a Bayesian equilibrium, the choice of each type is a best reply to its belief hierarchy, and this is common belief; therefore, no further condition has to be added to the definition of equilibrium. Unlike a Bayesian psychological game, a complete-information psychological game does not come equipped with a type structure; hence, there seems to be no way to automatically unfold belief hierarchies starting from a profile of strategies s .

Despite this, there is no substantial difference between the two equilibrium concepts; as a matter of fact, one is a special case of the other. To see this, note that a complete-information psychological game G can be interpreted as a trivial Bayesian game Γ with just one type for each player. In this case, a profile σ of maps from types to strategies is just a profile of strategies s in G . Hence, the unique type of each player i corresponds to the degenerate hierarchy of beliefs whereby beliefs of all orders are correct: the first-order belief of i assigns probability one to s_{-i} , and higher-order beliefs assign probability one to the lower-order beliefs of coplayers. Strategy profile s is a Bayesian equilibrium if each player's strategy is a best reply to the belief hierarchy of this player's unique type. Therefore, the profile of trivial maps $\sigma = s$ and associated belief hierarchies is an equilibrium of the trivial Bayesian game Γ if and only if it is an equilibrium of G according to the complete-information definition.

However, Harsanyi's methodology allows for a more flexible approach to complete-information equilibrium: a type structure in the sense of Harsanyi may have so-called "redundant" types, i.e., distinct types that nonetheless feature the same utility parameter and the same hierarchy of beliefs about utility

parameters. This means that a complete-information game G can be equipped with a type structure \mathcal{T} with multiple redundant types, thus obtaining a non-trivial Bayesian game $\Gamma = (G, \mathcal{T})$. Although all types in Γ have the same utility function and hierarchical beliefs about utility functions, now an equilibrium profile σ can map different types to different strategies. As a consequence, in a Bayesian equilibrium players may be uncertain about the strategies and hierarchical beliefs about strategies of the coplayers.¹⁶ In other words, even if the utility functions are common knowledge and hence exogenous hierarchies of beliefs are trivially unique, there may be multiple hierarchies of endogenous beliefs. If there are sequential moves, this implies that players can change their mind about the coplayers' intentions as the play unfolds on an equilibrium path, which is impossible according to the complete-information equilibrium concepts of Geanakoplos et al. (1989) and Battigalli and Dufwenberg (2009).¹⁷

3.8. Actual Distribution of Types and Predictions

We adapt to psychological games the general notion of Bayesian equilibrium of Harsanyi (1967, 1968a, b), which is inherently subjective, because players' beliefs are not derived from an objective distribution over types. In fact, our models do not assume that players know the objective statistical distribution of types and derive their exogenous beliefs from such distribution.¹⁸ If this were the case, different types could have different beliefs only if the types of A and B were correlated, as in the case (for example) with assortative matching. However, under the random-matching structure typical of lab experiments, the types of A and B are objectively independent, with marginal probabilities given by the frequency distribution of types in the population from which subjects are drawn at random. Hence, conditioning on his own type, a player cannot learn anything about the type of the coplayer. With this, in a Bayesian equilibrium, a player's first- and second-order beliefs about the coplayer would be type independent, contrary to the findings of the experimental literature, which suggests instead that such beliefs are very heterogeneous.¹⁹

We analyze equilibria of Bayesian games with subjective beliefs to allow for such heterogeneity. Hence,

¹⁶ With standard (i.e., belief-independent) preferences, this is equivalent to the subjective correlated equilibrium concept (Brandenburger and Dekel 1987).

¹⁷ See the notion of "polymorphic sequential equilibrium" of Battigalli et al. (2015, p. 8).

¹⁸ In other words, we do not assume an "objective" common prior on the state space.

¹⁹ See the references cited in §1.

our equilibrium analysis does not have to posit an objective statistical distribution on the type space. Of course, such a distribution is necessary to derive statistical predictions. This is apparent in our leading example. The relative frequency of the trusting strategy I and of “optimistic” first-order beliefs ($\alpha_A > 1/2$) is just the fraction of A agents in the population who believe that the high-guilt type θ^H is more likely than the low-guilt type θ^L . The relative frequency of the sharing strategy S coincides with the fraction of B agents whose guilt type is θ^H . Finally, the relative frequency of “optimistic” initial second-order beliefs ($\beta_B^{\otimes} > 1/2$) is the fraction of B agents who believe that with more than 50% probability A deems θ^H more likely than θ^L .

A subtler question is whether we should expect to observe a positive correlation between (conditional) second-order beliefs and the propensity to Share. It turns out that this is the case if the guilt and epistemic components of B 's type are statistically independent.²⁰ We extensively comment on statistical predictions in §6.

4. Role-Dependent Guilt

We analyze a model that generalizes the example of §3.3: player A is commonly known to be a material-payoff maximizer, but there are infinitely many possible guilt types of player B . If the agents playing in roles A and B are drawn at random from the same population, as in most experiments, then we are assuming that the “potential” guilt sensitivity of an agent becomes an actual tendency to live up to the other's expectations only if this agent plays in the role of the “trustee” B . See the discussion in §1.

4.1. Type Structure

Because player A is commonly known to be selfish, $\Theta_A = \{0\}$. The set of possible guilt types of B is a closed interval $\Theta_B = [\theta^L, \theta^H]$ with $0 \leq \theta^L < 1$ and $\theta^H > 2$. The beliefs of each player i about the type of coplayer $-i$ are solely determined by the epistemic type $e_i \in \mathcal{E}_i$, with the convenient parametrization $\mathcal{E}_i = [0, 1]$. Therefore, $T_A = \{0\} \times \mathcal{E}_A \cong \mathcal{E}_A$ and $T_B = [\theta^L, \theta^H] \times [0, 1]$.

Each type of player A believes that the guilt and epistemic types of B are statistically independent. We let $e_A = t_A \in T_A$ parametrize A 's subjective distribution over the guilt types of B : $\mathbb{P}_{t_A}[\mathfrak{D}_B < z] = G_{t_A}(z)$, where the cumulative distribution function (cdf) $G_{t_A}: \mathbb{R} \rightarrow [0, 1]$ has support $[\theta^L, \theta^H]$ and is continuous on (θ^L, θ^H) , and $G_{t_A}(z)$ is continuous in t_A

for each $z \in (\theta^L, \theta^H)$.²¹ The beliefs about \mathfrak{D}_B of higher types first-order stochastically dominate those of lower types

$$t'_A < t''_A \Rightarrow G_{t'_A}(z) > G_{t''_A}(z), \quad (6)$$

for all $z \in [\theta^L, \theta^H]$ and $t'_A, t''_A \in [0, 1]$.

We also assume that there exist thresholds $\underline{t}_A > 0$ and $\bar{t}_A < 1$ such that every type $t_A < \underline{t}_A$ believes that $\mathfrak{D}_B < 1$ with more than 50% probability, and every type $t_A > \bar{t}_A$ believes that $\mathfrak{D}_B > 2$ with more than 50% probability

$$\begin{aligned} \underline{t}_A &:= \sup\{t_A: G_{t_A}(1) > \frac{1}{2}\} > 0, \\ \bar{t}_A &:= \min\{t_A: G_{t_A}(2) \leq \frac{1}{2}\} < 1. \end{aligned} \quad (7)$$

By continuity of $G_{t_A}(z)$ in t_A and z , and the stochastic-order assumption (6), \underline{t}_A (respectively, \bar{t}_A) is the unique solution to $G_{t_A}(1) = 1/2$ (respectively, $G_{t_A}(2) = 1/2$), and $\underline{t}_A < \bar{t}_A$.

The marginal beliefs of each type t_A about the epistemic type of B are given by the same continuous cdf $F: \mathbb{R} \rightarrow [0, 1]$ with support $[0, 1]$.²² The resulting belief function $\tau_A: T_A \rightarrow \Delta(T_B)$ satisfies

$$\tau_A(t_A)[\mathfrak{D}_B \leq z \cap e_B \leq y] = G_{t_A}(z)F(y), \quad (8)$$

for all $t_A, y \in [0, 1]$ and $z \in [\theta^L, \theta^H]$.

Each epistemic type e_B of B has beliefs about A 's type given by a continuous cdf $F_{e_B}: \mathbb{R} \rightarrow [0, 1]$ with support $[0, 1]$. The resulting belief function $\tau_B: T_B \rightarrow \Delta(T_A)$ satisfies

$$\tau_B(\theta_B, e_B)[\mathbf{t}_A \leq x] = F_{e_B}(x), \quad (9)$$

for all $\theta_B \in [\theta^L, \theta^H]$ and $e_B, x \in [0, 1]$.²³ Furthermore, we assume that the following stochastic-order property holds: the conditional expectations $\mathbb{E}_{e_B}[\mathbf{t}_A | \mathbf{t}_A > x]$ are strictly increasing in e_B

$$\begin{aligned} e'_B < e''_B &\Rightarrow \frac{1}{1 - F_{e'_B}(x)} \int_x^1 t_A \, dF_{e'_B}(t_A) \\ &< \frac{1}{1 - F_{e''_B}(x)} \int_x^1 t_A \, dF_{e''_B}(t_A), \end{aligned} \quad (10)$$

²¹ The distribution G_{t_A} has support $[\theta^L, \theta^H]$ if it is strictly increasing on $[\theta^L, \theta^H]$, $G_{t_A}(z) = 0$ for $z < \theta^L$, and $G_{t_A}(z) = 1$ for $z \geq \theta^H$. We allow for atoms at θ^L and θ^H , as in the parametrized cdf

$$G_{t_A}(z) = \begin{cases} 0 & z < \theta^L, \\ 1 - t_A - \varepsilon + \varepsilon \frac{z - \theta^L}{\theta^H - \theta^L} & \theta^L \leq z < \theta^H, \\ 1 & z \geq \theta^H, \end{cases}$$

which essentially gives back the leading example for ε small: $\mathbb{P}_{t_A}[\mathfrak{D}_B = \theta^H] = t_A$, $\mathbb{P}_{t_A}[\mathfrak{D}_B = \theta^L] = 1 - t_A - \varepsilon$.

