

Dynamic psychological games

Pierpaolo Battigalli ^{a,*}, Martin Dufwenberg ^b

^a *Bocconi University, Via Sarfatti 25, 20136 Milan, Italy*

^b *University of Arizona, Tucson, AZ, USA*

Received 30 November 2007; accepted 31 January 2008

Available online 17 May 2008

Abstract

The motivation of decision makers who care for various emotions, intentions-based reciprocity, or the opinions of others may depend directly on beliefs (about choices, beliefs, or information). Geanakoplos, Pearce and Stacchetti [J. Geanakoplos, D. Pearce, E. Stacchetti, Psychological games and sequential rationality, *Games Econ. Behav.* 1 (1989) 60–79] point out that traditional game theory is ill-equipped to address such matters, and they pioneer a new framework which does. However, their toolbox – psychological game theory – incorporates several restrictions that rule out plausible forms of belief-dependent motivation. Building on recent work on dynamic interactive epistemology, we propose a more general framework. Updated higher-order beliefs, beliefs of others, and plans of action may influence motivation, and we can capture dynamic psychological effects (such as sequential reciprocity, psychological forward induction, and regret) that were previously ruled out. We develop solution concepts, provide examples, explore properties, and suggest avenues for future research.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C72; C73

Keywords: Psychological games; Belief-dependent motivation; Extensive-form solution concepts; Dynamic interactive epistemology

* Corresponding author. Fax: +39 02 5836 3332.

E-mail address: pierpaolo.battigalli@unibocconi.it (P. Battigalli).

1. Introduction

We develop a framework for analyzing strategic interaction when players have ‘belief-dependent’ motivations, generalizing the theory of extensive form *psychological games* proposed by Geanakoplos, Pearce and Stacchetti [36] (hereafter GPS, see also Gilboa and Schmeidler [37]). The rest of this Introduction motivates in more detail.

Traditional game theory is not a rich enough toolbox to adequately describe many psychological or social aspects of motivation and behavior. The traditional approach assumes utilities depend only on which actions are chosen, but if decision makers are emotional or care for the intentions, opinions, or emotions of others utilities may depend also on which beliefs (about choices, beliefs, or information) players harbor. The following examples illustrate:

1. When Abi takes a taxi ride she tips as much as she expects that the driver (Ben) expects to get. She suffers from guilt if she tips less.
2. Cleo suddenly pushes Dan over. Should Dan splash a bucket of water over Cleo in return? Maybe she actually tried to hug him? If so, Dan would rather forgive (maybe even hug) Cleo.
3. Eva is unemployed. Her neighbor, Fred, observes the effort with which she tries to get a job. Fred’s taxes pay for Eva’s unemployment benefits, so Eva’s choice has externalities the size of which depends on her talent translating effort to probability of getting a job (low effort is costlier to Fred if Eva is talented and could have gotten a job had she tried harder). Eva’s talent is known only to her, but Fred makes inferences observing her effort. This determines the social respect he bestows on Eva, and since she cares about respect this influences her effort.
4. Gwen is anxious about her health. Hal, her doctor, has diagnosed a serious illness. He is concerned with Gwen’s health and anxiety. Should he prescribe the most appropriate treatment and thus reveal to Gwen how bad is her situation?

Abi’s tip, Dan’s hug/soak choice, Eva’s effort, and Hal’s prescription each pins down an outcome. Yet the preferred choice depends on a belief.¹

The point that belief-dependent motivation may be important for strategic decision making is made by GPS, who present several intriguing examples involving various emotions. They show the inadequacy of traditional methods to represent the involved preferences, and develop an extension (in the normal as well as in the extensive form) of traditional game theory to deal with the matter. A literature has emerged which either draws on or which can be related to GPS’ framework. Contributions include applications to specific economic problems,² models of certain forms of belief-dependent preferences (like reciprocity or guilt),³ and experimental studies that

¹ Abi’s preference depends on her belief of Ben’s belief about the tip; Dan’s on his assessment of Cleo’s intentions; Eva’s on Fred’s inferences on her talent; Hal’s on his belief about Gwen’s anxiety – a belief-induced psychological state – and on his belief about Gwen’s sensitivity to anxiety.