²² That is, F is strictly increasing on $[0, 1]$ with $F(0) = 1 - F(1) = 0$.

²³ Conditions (8) and (9) imply that there is no perception of false consensus. See the discussion in §6.

²⁰ This is just a sufficient condition for all models based on the Trust Minigame with guilt aversion.

for all $e'_B, e''_B \in [0, 1]$ and $x \in [0, 1]$. Intuitively, this means that higher epistemic types of B have higher beliefs about the (epistemic) type of A .²⁴

All of the above is common knowledge. Because the beliefs of type (θ_B, e_B) depend only on the epistemic component e_B , to ease notation, we write $\alpha_B(e_B)$, $\beta_B^{\mathcal{O}}(e_B)$, $\beta_B^I(e_B)$ instead of, respectively, $\alpha_B(\theta_B, e_B)$, $\beta_B^{\mathcal{O}}(\theta_B, e_B)$, $\beta_B^I(\theta_B, e_B)$, as we did in the leading example.

4.2. Equilibrium Analysis

By Remark 1, there is a pooling equilibrium with no trust and no cooperation. Henceforth, we study the nondegenerate equilibria, i.e., those where a positive fraction of A types choose In. Our assumptions imply that all nondegenerate equilibria exhibit threshold decision functions whereby higher types choose the prosocial action (the threshold types cannot be atoms; therefore their choices are immaterial). We call such functions “monotone.”

DEFINITION 1. Decision function σ_A is *monotone (increasing)* if there is a threshold $\hat{t}_A \in [0, 1]$ such that, for every t_A ,

$$\begin{aligned} t_A < \hat{t}_A &\Rightarrow \sigma_A(t_A) = O, \\ t_A > \hat{t}_A &\Rightarrow \sigma_A(t_A) = I. \end{aligned}$$

Decision function σ_B is *monotone (increasing)* in θ_B if, for every $e_B \in [0, 1]$, there is a threshold $\hat{\theta}_B(e_B) \in [\theta^L, \theta^H]$ such that, for every θ_B ,

$$\begin{aligned} \theta_B < \hat{\theta}_B(e_B) &\Rightarrow \sigma_B(\theta_B, e_B) = K, \\ \theta_B > \hat{\theta}_B(e_B) &\Rightarrow \sigma_B(\theta_B, e_B) = S. \end{aligned}$$

Finally, decision function σ_B is *monotone (increasing)* in e_B if, for every $\theta_B \in [\theta^L, \theta^H]$, there is a threshold $\hat{e}_B(\theta_B)$ such that, for every e_B ,

$$\begin{aligned} e_B < \hat{e}_B(\theta_B) &\Rightarrow \sigma_B(\theta_B, e_B) = K, \\ e_B > \hat{e}_B(\theta_B) &\Rightarrow \sigma_B(\theta_B, e_B) = S. \end{aligned}$$

A nondegenerate equilibrium (σ_A, σ_B) determines the endogenous belief functions α_A , $\beta_B^{\mathcal{O}}$, and β_B^I as in Equations (3)–(5), so that $\sigma_A^{-1}(I)$ has a positive measure, $\sigma_A(t_A)$ is a best reply to $\alpha_A(t_A)$ for all t_A , and $\sigma_B(\theta_B, e_B)$ is a best reply to $\beta_B^I(e_B)$ for all θ_B and e_B .

²⁴ This assumption holds if the epistemic types of B are ordered by hazard rate. When every cdf F_{e_B} is differentiable with density f_{e_B} , this can be expressed as follows:

$$e'_B < e''_B \Rightarrow \frac{f_{e'_B}(t_A)}{1 - F_{e'_B}(t_A)} < \frac{f_{e''_B}(t_A)}{1 - F_{e''_B}(t_A)},$$

for all $e'_B, e''_B \in [0, 1]$ and $t_A \in [0, 1]$. See Shaked and Shantikumar (2007, pp. 16–17).

The incentive conditions give

$$\alpha_A(t_A) < \frac{1}{2} \Rightarrow \sigma_A(t_A) = O, \quad (11)$$

$$\alpha_A(t_A) > \frac{1}{2} \Rightarrow \sigma_A(t_A) = I,$$

$$2 < 4 - 2\theta_B \beta_B^I(e_B) \Rightarrow \sigma_B(\theta_B, e_B) = K, \quad (12)$$

$$2 > 4 - 2\theta_B \beta_B^I(e_B) \Rightarrow \sigma_B(\theta_B, e_B) = S,$$

for all $t_A \in T_A$, $(\theta_B, e_B) \in T_B$.

PROPOSITION 1. Every nondegenerate equilibrium of the model given by (6)–(10) has the following structure.

(a) σ_A is monotone with threshold $\hat{t}_A \in [\underline{t}_A, \bar{t}_A]$, which is the unique solution to equation

$$\alpha_A(\hat{t}_A) = \frac{1}{2},$$

where

$$\alpha_A(t_A) = \int_{[0, 1]} (1 - G_{t_A}(1/\beta_B^I(e_B))) dF(e_B),$$

for all $t_A \in [0, 1]$.

(b) σ_B is monotone in θ_B with threshold function $\hat{\theta}(e_B) = 1/\beta_B^I(e_B)$, and monotone in e_B with threshold function $\hat{e}(\theta_B) = (\beta_B^I)^{-1}(1/\theta_B)$.

(c) The endogenous beliefs of B satisfy

$$\alpha_B(e_B) = 1 - F_{e_B}(\hat{t}_A) > 0,$$

$$\beta_B^{\mathcal{O}}(e_B) = 1 - \int_{[0, 1]} \int_{[0, 1]} G_{t_A}(1/\beta_B^I(x)) dF(x) dF_{e_B}(t_A),$$

$$\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | t_A > \hat{t}_A] \geq \frac{1}{2},$$


for all $e_B \in [0, 1]$.

Proposition 1 characterizes the structure of all nondegenerate equilibria of the model. Such equilibria have monotone decision functions. By condition (6), the higher the epistemic type of A , the higher A 's belief that the guilt type of B is high. Because higher guilt types of B have a higher propensity to Share, A 's first-order belief is increasing in t_A ; hence, according to the incentive condition (11), all epistemic types of A higher than \hat{t}_A choose In.

By (10), higher epistemic types of B hold higher second-order beliefs on the epistemic type of A conditional on I . Because all epistemic types of A higher than \hat{t}_A choose In, $\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | t_A > \hat{t}_A]$ is increasing in e_B . Hence, incentive condition (12) implies that the decision function of B is monotone both in θ_B and in e_B . Such decision function is characterized by a decreasing threshold function $\hat{e}(\theta_B)$, as shown in Figure 3.

As anticipated in §3.4, our assumptions imply that nondegenerate equilibria are the only ones consistent with forward-induction reasoning. Choice I signals that $\alpha_A \geq 1/2$. Therefore, $\beta_B^I(e_B) \geq 1/2$, and A predicts

Figure 3 (Color online) Equilibrium Choice of B



that all of the types with $\theta_B > 2$ would Share; thus, $\alpha_A(t_A) \geq \mathbb{P}_{t_A}[\mathfrak{D}_B > 2]$. By assumption, $\mathbb{P}_{t_A}[\mathfrak{D}_B > 2] > 1/2$ for all $t_A \in (\bar{t}_A, 1]$; hence, the measure of the set of A types trusting B is at least $1 - \bar{t}_A > 0$.

On the other hand, $1 - \alpha_A(t_A) \geq \mathbb{P}_{t_A}[\mathfrak{D}_B < 1]$. By assumption, $\mathbb{P}_{t_A}[\mathfrak{D}_B < 1] < 1/2$ for all $t_A \in [0, \bar{t}_A]$; hence, the measure of the set of A types that do not trust B is at least $\bar{t}_A > 0$. Therefore, there is heterogeneity of behavior and of endogenous beliefs among A types. As for B , heterogeneity of behavior is quite obvious from incentive condition (12), given that B uses forward induction, and that there are types with $\theta_B > 2$ and types with $\theta_B < 1$. Heterogeneity of endogenous beliefs follows from forward-induction reasoning and our assumptions about exogenous beliefs.

5. Role-Independent Guilt

We now analyze a model of situations where the agents playing in roles A and B are drawn from the same population, as in most games played in the laboratory. If guilt aversion is not affected by the role played in the game, the type structure must be symmetric. To anticipate, the main difference from the model of §4 is that here also player A may experience guilt feelings triggered by the expectation of B 's disappointment. This is related to other differences between the assumed type structures.


Player B can only be disappointed after the terminal history O , in which case the extent of his disappointment also depends on what he plans to do in the subgame, i.e., his strategy. To derive B 's disappointment, first note that his expected material payoff is

$$\mathbb{E}_B[\mathbf{m}_B] = \begin{cases} 1 \cdot (1 - \alpha_B) + 2 \cdot \alpha_B & \text{if } s_B = S, \\ 1 \cdot (1 - \alpha_B) + 4 \cdot \alpha_B & \text{if } s_B = K. \end{cases}$$

Because $m_B(O) = 1$ is the lowest material payoff for B , $\mathbb{E}_B[\mathbf{m}_B] \geq m_B(O)$, and B 's disappointment after O is

$$\begin{aligned} & \max\{0, \mathbb{E}_B[\mathbf{m}_B] - m_B(O)\} \\ & = \mathbb{E}_B[\mathbf{m}_B] - 1 = \begin{cases} \alpha_B & \text{if } s_B = S, \\ 3\alpha_B & \text{if } s_B = K. \end{cases} \end{aligned} \tag{13}$$

Figure 4 The Trust Minigame with Psychological Utilities of A and B



We can represent this strategic situation with a psychological game parametrized by the guilt sensitivities θ_A and θ_B . To analyze such version of the Trust Minigame with guilt aversion, we need to expand our notation about beliefs by introducing a feature of Ann's second-order beliefs, her expectation of Bob's disappointment if she goes Out

$$\bar{\beta}_A = \mathbb{E}_A[\mathbb{E}_B[\mathbf{m}_B] - m_B(O)].$$

The psychological game with role-independent guilt aversion is more easily represented in a sort of reduced form where each player's psychological utility depends on each player's own endogenous second-order belief rather than the coplayer's endogenous first-order belief, as shown in Figure 4.²⁵

5.1. Type Structure

The possibility that both players can feel guilt complicates the analysis. Therefore, we introduce some restrictive assumptions on the type space to maintain the tractability of the model. In particular, whereas keeping the continuum of epistemic types on both sides, we now assume that each player i 's guilt type can only take two values, low or high; i.e., $\theta_i \in \{\theta^L, \theta^H\}$, with $\theta^L = 0$ and $\theta^H > 1$. As in the previous section, we assume that $\mathcal{E}_i = [0, 1]$, and that each epistemic type of each player i believes that the guilt and epistemic types of player $-i$ are independent.²⁶ Specifically, we model the exogenous beliefs of both players as we did for player A in the leading example of §3.3: e_i parametrizes i 's subjective probability of the high-guilt type of the coplayer: $\tau_i(\theta_i, e_i)[\mathfrak{D}_{-i} = \theta^H] = e_i$. This implies that, for each i , $t_i = (\theta_i, e_i) \in \{\theta^L, \theta^H\} \times [0, 1] = T_i$ and that we can write $\tau_i: [0, 1] \rightarrow \Delta(\{\theta^L, \theta^H\} \times [0, 1])$. As a consequence, the endogenous second-order beliefs of players A and B are independent of their guilt type. We also assume

²⁵ Here we use an observation by Battigalli and Dufwenberg (2009): A psychological utility function of the form $u_i: \Theta_i \times Z \times \mathcal{H}_i \times \mathcal{H}_{-i} \rightarrow \mathbb{R}$ can be replaced by a utility function $\tilde{u}_i: \Theta_i \times Z \times \mathcal{H}_i \rightarrow \mathbb{R}$ inducing the same best-reply correspondence, which depends only on the endogenous beliefs of i . As the example shows, this may require replacing the low-order initial beliefs of others with one's own higher-order conditional beliefs.

²⁶ In the model with role-dependent guilt this assumption holds trivially for the beliefs of B about A , because there is only one possible guilt type of A .

that each type of each player has the same marginal beliefs about the epistemic type of the coplayer given by a continuous cdf F with support $[0, 1]$. Thus,

$$\forall e_i \in [0, 1], \forall x, \tau_i(e_i)[\mathfrak{D}_{-i} = \theta^H \cap \mathbf{e}_{-i} \leq x] = e_i F(x), \quad (14)$$

where F is strictly increasing on $[0, 1]$ and admits a density f . Note that here the epistemic type of B parametrizes an exogenous first-order belief, i.e., B 's subjective probability that the guilt type of A is high. By contrast, in the model of §4, there is only one possible guilt type of A and e_B parametrizes the exogenous second-order belief of B .

By Equation (14), i 's expectation of \mathbf{e}_{-i} is independent of e_i ; hence, we write $\mathbb{E}_{e_i}[\mathbf{e}_{-i}] = \mathbb{E}[\mathbf{e}_{-i}]$.²⁷ To simplify the exposition and avoid tedious discussions of subcases in the equilibrium analysis, we assume that this expectation is not too low

$$\mathbb{E}[\mathbf{e}_{-i}] > \frac{1}{3}. \quad (15)$$

5.2. Equilibrium Analysis

Remark 1 applies to this model as well; hence, there is a pooling equilibrium with no trust and no cooperation. Henceforth, we focus on the characterization of equilibria (σ_A, σ_B) such that $\sigma_A^{-1}(I)$ has a positive measure, i.e., the nondegenerate equilibria. The most important difference from the model of §4 is that, here, choice I is not a clear signal of A 's trust, because A 's guilt type may be high, and high-guilt type (θ^H, e_A) may choose I in equilibrium even if $\alpha_A(e_A) < 1/2$, to avoid disappointing B . Specifically, Equation (13) shows that B 's disappointment is maximal when he plans to Keep. Therefore, choice I may be interpreted as a signal that A 's guilt type is high, A expects B to Keep and goes In to prevent B 's disappointment. The following result is a stark illustration of this phenomenon.

PROPOSITION 2. *In the model given by (14) and (15) the following is an equilibrium: $\sigma_A(\theta^L, e_A) = O$, $\sigma_A(\theta^H, e_A) = I$, and $\sigma_B(\theta^L, e_B) = \sigma_B(\theta^H, e_B) = K$ for all e_A and e_B , which yields the following endogenous beliefs: $\alpha_A(e_A) = 0$, $\bar{\beta}_A(e_A) = 3\mathbb{E}[\mathbf{e}_B]$, $\alpha_B(e_B) = e_B$, and $\beta_B^\emptyset(e_B) = \beta_B^I(e_B) = 0$ for all e_A and e_B .*

The equilibrium described by Proposition 2 is structurally different from the nondegenerate equilibria of §4: here, every B -type Keeps and, despite this, high-guilt types of A choose I , because the prospective guilt from disappointing B prevails over the monetary incentive. The following proposition characterizes the other nondegenerate equilibria.

²⁷ All expectations not indexed by the epistemic type e_i are determined by the common marginal cdf F on $\mathcal{E}_{-i} = [0, 1]$.

PROPOSITION 3. *In the model given by (14) and (15) the nondegenerate equilibria (σ_A, σ_B) with heterogeneous behavior by B types have the following structure: $\sigma_B(\theta^L, e_B) = K$ for every e_B ; the decision functions $\sigma_A(\theta^n, \cdot)$ ($n \in \{L, H\}$) are monotone increasing; $\sigma_B(\theta^H, \cdot)$ is monotone decreasing; and the corresponding thresholds \hat{e}_A^L , \hat{e}_A^H , and \hat{e}_B^H , with $0 \leq \hat{e}_A^H < \hat{e}_A^L \leq 1$ and $0 \leq \hat{e}_B^H < 1$, are such that*

(a)

$$\hat{e}_A^L = \min \left\{ 1, \frac{1}{2F(\hat{e}_B^H)} \right\},$$

$$\hat{e}_A^H > 0 \Rightarrow 1 - \theta^H \bar{\beta}_A(\hat{e}_A^H) = 2\alpha_A(\hat{e}_A^H),$$

and

$$\alpha_A(e_A) = e_A F(\hat{e}_B^H),$$

$$\begin{aligned} \bar{\beta}_A(e_A) &= (1 - F(\hat{e}_A^L))(3 - 2F(\hat{e}_B^H)e_A) \\ &\quad + (F(\hat{e}_A^L) - F(\hat{e}_A^H))3\mathbb{E}[\mathbf{e}_B] \\ &\quad - 2(F(\hat{e}_A^L) - F(\hat{e}_A^H))e_A F(\hat{e}_B^H) \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B \mid \sigma_B = S \cap \mathfrak{D}_B = \theta^H], \end{aligned}$$

for every e_A . Hence, $\alpha_A(\cdot)$ is increasing, and $\bar{\beta}_A(\cdot)$ is decreasing.

(b)

$$0 < \hat{e}_B^H < 1 \Rightarrow \beta_B^I(\hat{e}_B^H) = \frac{1}{\theta^H},$$

and

$$\alpha_B(e_B) = (1 - F(\hat{e}_A^L)) + (F(\hat{e}_A^L) - F(\hat{e}_A^H))e_B,$$

$$\beta_B^\emptyset(e_B) = F(\hat{e}_B^H)\mathbb{E}[\mathbf{e}_A],$$

$$\begin{aligned} \beta_B^I(e_B) &= F(\hat{e}_B^H)(\mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H]e_B \\ &\quad + \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L](1 - e_B)), \end{aligned}$$

for every e_B . Hence, $\alpha_B(\cdot)$ is increasing, $\beta_B^\emptyset(\cdot)$ is constant, and $\beta_B^I(\cdot)$ is decreasing.

The equilibria of Proposition 3 are more similar to those of §4. In both models, higher epistemic types of A are more confident that the guilt type of B is high and hence B would Share. As for player B , $\beta_B^I(e_B)$ is decreasing in e_B , and so is the decision function $\sigma_B(\theta^H, e_B)$. The reason that $\beta_B^I(e_B)$ is decreasing, unlike the model of §4 where it is increasing, is that here e_B parametrizes the first-order beliefs, not the second-order beliefs, of B .²⁸ A high (respectively, low) epistemic type of B thinks that A 's guilt type is likely to be high (respectively, low). Therefore, high epistemic types of B tend to explain choice I as the result of A 's high guilt aversion, despite the fact that α_A is low, whereas low epistemic types of B think that A is selfish and explain I with a high α_A .

²⁸ In §4 the exogenous first-order beliefs of B are trivial, because B knows that A is selfish.

6. Discussion

We discuss the relevance of our models for experimental work, and then we offer our methodological perspective on the use of the Bayesian equilibrium concept.