² See Huang and Wu [42], Dufwenberg [24,25], Geanakoplos [35], Ruffle [71], Dufwenberg and Kirchsteiger [28], Huck and Kübler [43], Caplin and Eliaz [20], Caplin and Leahy [22], Li [55], Kőszegi [48], Kőszegi and Rabin [49,50], Sebald [72], Bernheim [16] and Dufwenberg and Lundholm [30] also consider belief-dependent motivations although the authors do not draw a connection to psychological games.

³ See Rabin [67], Dufwenberg and Kirchsteiger [29], Falk and Fischbacher [31], Battigalli and Dufwenberg [8], Segal and Sobel [73]; see Sobel [75] for some related discussion.

have provided support for such models.⁴ By now a quite large set of economists argue that belief-dependent motivation is relevant to economic behavior.⁵

While GPS' paper is a source of inspiration for all work on belief-dependent motivation, and an applicable toolbox for some work, a careful scrutiny reveals that their approach is too restrictive to handle many plausible forms of belief-dependent motivation (this is acknowledged by GPS themselves; see pp. 70, 78–79). There are several reasons including the following four:

- R1** (*updated beliefs*): GPS only allow *initial* beliefs to enter the domain of a player's utility, while many seemingly important forms of belief-dependent motivation require *updated* beliefs to matter.
- R2** (*others' beliefs*): GPS only allow a player's *own* beliefs to enter the domain of his utility function, while there are conceptual and technical reasons to let *others'* beliefs matter.
- R3** (*dependence on plans*): GPS follow the traditional extensive games approach of letting strategies influence utilities only insofar as they influence terminal histories, but many forms of belief-dependent motivation become compelling in particular in conjunction with preferences that depend on strategies in ways not captured by terminal histories.
- R4** (*non-equilibrium analysis*): GPS restrict attention to equilibrium analysis, but in many strategic situations there is little compelling reason to expect players to coordinate on an equilibrium and one may wish to explore alternative assumptions.

This list deserves backup by examples, but we postpone this until the next section. Here we just note that items in the list have lead some researchers to deviate from GPS' framework, in developing specific examples or models with belief-dependent motivation. However almost no papers are concerned with developing the overall framework of psychological game theory.⁶ We attempt to fill this gap, using **R1–R4** as guiding principles.

Our approach crucially draws on Battigalli and Siniscalchi's [10] work on how to represent hierarchies of conditional beliefs. This is essential for **R1**, and figures in the background of **R2–R4** which are all related to updated beliefs. We define a large class of psychological games, which contains (in a particular sense) GPS' games and traditional games as special cases. Our main goal is to develop this basic framework, and to illustrate some solution concepts that can be meaningfully developed for it. While one could imagine a variety of interesting solution concepts, we choose to extend two basic concepts of classical game theory to our setting: sequential equilibrium and (extensive form) rationalizability. We prove related theorems, and illustrate how the concepts work in examples.

⁴ See Dufwenberg and Gneezy [27], Guerra and Zizzo [39], Bacharach, Guerra and Zizzo [5], Charness and Dufwenberg [23], Bouckaert and Dhaene [17], Dufwenberg, Gaechter and Hennig-Schmidt [26], Tadelis [77], as well as the survey Attanasi and Nagel [3].

⁵ Add some (not many) decision-theorists to this list: Machina [58,59] presents examples which concern belief-dependent forms of disappointment. Robin Pope has written extensively (since the early 80's) about how conventional decision theory excludes various forms of belief-dependent motivation; Pope [66] expounds her program and give further references and Albers, Pope, Selten and Vogt [1] report on a related experiment. Bell [13], Loomes and Sugden [57], Karni [44], Karni and Schlee [45], and Caplin and Leahy [21] develop single decision-maker models in which utility may depend directly on beliefs.