6.1. Empirical Predictions

An equilibrium specifies actions, beliefs about actions (endogenous first-order beliefs), and beliefs about beliefs about actions (endogenous second-order beliefs) for each type of each player. We focus on non-degenerate equilibria of the Trust Minigame where a positive fraction of A types trust the second mover, B . Qualitative predictions about behavior and hierarchical beliefs about behavior can be obtained by assuming that the actual distribution of types satisfies some mild assumptions.²⁹ Such predictions can be used to organize experimental data.

6.1.1. Heterogeneity and Correlations. If the distribution of types has a rich support and the upper bound on guilt aversion is sufficiently high, we should expect not only heterogeneous behavior but also heterogeneous hierarchical beliefs about behavior, with a large fraction of subjects who exhibit intermediate beliefs.³⁰ Furthermore, if the epistemic component of players' types, e_i , is statistically independent of the guilt component, \mathfrak{d}_i , then we should observe a positive correlation between prosocial actions and endogenous second-order beliefs, for player B and—in the case of role-independent guilt—also for player A . Indeed, the willingness to choose the prosocial action, in particular the willingness to Share of B , is an increasing function of the guilt type and of the endogenous (conditional) second-order belief. In our model, the latter depends only on the epistemic type. If epistemic and guilt types are statistically independent, then the prosocial action must be positively correlated with the endogenous second-order belief. Without relying on equilibrium analysis, Charness and Dufwenberg (2006) derive such positive correlation for B subjects (supported by the data) from the assumption that, in our notation, \mathfrak{d}_B is statistically independent of the conditional second-order belief β_B^i . This is consistent with our equilibrium analysis, but instead of *assuming* independence between \mathfrak{d}_B and the endogenous belief β_B^i , we derive it from the independence between \mathfrak{d}_B and e_B , which are both exogenous.

²⁹ Recall that our models specify interactive beliefs for each type but are silent on the actual distribution of types.

³⁰ Attanasi et al. (2013) instead implement in an experimental treatment a situation close to complete information, i.e., one where the support of exogenous beliefs for each matched pair of subjects is very small. They show that in this treatment endogenous beliefs tend to be extreme, as predicted by the complete information theory, whereas in the control treatment (incomplete information) they are indeed heterogeneous and mostly have intermediate values.

Statistical independence between the guilt and epistemic components of types is a natural benchmark. However, it is also plausible to assume that, by a kind of false-consensus effect (see Ross et al. 1977), types with higher guilt aversion tend to have higher beliefs about the aversion to guilt of the coplayer. Adding such positive correlation to the model with role-independent guilt yields a negative correlation between the conditional second-order beliefs of B subjects and their guilt type: high-guilt types of B tend to believe that the guilt type of A is high and to explain A 's trust as a desire to not disappoint B rather than to obtain a higher material payoff. This tends to decrease the correlation between the prosocial action and the conditional second-order belief. On the other hand, when A is known to be selfish (role-dependent guilt), we may have a different kind of false consensus: the higher the guilt type of B , the higher (in the stochastic sense) B 's belief about A 's belief that B 's guilt type is high. In this case, positive correlation between the guilt and epistemic components tends to strengthen the positive correlation between the prosocial action and the conditional second-order belief.³¹

6.1.2. Further Analysis of Beliefs. The theoretical insights of our models can be used by experimental economists to extend the elicitation and analysis of players' beliefs, design new experiments, and explain previous experimental results.

First, our theoretical analysis highlights the importance of beliefs that are not considered in the experimental literature on the Trust Game, specifically, the second-order beliefs of A , and the initial beliefs of B .³² The former are relevant when also A may be guilty averse: as discussed above, the presence of such guilt

³¹ The actual existence of a false-consensus effect does not imply that players' subjective beliefs must display a perception of false consensus for the coplayer. Such perceptions are modeled by the type structure. In our models there is no perception of false consensus because of the twin assumptions that the belief maps do not depend on the guilt component of players' types, and that each player deems the epistemic component of the coplayer type to be independent of the guilt component. However, taking into account what we just said about the actual false-consensus effect, we can speculate about the effect of introducing the perception of false consensus in our models. If, in the model with role-dependent guilt, we let A perceive a positive correlation between the guilt and epistemic components of B 's type, the qualitative results do not change: now A expects high-guilt types of B to be even more cooperative because A expects them to hold on average higher endogenous second-order beliefs. On the other hand, the effects of introducing a strong perception of false consensus in the model with role-independent guilt are not clear: here higher-guilt types of B should be expected to hold on average lower endogenous second-order beliefs.

³² See, e.g., Charness and Dufwenberg (2006), Ellingsen et al. (2010), and Chang et al. (2011). More precisely, Chang et al. (2011) elicit B 's first-order beliefs, although they do not use them in the analysis.

aversion should yield a positive correlation between A 's second-order beliefs and the propensity to go In. New experiments eliciting such beliefs could check for such correlation, thus providing indirect evidence on the role dependence of guilt aversion. As for B 's beliefs, choice should correlate with conditional beliefs, but also initial beliefs are relevant because they reflect B 's strategic reasoning before playing the game.³³

Our analysis also provides a potential explanation of why the conditional second-order beliefs of B subjects do not conform to the classical forward-induction argument: if A is known to be selfish, then $\beta_B^I = \mathbb{E}_B(\alpha_A | I) \geq 1/2$ because A goes In only if $\alpha_A \geq 1/2$. However, as shown in §5, if A is perceived by B as potentially guilt averse, action In may be interpreted as a desire not to disappoint B . Hence, it may well be the case that $\beta_B^I < 1/2$. Indeed, experimental data show that a significant fraction of B subjects hold such low conditional second-order beliefs.³⁴

6.2. Applicability of Subjective Bayesian Equilibrium

Our use of Bayesian equilibrium to model behavior and endogenous beliefs in games with belief-dependent preferences deserves discussion. Here we first explain why the traditional justification that equilibrium is attained through learning does not apply, then we elaborate on our interpretation of Bayesian equilibrium analysis.

6.2.1. Equilibrium and Learning. It is frequently argued that equilibrium analysis is appropriate to organize data because agents learn equilibrium behavior by playing a game many times against randomly matched coplayers. Here we do not rely on such justification for several reasons.

First and most importantly, in so far as we aim at organizing experimental data, we must take into account that, in most experiments on the Trust Game and other social dilemmas, subjects play the game just once or perhaps a few times; hence, they cannot learn.

Second, as noted by Battigalli and Dufwenberg (2009), once behavior has stabilized in a recurrent game, strategy distributions should look like a self-confirming equilibrium, which may be different from

a Nash or Bayesian equilibrium if agents have belief-dependent preferences.

A third related issue is that we use the general, subjective notion of Bayesian equilibrium, because we assume that players do not know the objective distribution of types. Then, even with standard preferences, Bayesian equilibrium is not the right tool to capture self-confirming patterns of behavior. The reason is that Bayesian equilibrium postulates that players have correct conjectures about the true (type-dependent) decision functions of coplayers, but this assumption can be justified by learning only in those rare situations where agents obtain sufficient information feedback to correctly identify the actual decision functions. However, such fine information feedback typically allows one to also identify the distribution of types, which yields an objective Bayesian–Nash equilibrium (see Dekel et al. 2004). When instead such strong assumptions about information feedback do not hold and we model steady states of learning dynamics, beliefs about parameters are not exogenous, as in subjective Bayesian equilibrium analysis, because they have to be consistent with the long-run frequencies of observations implied—via information feedback—by the frequencies of parameter values and choices (e.g., observations of realized monetary payoffs).

In summary, many experiments about social dilemmas are not designed so as to make subjects choose recurrently; hence, there cannot be any equilibrating process through learning. However, even if equilibrium analysis aims at organizing data about stabilized behavior in situations of recurrent interactions, subjective Bayesian equilibrium is not the appropriate tool, and a different approach is called for.

6.2.2. Equilibrium and Strategic Reasoning. We instead use Bayesian equilibrium analysis to provide an orderly and consistent description of strategic reasoning in an incomplete-information environment without assuming that behavior has stabilized, e.g., because subjects play just once. It has been shown that, if one drops—as we do—the assumption that exogenous beliefs are derived from an objective distribution, then the Bayesian equilibrium assumption that players hold correct conjectures about the coplayers' decision functions just ensures that behavior and endogenous beliefs are consistent with a common certainty of rationality, which is characterized by incomplete-information rationalizability (Battigalli and Siniscalchi 2003).³⁵ This result refers to equilibrium outcomes encompassing *all* of the subjective Bayesian equilibrium models based on a given

³³ The connection between strategic reasoning and hierarchies of initial beliefs in dynamic games is clarified by the literature on epistemic game theory. See Dekel and Siniscalchi (2015) and references therein.

³⁴ For example, in Charness and Dufwenberg (2006)—where the forward-induction threshold is $1/2$ ($7/10$) in treatments with 5–5 (7–7) outside options—only 42% ($19/45$) of B subjects in the control treatment with 5–5 outside options and only 31% ($15/48$) of B subjects in the control treatment with 7–7 outside options have a conditional second-order belief above the forward-induction threshold.

³⁵ This result relies on the existence of redundant types. Its earliest version is due to Brandenburger and Dekel (1987), who analyzed subjective correlated equilibria of games with complete information.

game form with parametrized utility functions. Of course, our specific assumptions about exogenous beliefs yield equilibrium implications that go beyond mere rationalizability. Therefore, we offer an analysis between objective Bayesian–Nash equilibrium and the most general notion of incomplete-information rationalizability. In particular, our discussion of the model with role-dependent guilt emphasizes that some key results about nondegenerate equilibria follow from a forward-induction logic: all of the guilt types $\theta_B > 2$ (those who would Share if they were sure that $\alpha_A \geq 1/2$) do indeed Share given A 's trusting action In, because they rationalize such action and infer that $\alpha_A \geq 1/2$.³⁶ Predicting this, all of the types of A who assign more than 50% probability to $\mathfrak{D}_B > 2$ play the trusting action. On the other hand, the types who assign more than 50% probability to $\mathfrak{D}_B < 1$ stay Out, because they understand that all of the guilt types with $\theta_B < 1$ would Keep, independently of their beliefs. Given the heterogeneity of exogenous beliefs, forward induction is enough to imply heterogeneity of behavior and of endogenous beliefs.

Our equilibrium analysis goes beyond these key insights, yielding monotone decision functions and the correlations discussed in the previous subsection. Furthermore, in the model with role-independent guilt, the forward-induction logic does not have such clear-cut implications, because action In may be rationalized by a desire not to disappoint B rather than get a high monetary payoff.

It would be interesting to further depart from traditional equilibrium analysis and explore a rationalizability approach to guilt aversion in social dilemmas whereby some “natural” restrictions on beliefs are taken as given and commonly understood.³⁷ Battigalli et al. (2013) use this approach in the analysis of a cheap-talk sender–receiver game.

Acknowledgments

The authors thank two anonymous referees and the anonymous associate editor for helpful comments conducive to a substantially improved draft. The authors also thank Nicodemo De Vito, Martin Dufwenberg, Alejandro Francetich, Astrid Gamba, Arsen Palestini, Cintia Retz Lucci, Marco Scarsini, and Severine Toussaert, as well as seminar participants in Toulouse, the Catholic University of Milan, the Polytechnic University of Milan, and Sapienza University of Rome for their comments. The first author gratefully acknowledges the financial support of the European Research Council (ERC) [Starting Grant DU 283953] and of the University of Strasbourg through “Attractivité” IdEx 2013. The second author gratefully acknowledges the

financial support of the ERC [Advanced Grant 324219]. The third author gratefully acknowledges the financial support of PRIN 2010–2011 “New approaches to political economy: positive political theories, empirical evidence, and experiments in laboratory.”

Appendix

Proof of Proposition 1

The statement is implied by the following claims, which hold for every nondegenerate Bayesian equilibrium (σ_A, σ_B) with endogenous belief functions $\alpha_A, \alpha_B, \beta_B^S$, and β_B^I .

CLAIM 1. For all $e_B \in \mathcal{E}_B, \beta_B^I(e_B) \geq \frac{1}{2}$.

PROOF. In a nondegenerate equilibrium, $\sigma_A(t_A) = I$ for a set of types with positive measure. By assumption, for each e_B, F_{e_B} has full support; hence $\mathbb{P}_{e_B}[\sigma_A = I] > 0$ and $\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | \sigma_A = I]$ are well defined. Because $\sigma_A(t_A) = I$ only if $\alpha_A(t_A) \geq \frac{1}{2}$, then

$$\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | \sigma_A = I] \geq \mathbb{E}_{e_B}[\alpha_A | \alpha_A \geq \frac{1}{2}] \geq \frac{1}{2}. \quad \square$$

CLAIM 2. For every $(\theta_B, e_B) \in T_B$ with $\theta_B > 2, \sigma_B(\theta_B, e_B) = S$.

PROOF. By Claim 1, $\beta_B^I(e_B) \geq \frac{1}{2}$; therefore, $\theta_B > 2$ implies $2 > 4 - 2\theta_B\beta_B^I(e_B)$ and $\sigma_B(\theta_B, e_B) = S. \quad \square$

CLAIM 3. For every $t_A > \bar{t}_A, \sigma_A(t_A) = I$.

PROOF. By definition of \bar{t}_A , if $t_A > \bar{t}_A$ then $\tau_A(t_A)[\mathfrak{D}_B > 2] > 1/2$. By the stochastic-order assumption (6), it follows that

$$t_A > \bar{t}_A \Rightarrow \alpha_A(t_A) = \tau_A(t_A)[\sigma_B = S] \geq \tau_A(t_A)[\mathfrak{D}_B > 2] > \frac{1}{2}.$$

Therefore, $\sigma_A(t_A) = I$ for every $t_A > \bar{t}_A. \quad \square$

CLAIM 4. Decision function σ_B is monotone in θ_B with threshold $\hat{\theta}(e_B) = 1/\beta_B^I(e_B) \leq 2$.

PROOF. Fix e_B arbitrarily. By incentive condition (12),

$$\sigma_B(\theta_B, e_B) = \begin{cases} K & \text{if } \theta_B < \hat{\theta}(e_B), \\ S & \text{otherwise,} \end{cases}$$

where the threshold $\hat{\theta}(e_B) = 1/\beta_B^I(e_B)$. By Claim 1, $\hat{\theta}(e_B) \leq 2. \quad \square$

CLAIM 5. Decision function σ_A is monotone with threshold $\hat{t}_A \in [\bar{t}_A, \bar{t}_A]$, which is the unique solution to equation

$$\int_{[0,1]} G_{\hat{t}_A}(1/\beta_B^I(e_B)) dF(e_B) = \frac{1}{2}. \quad (16)$$

PROOF. By Claim 4,

$$\begin{aligned} \alpha_A(t_A) &= \tau_A(t_A) \left[\mathfrak{D}_B > \frac{1}{\beta_B^I(e_B)} \right] \\ &= \int_{[0,1]} (1 - G_{t_A}(1/\beta_B^I(e_B))) dF(e_B). \end{aligned}$$

Therefore, by the stochastic-order assumption (6), $t'_A < t''_A$ implies

$$\begin{aligned} \alpha_A(t'_A) &= \int_{[0,1]} (1 - G_{t'_A}(1/\beta_B^I(e_B))) dF(e_B) \\ &< \int_{[0,1]} (1 - G_{t''_A}(1/\beta_B^I(e_B))) dF(e_B) = \alpha_A(t''_A). \end{aligned}$$

³⁶ As we did in §§2–5, we use “Share,” “Keep,” “Out,” and “In” both as action labels and as words in the natural language.

³⁷ See, e.g., Battigalli and Siniscalchi (2003) for such an approach to incomplete-information games with standard preferences.

Because $\alpha_A(\hat{t}_A) \leq 1/2$ and $\alpha_A(\bar{t}_A) \geq 1/2$, threshold \hat{t}_A is the unique solution $\hat{t}_A \in [\bar{t}_A, \hat{t}_A]$ to Equation (16). Incentive condition (11) implies that σ_A is monotone with threshold \hat{t}_A . \square

Claim 5 implies that

$$\beta_B^{\sigma}(e_B) = 1 - \int_{[0,1]} \int_{[0,1]} G_{t_A}(1/\beta_B^I(x)) dF(x) dF_{e_B}(t_A).$$

CLAIM 6. β_B^I is strictly increasing.

PROOF. By Claim 5 and the assumption that each $F_{e_B}(\cdot)$ is continuous and strictly increasing on $[0, 1]$,

$$\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | \sigma_A = I] = \mathbb{E}_{e_B}[\alpha_A | t_A > \hat{t}_A],$$

where \hat{t}_A is the threshold of Claim 5 and $\mathbb{E}_{e_B}[\alpha_A | t_A > \hat{t}_A]$ is strictly increasing in e_B by Equation (10). \square

CLAIM 7. Decision function σ_B is monotone in e_B with threshold function $\hat{e}(\theta_B) = (\beta_B^I)^{-1}(1/\theta_B)$.

PROOF. By Claims 4 and 6, $\beta_B^I(e_B) = 1/\hat{\theta}_B(e_B)$ is increasing, hence invertible. Thus, incentive condition (12) implies that σ_B is monotone in e_B with threshold function $\hat{e}(\theta_B) = (\beta_B^I)^{-1}(1/\theta_B)$. \square

CLAIM 8. The endogenous first-order belief of B is

$$\alpha_B(e_B) = 1 - F_{e_B}(\hat{t}_A) > 0.$$

PROOF. Claim 5 implies that

$$\alpha_B(e_B) = 1 - F_{e_B}(\hat{t}_A).$$

The fact that $\hat{t}_A \leq \bar{t}_A$, together with the assumptions that $\bar{t}_A < 1$ and that each $F_{e_B}(\cdot)$ is continuous and strictly increasing on $[0, 1]$ imply that $\alpha_B(e_B) > 0$. \square

Proof of Proposition 2

We must show that the candidate equilibrium where $\sigma_A(\theta^L, e_A) = O$, $\sigma_A(\theta^H, e_A) = I$, and $\sigma_B(\theta_B, e_B) = K$ for all e_A , θ_B , and e_B satisfies the incentive constraints. These decision functions imply that $\alpha_A(e_A) = 0$ and $\alpha_B(e_B) = \mathbb{P}_{e_B}[\mathfrak{D}_A = \theta^H] = e_B$ for all e_A and e_B . Therefore,

$$\beta_B^I(e_B) = \beta_B^{\sigma}(e_B) = 0,$$

and

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mathbb{E}_{e_A}[\mathbb{E}_{e_B}[\mathbf{m}_B] - m_B(O)] = \mathbb{E}_{e_A}[1 - \alpha_B + 4\alpha_B - 1] \\ &= 3\mathbb{E}_{e_A}[\alpha_B] = 3\mathbb{E}[e_B] > 1, \end{aligned}$$

where we used the definition of β_B^I and Bayes' rule, the definition of $\bar{\beta}_A$, Equation (13), and condition (15).