⁶ Kolpin [47] explores an alternative route to GPS' games, where players 'choose beliefs.' Segal and Sobel [73] analyze simultaneous move games, and assume preferences over material consequences depend on the equilibrium probability distribution over actions. They observe that their approach can be regarded as a reformulation of GPS' normal form games.

R1–R4 do not exhaust the good reasons to generalize GPS, but in the name of pedagogical clarity we only deal with **R1–R4** in most sections of this paper (2–5). Section 2 surveys these conceptual issues. Section 3 develops the general framework, up to the definition of a psychological game. Section 4 concerns sequential equilibrium. Section 5 concerns interactive epistemology and rationalizability. Section 6 contains a discussion as well as extensions beyond **R1–R4**. We compare our approach to that of GPS in more depth, consider imperfect information, chance moves, asymmetric information, own-plan dependence, dynamic inconsistency, and multi-self utility, and finally offer remarks regarding solutions concept we did not develop. Appendix A collects most of the proofs.

2. Overview of some conceptual issues

This section surveys the conceptual issues that motivate us. We first describe what GPS' do, and why this is 'non-standard' *vis-a-vis* traditional game theory (2.1). We then explain what is our own contribution, going through **R1–R4** in more detail (2.2). The style is 'semi-technical,' we introduce some notation, but postpone proper treatment of details for later.

2.1. What GPS do

The traditional approach to analyzing extensive games (with complete information) describes a player's preferences using a utility function of the form

$$u_i : Z \rightarrow \mathbb{R}$$

where Z is the set of terminal histories (endnodes).

Psychological games capture richer motivations than traditional games, and the payoff functions have richer domains. GPS define a set of i 's initial (pre-play) beliefs about others' strategies and initial beliefs, here referred to as $\bar{\mathbf{M}}_i$, which does not rule out any hierarchy of initial beliefs. Each element of $\bar{\mathbf{M}}_i$ is a sequence $\bar{\mu}_i = (\bar{\mu}_i^1, \bar{\mu}_i^2, \dots)$ where $\bar{\mu}_i^1$ represents i 's beliefs about the opponents' strategies, or first-order beliefs, $\bar{\mu}_i^2$ represents i 's joint beliefs about the opponents' strategies and first-order beliefs, and so on.⁷

GPS model preferences using utility functions of the form

$$u_i : Z \times \bar{\mathbf{M}}_i \rightarrow \mathbb{R}.$$

This structure bears some superficial similarities to games of incomplete information. It is worth clarifying the differences. In a game of incomplete information some payoff-relevant exogenous parameters (*e.g.* players' abilities or tastes) are not commonly known. Let $\theta \in \Theta$ denote the vector of such parameters. Players' payoffs are represented by parametrized utility functions $v_i : Z \times \Theta \rightarrow \mathbb{R}$. Note that θ does *not* specify strategic choices. A player has beliefs about θ (comprising her private information about θ), beliefs about the beliefs of others concerning θ , etc. Following Harsanyi [41], such first- and higher-order beliefs can be represented in an elegant, albeit implicit, form by assuming that each player i is characterized by a 'type' $t_i \in T_i$ and each t_i corresponds to a probability measure p_{t_i} over the set of payoff-relevant parameters and opponents' types, *i.e.* $p_{t_i} \in \Delta(\Theta \times T_{-i})$. It can be shown that p_{t_i} corresponds to an infinite

⁷ More formally, first-order beliefs are elements of $\Delta(S_{-i})$ (where S_{-i} is the set of strategy profiles of i 's co-players), second-order beliefs elements of $\Delta(S_{-i} \times \prod_{j \neq i} \Delta(S_{-j}))$, etc. (see [18] and [61]). Upper-bars distinguish initial beliefs from systems of conditional beliefs, the main object of our analysis. We will be more precise in Section 3.