The beliefs equations for A and the fact that $\theta^L = 0$ and $\theta^H > 1$ imply that incentive condition (11) holds for every type of A, whereas incentive condition (12) holds for every type of B because $\beta_B^I = 0$. \square

Proof of Proposition 3

We start from the conjecture that a strictly positive fraction of A types choose I and provide a characterization of the equilibria that verify this property. We analyze the equilibrium decision functions $\sigma_A(\theta^n, \cdot): [0, 1] \rightarrow \{I, O\}$, $\sigma_B(\theta^n, \cdot): [0, 1] \rightarrow \{S, K\}$, with $n = H, L$, and we show that $\sigma_A(\theta^n, \cdot)$ is monotone increasing and $\sigma_B(\theta^n, \cdot)$ is monotone decreasing. We also provide a characterization of some properties of the endogenous beliefs. We do so by proceeding through a series of claims.

The observation that low-guilt types of B choose K (Remark 2) is used to prove that $\sigma_A(\theta^L, \cdot)$ is monotone (Claim 9). Then, we show that in every nondegenerate equilibrium A's expectation of B's disappointment, $\bar{\beta}_A$, is strictly positive (Claim 10). Next, we prove that $\bar{\beta}_A > 0$ implies that B believes a high-guilt A to be strictly more likely to go In than a low-guilt A (Claim 11). We use this implication to characterize the monotonicity of the endogenous beliefs of B (Claims 12 and 13), and to prove that $\sigma_B(\theta^H, \cdot)$ is monotone decreasing (Claim 13). Next, we prove that $\sigma_A(\theta^H, \cdot)$ is monotone (Claim 14).

Remark 2, Claim 9, and Claims 12–14 taken together lead to Proposition 3.

REMARK 2. Player B's incentive condition, characterized by Equation (12) also in this model, implies that $\sigma_B(\theta^L, e_B) = K$ for every e_B .

Let μ_A be the common marginal belief of each type of A about the epistemic type of B, and let $E_B^{HS} = \{e_B: \sigma_B(\theta^H, e_B) = S\}$. The next claim shows that both A's initial first-order belief and her θ^L decision function are increasing in her epistemic type.

CLAIM 9. For every e_A ,

$$\begin{aligned} \alpha_A(e_A) &= e_A \mu_A[E_B^{HS}], \\ \sigma_A(\theta^L, e_A) &= \begin{cases} I & \text{if } e_A > \hat{e}_A^L, \\ O & \text{otherwise,} \end{cases} \end{aligned}$$

where $\hat{e}_A^L = \min\{1, 1/(2\mu_A[E_B^{HS}])\} \in [\frac{1}{2}, 1]$.

PROOF. Remark 2 and B's incentive condition imply

$$[\sigma_B = S] = \{(\theta_B, e_B): \theta_B = \theta^H, \theta^H \beta_B^I(e_B) > 1\}.$$

By assumption, $\tau_A(e_A)[\mathfrak{D}_B = \theta^H \cap e_B \leq y] = e_A F(y)$ for each y . Therefore,

$$\begin{aligned} \alpha_A(e_A) &= \tau_A(e_A)[\sigma_B = S] = \tau_A(e_A) \left[\mathfrak{D}_B = \theta^H \cap \beta_B^I > \frac{1}{\theta^H} \right] \\ &= e_A \mu_A[E_B^{HS}], \end{aligned}$$

which is increasing in e_A . The incentive condition (11) for A when the guilt type is low implies that $\hat{e}_A^L = \min\{1, 1/(2\mu_A[E_B^{HS}])\}$; note that $\hat{e}_A^L \in [\frac{1}{2}, 1]$. \square

Next note that, in a nondegenerate equilibrium, A necessarily expects to disappoint B by going Out. We formally state this in the following claim.

CLAIM 10. In every nondegenerate equilibrium, $\bar{\beta}_A(e_A) > 0$ for each e_A .

PROOF. In a nondegenerate equilibrium a positive fraction of A types go In; i.e., the set

$$\{e_A: \sigma_A(\theta^L, e_A) = I\} \cup \{e_A: \sigma_A(\theta^H, e_A) = I\}$$

has a positive Lebesgue measure. Let μ_B denote the probability measure on $\mathcal{E}_A = [0, 1]$ induced by cdf F , an exogenous marginal belief of player B . To ease notation, let

$$E_A^{LI} = \{e_A: \sigma_A(\theta^L, e_A) = I\},$$

$$E_A^{HI} = \{e_A: \sigma_A(\theta^H, e_A) = I\}.$$

A positive fraction of A types go In and μ_B has full support; therefore, $\mu_B[E_A^{LI}] + \mu_B[E_A^{HI}] > 0$. Hence, each epistemic type $e_B \in (0, 1)$ expects A to go In with positive probability

$$\begin{aligned} \alpha_B(e_B) &= \tau_B(e_B)(\sigma_A = I \mid \mathfrak{D}_A = \theta^L) \tau_B(e_B)(\mathfrak{D}_A = \theta^L) \\ &\quad + \tau_B(e_B)(\sigma_A = I \mid \mathfrak{D}_A = \theta^H) \tau_B(e_B)(\mathfrak{D}_A = \theta^H) \\ &= \mu_B[E_A^{LI}](1 - e_B) + \mu_B[E_A^{HI}]e_B > 0. \end{aligned}$$

Therefore, for each epistemic type $e_B \in (0, 1)$,

$$\mathbb{E}_{e_B}[\mathbf{m}_B] = 1 \cdot (1 - \alpha_B(e_B)) + 2 \cdot \alpha_B(e_B) > 1.$$

Because $\mu_A[(0, 1)] = 1$, for each e_A ,

$$\tilde{\beta}_A(e_A) = \mathbb{E}_{e_A}[\max\{0, \mathbb{E}_{e_B}[\mathbf{m}_B] - 1\}] > 0. \quad \square$$

CLAIM 11. According to B 's beliefs, a high-guilt A is strictly more likely to go In than a low-guilt A : $\mu_B[E_A^{HI}] > \mu_B[E_A^{LI}]$. Furthermore, whenever the conditional expectations $\mathbb{E}_{e_B}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H]$ and $\mathbb{E}_{e_B}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L]$ are well defined, they are independent of e_B and satisfy

$$\begin{aligned} \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H] &= \frac{1}{\mu_B[E_A^{HI}]} \int_{E_A^{HI}} e_A d\mu_B(e_A) \\ &< \frac{1}{\mu_B[E_A^{LI}]} \int_{E_A^{LI}} e_A d\mu_B(e_A) \\ &= \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L]. \end{aligned}$$

PROOF. $\sigma_A(\theta^L, e_A) = I$ if and only if (iff) $2\mu_A[E_B^{HS}]e_A > 1$, and $\sigma_A(\theta^H, e_A) = I$ iff $2\mu_A[E_B^{HS}]e_A > 1 - \theta^H \tilde{\beta}_A(e_A)$. Note that $\theta^H \tilde{\beta}_A(e_A) > 0$ because $\theta^H > 1$ by assumption, and $\tilde{\beta}_A(e_A) > 0$ by Claim 10. Because μ_B has full support

$$\begin{aligned} \mu_B[E_A^{HI}] - \mu_B[E_A^{LI}] \\ = \mu_B[\{e_A: 1 - \theta^H \tilde{\beta}_A(e_A) < 2\mu_A[E_B^{HS}]e_A \leq 1\}] > 0. \end{aligned}$$

Recall that, according to B 's beliefs, \mathfrak{D}_A and \mathbf{e}_A are independent. Therefore, for every $x \in [0, 1]$,

$$\begin{aligned} \mathbb{P}_{e_B}[\mathbf{e}_A < x \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L] &= \mathbb{P}_{e_B}[\mathbf{e}_A < x \mid \mathbf{e}_A \in E_A^{LI} \cap \mathfrak{D}_A = \theta^L] \\ &= \mathbb{P}[\mathbf{e}_A < x \mid \mathbf{e}_A \in E_A^{LI}], \end{aligned}$$

whenever the conditional probability is well defined (i.e., for $\mu_B[E_A^{LI}] > 0$ and $e_B < 1$). The conditional probability $\mathbb{P}[\mathbf{e}_A < x \mid \mathbf{e}_A \in E_A^{LI}]$ is independent of e_B because it is determined by the common marginal belief μ_B on $\mathcal{E}_A = [0, 1]$ generated by cdf F

$$\mathbb{P}[\mathbf{e}_A < x \mid \mathbf{e}_A \in E_A^{LI}] = \frac{\mu_B[\{e_A \in E_A^{LI}: e_A < x\}]}{\mu_B[E_A^{LI}]}.$$

Similarly

$$\begin{aligned} \mathbb{P}_{e_B}[\mathbf{e}_A < x \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H] \\ = \mathbb{P}[\mathbf{e}_A < x \mid \mathbf{e}_A \in E_A^{HI}] = \frac{\mu_B[\{e_A \in E_A^{HI}: e_A < x\}]}{\mu_B[E_A^{HI}]}, \end{aligned}$$

whenever the conditional probability is well defined (i.e., for $e_B > 0$, since we know that $\mu_B[E_A^{HI}] > 0$). Note that $E_A^{LI} = (\hat{e}_A^L, 1] \subset E_A^{HI} \subseteq [0, 1]$. Therefore, for each $e_B \in (0, 1)$,

$$\begin{aligned} \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H] &= \frac{1}{\mu_B[E_A^{HI}]} \int_{E_A^{HI}} e_A d\mu_B(e_A) \\ &< \frac{1}{\mu_B[E_A^{LI}]} \int_{E_A^{LI}} e_A d\mu_B(e_A) = \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L], \end{aligned}$$

where the second conditional expectation is well defined if $\mu_B[E_A^{LI}] > 0$, i.e., if $\hat{e}_A^L < 1$. \square

CLAIM 12. The endogenous first-order belief of B is

$$\alpha_B(e_B) = \mu_B[E_A^{LI}] + e_B(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]),$$

which is strictly increasing in e_B .

PROOF. The endogenous first-order belief of B is

$$\begin{aligned} \alpha_B(e_B) &= \mathbb{P}[\sigma_A = I] = \mu_B[E_A^{LI}](1 - e_B) + \mu_B[E_A^{HI}]e_B \\ &= \mu_B[E_A^{LI}] + e_B(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]). \end{aligned}$$

Note that α_B is strictly increasing in e_B given that $\mu_B[E_A^{HI}] > \mu_B[E_A^{LI}]$, as shown in Claim 11. \square

CLAIM 13. The endogenous second-order belief of B is such that

$$\beta_B^{\mathcal{G}}(e_B) = \mu_A[E_B^{HS}]\mathbb{E}(\mathbf{e}_A),$$

which is constant, and

$$\begin{aligned} \beta_B^I(e_B) &= \mu_A[E_B^{HS}](\mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^L](1 - e_B) \\ &\quad + \mathbb{E}[\mathbf{e}_A \mid \sigma_A = I \cap \mathfrak{D}_A = \theta^H]e_B), \end{aligned}$$

which is decreasing (strictly, if $\mu_A[E_B^{HS}] > 0$). Moreover, $\sigma_B(\theta^H, \cdot)$ is monotone decreasing; i.e.,

$$\sigma_B(\theta^H, e_B) = \begin{cases} S & \text{if } e_B < \hat{e}_B^H, \\ K & \text{otherwise,} \end{cases}$$

where \hat{e}_B^H satisfies the incentive conditions

$$\begin{aligned} \hat{e}_B^H = 0 &\implies \beta_B^I(\hat{e}_B^H) \leq \frac{1}{\theta^H}, \\ \hat{e}_B^H \in (0, 1) &\implies \beta_B^I(\hat{e}_B^H) = \frac{1}{\theta^H}, \\ \hat{e}_B^H = 1 &\implies \beta_B^I(\hat{e}_B^H) \geq \frac{1}{\theta^H}. \end{aligned}$$

PROOF. The endogenous second-order belief of B is independent of e_B because, by assumption, α_A depends only on e_A and each type of B has the same marginal belief μ_B (the measure generated by cdf F) on $\mathcal{E}_B = [0, 1]$. Specifically

$$\beta_B^{\mathcal{G}}(e_B) = \mathbb{E}_{e_B}[\alpha_A] = \mathbb{E}_{e_B}(\mu_A[E_B^{HS}]\mathbf{e}_A) = \mu_A[E_B^{HS}]\mathbb{E}(\mathbf{e}_A).$$

Given that $\beta_B^I(e_B) = \mathbb{E}_{e_B}[\alpha_A | \sigma_A = I]$ and using Claims 9 and 11, we obtain

$$\begin{aligned} \beta_B^I(e_B) &= \mu_A[E_B^{HS}] \mathbb{E}_{e_B}[\mathbf{e}_A | \sigma_A = I] \\ &= \mu_A[E_B^{HS}] (\mathbb{E}[\mathbf{e}_A | \sigma_A = I \cap \mathfrak{D}_A = \theta^L] (1 - e_B) \\ &\quad + \mathbb{E}[\mathbf{e}_A | \sigma_A = I \cap \mathfrak{D}_A = \theta^H] e_B), \end{aligned}$$

whenever the conditional probabilities are well defined. Therefore, $\beta_B^I(\cdot)$ is decreasing in e_B , given that

$$\begin{aligned} \frac{\partial \beta_B^I(e_B)}{\partial e_B} &= \mu_A[E_B^{HS}] (\mathbb{E}[\mathbf{e}_A | \sigma_A = I \cap \mathfrak{D}_A = \theta^H] \\ &\quad - \mathbb{E}[\mathbf{e}_A | \sigma_A = I \cap \mathfrak{D}_A = \theta^L]) \leq 0, \end{aligned}$$

by Claim 11 (note that $\mu_A[E_B^{HS}]$ may be zero). The incentive condition (12) for the high-guilt type of B implies that he chooses S iff $\beta_B^I(e_B) > 1/\theta^H$. Therefore, B 's decision function $\sigma_B(\theta^H, \cdot)$ is monotone decreasing in e_B , and characterized by a threshold \hat{e}_B^H that satisfies the incentive conditions stated in this claim. \square

CLAIM 14. *The endogenous second-order belief of A is such that*

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mu_B[E_A^{LI}] (3 - 2\mu_A[E_B^{HS}] e_A) \\ &\quad + 3(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \mathbb{E}[\mathbf{e}_B] \\ &\quad - 2(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) e_A \mu_A[E_B^{HS}] \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H], \end{aligned}$$

which is decreasing. Moreover $\sigma_A(\theta^H, \cdot)$ is monotone (increasing); i.e.,

$$\sigma_A(\theta^H, e_A) = \begin{cases} I & \text{if } e_A \geq \hat{e}_A^H, \\ O & \text{otherwise,} \end{cases}$$

where $\hat{e}_A^H \in [0, 1)$ satisfies the following incentive conditions

$$\begin{aligned} \hat{e}_A^H = 0 &\implies 1 - \theta^H \bar{\beta}_A(\hat{e}_A^H) \leq 2\alpha_A(e_A), \\ \hat{e}_A^H > 0 &\implies 1 - \theta^H \bar{\beta}_A(\hat{e}_A^H) = 2\alpha_A(e_A). \end{aligned}$$

PROOF. Remember that B 's disappointment depends on whether B plans to choose S or K after I . Therefore, A 's expectation of B 's disappointment, $\bar{\beta}_A(e_A)$, depends on whether A expects B to choose S or K

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mathbb{E}_{e_A}[3\alpha_B | \sigma_B = K] \mathbb{P}_{e_A}[\sigma_B = K] \\ &\quad + \mathbb{E}_{e_A}[\alpha_B | \sigma_B = S] \mathbb{P}_{e_A}[\sigma_B = S]. \end{aligned}$$

Decomposing expected values and taking into account that $\mathbb{P}[\sigma_B = S | \mathfrak{D}_B = \theta^L] = 0$, we obtain

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mathbb{E}[3\alpha_B | \sigma_B = K \cap \mathfrak{D}_B = \theta^L] \mathbb{P}[\sigma_B = K | \mathfrak{D}_B = \theta^L] \mathbb{P}_{e_A}[\mathfrak{D}_B = \theta^L] \\ &\quad + \mathbb{E}[3\alpha_B | \sigma_B = K \cap \mathfrak{D}_B = \theta^H] \mathbb{P}[\sigma_B = K | \mathfrak{D}_B = \theta^H] \mathbb{P}_{e_A}[\mathfrak{D}_B = \theta^H] \\ &\quad + \mathbb{E}[\alpha_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H] \mathbb{P}[\sigma_B = S | \mathfrak{D}_B = \theta^H] \mathbb{P}_{e_A}[\mathfrak{D}_B = \theta^H]. \end{aligned}$$

Replacing probabilities with their specific expressions and using Claim 12, we obtain

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mu_B[E_A^{LI}] (3(1 - e_A) + 3(1 - \mu_A[E_B^{HS}]) e_A + e_A \mu_A[E_B^{HS}]) \\ &\quad + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) 3 \mathbb{E}[\mathbf{e}_B] (1 - e_A) \\ &\quad + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) e_A 3(1 - \mu_A[E_B^{HS}]) \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = K \cap \mathfrak{D}_B = \theta^H] \\ &\quad + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) e_A \mu_A[E_B^{HS}] \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]. \end{aligned}$$

Now observe that, since the random variables \mathbf{e}_B and \mathfrak{D}_B are independent, we can write

$$\begin{aligned} \mathbb{E}[\mathbf{e}_B] &= \mathbb{E}[\mathbf{e}_B | \mathfrak{D}_B = \theta^H] \\ &= (1 - \mu_A[E_B^{HS}]) \mathbb{E}[\mathbf{e}_B | \sigma_B = K \cap \mathfrak{D}_B = \theta^H] \\ &\quad + \mu_A[E_B^{HS}] \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]. \end{aligned}$$

Regrouping terms in the expression of $\bar{\beta}_A(e_A)$, this simplifies to

$$\begin{aligned} \bar{\beta}_A(e_A) &= \mu_B[E_A^{LI}] (3 - 2\mu_A[E_B^{HS}] e_A) + 3(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \mathbb{E}[\mathbf{e}_B] \\ &\quad - 2(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) e_A \mu_A[E_B^{HS}] \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]. \end{aligned}$$

Note that $\bar{\beta}_A(\cdot)$ is decreasing

$$\begin{aligned} \frac{\partial \bar{\beta}_A}{\partial e_A} &= -2\mu_A[E_B^{HS}] (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]) \leq 0. \end{aligned}$$

This completes the proof of the first part of the claim.