hierarchy of beliefs $(p_{t_i}^1, p_{t_i}^2, \dots)$ where $p_{t_i}^1 \in \Delta(\Theta)$ is the marginal of p_{t_i} on Θ , $p_{t_i}^2$ is a joint belief about θ and the opponents' beliefs about θ , and so on. Taking conditional expectations, the payoff functions of the incomplete information game can be represented as $V_i : Z \times \prod_j T_j \rightarrow \mathbb{R}$, where

$$V_i(z, t_i, t_{-i}) = \int v_i(z, \theta) p_{t_i}(d\theta | t_{-i}).$$

Thus, both psychological games and incomplete information games can be described so that payoffs depend not only on how the game is played ($z \in Z$) but also on hierarchical beliefs. However, we are talking about different beliefs in the two cases. In psychological games payoffs at endnodes depend on beliefs about strategies, beliefs about such beliefs, and so on. The modeler explains/predicts, such beliefs *via* some solution concept. Hence payoffs at a given endnode are *endogenous*. On the other hand, players' hierarchical beliefs about the parameter vector θ are as *exogenous* as θ itself. Hence payoffs at a given terminal history of an incomplete information game are exogenous as well.

2.2. Extension of GPS

GPS' approach can capture interesting forms of belief-dependent motivation. Example 1 of the Introduction, *e.g.*, could be handled by assuming that Abi's utility equals $w - m - 2 \max\{\bar{\mu} - m, 0\}$, where w is her pre-tip wealth, $m \in \{0, 1, \dots, w\}$ her tip, and $\bar{\mu}$ her expectation of Ben's expectation of m , a function of her second-order belief. Abi maximizes her utility by choosing $m = \bar{\mu}$.

However, the issues **R1–R4** lead us to enrich the domain of utilities further. We consider payoff functions of the form

$$u_i : Z \times \mathbf{M}_i \times \prod_{j \neq i} (\mathbf{M}_j \times S_j) \rightarrow \mathbb{R}$$

where \mathbf{M}_j (with $j = i$ or $j \neq i$) is the set of j 's possible *conditional* beliefs about others' strategies and conditional beliefs, S_j is the set of (pure) strategies of j . The conditioning in \mathbf{M}_j is done for every history, building on Battigalli and Siniscalchi [10] who show how to represent hierarchies of conditional beliefs without ruling out any hierarchy. $\bar{\mathbf{M}}_j$ is (isomorphic to) a subspace of \mathbf{M}_j , so the payoff functions we consider are more general than those assumed by GPS.⁸

Issues **R1–R4** will be related to different arguments of u_i as we go.

R1: updated beliefs

Rabin's [67] reciprocity theory, in which a player's preferences over material payoff distributions depends on the co-players intentions, is perhaps the most well-known application of GPS' theory. Rabin works in the normal form. His goal is to highlight key qualitative features of reciprocity, and he does not address issues of dynamic decision making although he points out that this is important for applied work (p. 1296).⁹ Dufwenberg and Kirchsteiger [29] pick up from there, and develop a reciprocity theory for extensive games. In motivating their exercise, they argue that it is necessary to deviate from GPS' extensive form framework: GPS only allow initial beliefs to enter the domain of a player's utility, while the modeling of reciprocal response at

⁸ For a more precise comparison between our framework and GPS see Subsection 6.1.

⁹ Rabin [68] (see p. 23 and footnote 16) also provides a hilarious autobiographical sequential-move example reminiscent of our Example 2 of the Introduction.