To show that $\sigma_A(\theta^H, \cdot)$ is monotone, we consider A 's incentive condition; i.e., type (θ^H, e_A) of A chooses I when

$$2\alpha_A + \theta^H \bar{\beta}_A(e_A) > 1.$$

Recall that Claim 9 shows that $\alpha_A = \mu_A[E_B^{HS}] e_A$; hence, the incentive condition can be rewritten as

$$2\mu_A[E_B^{HS}] e_A + \theta^H \bar{\beta}_A(e_A) > 1.$$

Next, we show that either (i) the left-hand side (LHS) is increasing in e_A , hence $\sigma_A(\theta^H, \cdot)$ is monotone (increasing) or constant; or (ii) the LHS is larger than 1, hence $\sigma_A(\theta^H, \cdot)$ is constant at I . Differentiating the LHS and using the expression for $\partial \bar{\beta}_A / \partial e_A$, we obtain

$$\begin{aligned} 2\mu_A[E_B^{HS}] + \theta^H \frac{\partial \bar{\beta}_A}{\partial e_A} &= 2\mu_A[E_B^{HS}] (1 - \theta^H (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H])). \end{aligned}$$

Therefore, the LHS is increasing iff

$$\theta^H (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]) \leq 1.$$

Suppose the LHS is strictly decreasing, i.e.,

$$\begin{aligned} \theta^H (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \\ \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]) > 1. \end{aligned} \quad (17)$$

Note that, by Claim 13, $\mathbb{E}[\mathbf{e}_B] \geq \mathbb{E}[\mathbf{e}_B | \mathbf{e}_B < \hat{e}_B^H] = \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]$; therefore, we obtain the following inequalities, which imply that the LHS is larger than 1:

$$\theta^H \bar{\beta}_A(e_A) \geq \theta^H (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]) (3 - 2\mu_A[E_B^{HS}]e_A) > 1.$$

In particular, the first inequality holds because the expression for $\bar{\beta}_A(e_A)$ and the fact that $\mathbb{E}[\mathbf{e}_B] \geq \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]$ imply that

$$\begin{aligned} \theta^H \bar{\beta}_A(e_A) &\geq \mu_B[E_A^{LI}] (3 - 2\mu_A[E_B^{HS}]e_A) + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \\ &\quad \cdot 3\mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H] \\ &\quad - 2(\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}])e_A\mu_A^H \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H] \\ &= (\mu_B[E_A^{LI}] + (\mu_B[E_A^{HI}] - \mu_B[E_A^{LI}]) \\ &\quad \cdot \mathbb{E}[\mathbf{e}_B | \sigma_B = S \cap \mathfrak{D}_B = \theta^H]) (3 - 2\mu_A[E_B^{HS}]e_A). \end{aligned}$$

The second inequality holds by Equation (17) and because $(3 - 2\mu_A[E_B^{HS}]e_A) > 1$.

Therefore,

$$\sigma_A(\theta^H, e_A) = \begin{cases} I & \text{if } e_A \geq \hat{e}_A^H, \\ O & \text{otherwise,} \end{cases}$$

where $\hat{e}_A^H \in [0, 1)$ satisfies the incentive conditions

$$\begin{aligned} \hat{e}_A^H = 0 &\Rightarrow 1 - \theta^H \bar{\beta}_A(\hat{e}_A^H) \leq 2\alpha_A(e_A), \\ \hat{e}_A^H > 0 &\Rightarrow 1 - \theta^H \bar{\beta}_A(\hat{e}_A^H) = 2\alpha_A(e_A). \quad \square \end{aligned}$$

Given that the decision functions $\sigma_B(\theta^H, \cdot)$, $\sigma_A(\theta^L, \cdot)$, and $\sigma_A(\theta^H, \cdot)$ are monotone and described respectively by thresholds \hat{e}_B^H , \hat{e}_A^L , and \hat{e}_A^H , we have that

$$\begin{aligned} \mu_A[E_B^{HS}] &= F(\hat{e}_B^H), \\ \mu_B[E_A^{LI}] &= 1 - F(\hat{e}_A^L), \\ \mu_B[E_A^{HI}] &= 1 - F(\hat{e}_A^H). \end{aligned}$$

This, together with Remark 2, Claim 9, and Claims 12–14, delivers the result stated in Proposition 3. \square

References

Attanasi G, Nagel R (2008) A survey of psychological games: Theoretical findings and experimental evidence. Innocenti A, Sbriglia P, eds. *Games, Rationality and Behavior: Essays on Behavioral Game Theory and Experiments* (Palgrave Macmillan, London), 204–232.

Attanasi G, Battigalli P, Nagel R (2013) Disclosure of belief-dependent preferences in the Trust Game. IGIER Working Paper 506, Bocconi University, Milan.

Barrett LF (2006) Solving the emotion paradox: Categorization and the experience of emotion. *Personality Soc. Psych. Rev.* 10:20–46.

Battigalli P, Dufwenberg M (2007) Guilt in games. *Amer. Econom. Rev. Papers Proc.* 97:170–176.

Battigalli P, Dufwenberg M (2009) Dynamic psychological games. *J. Econom. Theory* 144:1–35.

Battigalli P, Siniscalchi M (2003) Rationalization and incomplete information. *B.E. J. Theoret. Econom.* 3:Article 3.

Battigalli P, Charness G, Dufwenberg M (2013) Deception: The role of guilt. *J. Econom. Behav. Organ.* 93:227–232.

Battigalli P, Dufwenberg M, Smith A (2015) Frustration and anger in games. IGIER Working paper 539, Bocconi University, Milan.

Bellemare C, Sebald A, Strobel M (2011) Measuring the willingness to pay to avoid guilt: Estimation using equilibrium and stated belief models. *J. Appl. Econometrics* 26:437–453.

Berg J, Dickhaut J, McCabe K (1995) Trust, reciprocity, and social-history. *Games Econom. Behav.* 10:122–142.

Binmore K, Gale J, Samuelson L (1995) Learning to be imperfect: The ultimatum game. *Games Econom. Behav.* 8:56–90.

Brandenburger A, Dekel E (1987) Rationalizability and correlated equilibria. *Econometrica* 55:1391–1402.

Buskens V, Raub W (2013) Rational choice research on social dilemmas: Embeddedness effects on trust. Wittek R, Snijders TAB, Nee V, eds. *Handbook of Rational Choice Social Research* (Stanford University Press, Redwood City, CA), 113–150.

Caplin A, Leahy J (2004) The supply of information by a concerned expert. *Econom. J.* 114:487–505.

Carter CS, Keverne EB (2002) The neurobiology of social affiliation and pair bonding. Pfaff D, Arnold AP, Etgen AM, Fahrbach SE, Rubin RT, eds. *Hormones, Brain, and Behavior* (Academic Press, San Diego), 299–337.

Chang LJ, Smith A, Dufwenberg M, Sanfey A (2011) Triangulating the neural, psychological and economic bases of guilt aversion. *Neuron* 70:560–572.

Charness G, Dufwenberg M (2006) Promises and partnership. *Econometrica* 74:1579–1601.

Charness G, Dufwenberg M (2011) Participation. *Amer. Econom. Rev.* 101:1213–1239.

Cooper D, Kagel J (2013) Other-regarding preferences: A selective survey of experimental results. Kagel J, Roth A, eds. *Handbook of Experimental Economics*, Vol. 2 (Princeton University Press, Princeton, NJ).

Dekel E, Siniscalchi M (2015) Epistemic game theory. Young P, Zamir S, eds. *Handbook of Game Theory*, Vol. 4 (North Holland, Amsterdam), 619–702.

Dekel E, Fudenberg D, Levine D (2004) Learning to play Bayesian games. *Games Econom. Behav.* 46:282–303.

Dufwenberg M (2002) Marital investment, time consistency and emotions. *J. Econom. Behav. Organ.* 48:57–69.

Dufwenberg M (2008) Psychological games. Durlauf SN, Blume LE, eds. *The New Palgrave Dictionary of Economics*, 2nd ed. Vol. 6 (Palgrave Macmillan, London), 714–718.

Dufwenberg M, Gneezy U (2000) Measuring beliefs in an experimental lost wallet game. *Games Econom. Behav.* 30:163–182.

Ellingsen T, Johannesson M, Tjotta S, Torsvik G (2010) Testing guilt aversion. *Games Econom. Behav.* 68:95–107.

Geanakoplos J, Pearce D, Stacchetti E (1989) Psychological games and sequential rationality. *Games Econom. Behav.* 1:60–79.

Gneezy U (2005) Deception: The role of consequences. *Amer. Econom. Rev.* 95:384–394.

Guerra G, Zizzo DJ (2004) Trust responsiveness and beliefs. *J. Econom. Behav. Organ.* 55:25–30.

Harsanyi JC (1967) Games of incomplete information played by Bayesian players. Part I: The basic model. *Management Sci.* 14(3):159–182.

Harsanyi JC (1968a) Games of incomplete information played by Bayesian players. Part II, Bayesian equilibrium points. *Management Sci.* 14(5):320–334.

Harsanyi JC (1968b) Games of incomplete information played by Bayesian players. Part III: The basic probability distribution of the game. *Management Sci.* 14(7):486–502.

Haselton MG, Ketelaar T (2006) Irrational emotions or emotional wisdom? The evolutionary psychology of emotions and behavior. Forgas J, ed. *Hearts and Minds: Affective Influences on Social Cognition and Behavior*, Frontiers of Psychology Series (Psychology Press, New York), 21–40.

- Kosfeld M, Heinrichs M, Zak P, Fischbacher U, Fehr E (2005) Oxytocin increases trust in humans. *Nature* 435:673–676.
- Ong D (2011) Fishy gifts: Bribing with shame and guilt. Working paper, Peking University, Shenzhen, China.
- Reuben E, Sapienza P, Zingales L (2009) Is mistrust self-fulfilling? *Econom. Lett.* 104:89–91.
- Ross L, Greene D, House P (1977) The false consensus effect: An egocentric bias in social perception and attribution processes. *J. Experiment. Soc. Psych.* 13:279–301.
- Shaked M, Shantikumar JG (2007) *Stochastic Orders* (Springer, New York).
- Tadelis S (2011) The power of shame and the rationality of trust. Working paper, University of California, Berkeley.
- Vanberg C (2008) Why do people keep their promises? An experimental test of two explanations. *Econometrica* 76:1467–1480.
- Zak PJ (2008) The neurobiology of trust. *Scientific Amer.* 298:88–95.
- Zak PJ, Kurzban R, Matzner WT (2005) Oxytocin is associated with human trustworthiness. *Hormones Behav.* 48:522–527.