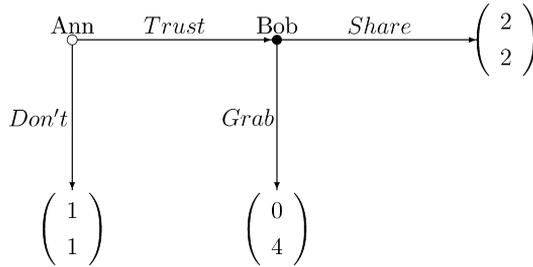


Fig. 1. Trust game Γ_1 with material payoffs.

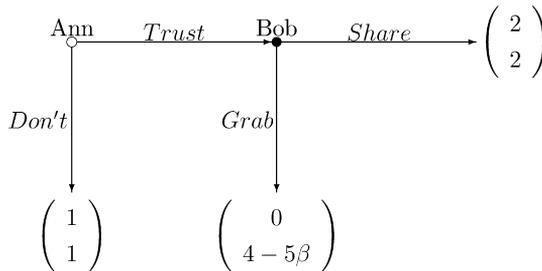


Fig. 2. Psychological trust game Γ_2 .

various ventures of a game tree requires that (intentions-based) kindness be re-evaluated using updated belief. The argument is an instance of **R1**.

Reciprocity theory does not provide the easiest route to illustrating the key issues involved though. Instead, we consider the motivation of guilt aversion, applied to the trust game Γ_1 in Fig. 1.¹⁰ Payoffs are in dollars and do not necessarily represent preferences. Therefore, we call them ‘material payoffs.’

We now modify Γ_1 to incorporate a guilt sentiment of Bob’s: To make our point, let us first specify what it means that ‘Bob lets Ann down.’ Ann is let down if the material payoff she gets is less than what she expected. Let α be the probability that Ann (initially) assigns to Bob’s strategy *Share* if *Trust*. Bob suffers from guilt to the extent that he believes he lets Ann down. The higher is α the more let down she will be if he chooses *Grab*. Bob does not know what α is, as this belief is in the mind of Ann. However, he has a belief about α . Let β be Bob’s expectation of α , conditional on Ann choosing *Trust*. We can model guilt aversion assuming that Bob’s utility at the terminal history (*Trust*, *Grab*) is decreasing in β . See Γ_2 in Fig. 2. What appears at the terminal histories should be thought of as utilities, not material payoffs although the notions coincide for all but one terminal histories.¹¹

Γ_2 is not a psychological game in GPS’ class, because β (being an updated belief) is not captured by any element of \bar{M}_i . This in itself illustrates **R1**. However, in order to appreciate the

¹⁰ Battigalli and Dufwenberg [8] use the framework of the present paper to develop a general model of (two forms of) guilt aversion for extensive game forms. For the specific context of trust games, related sentiments have previously been considered by Huang and Wu [42], Dufwenberg [24,25], Dufwenberg and Gneezy [27], Guerra and Zizzo [39], Charness and Dufwenberg [23], and Bacharach, Guerra and Zizzo [5].

¹¹ There is no special significance to the “5” in Fig. 2; we could have chosen many other numbers to make the upcoming point. Similar remarks apply to all examples below.

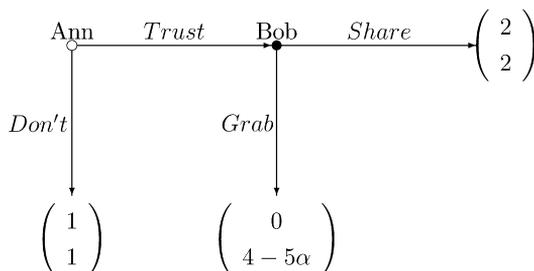


Fig. 3. Psychological Trust Game Γ_3 .

significance of this issue, it is useful to note that one can draw compelling (we think) conclusions about behavior that hinge crucially on the fact that β is an updated belief.

Following Dufwenberg [24,25], consider the following (for the time being intuitive) ‘psychological forward induction’ argument: Suppose Ann chooses *Trust*. If she is rational, she must believe the probability that Bob would choose *Share* (after *Trust*) is at least $\frac{1}{2}$, i.e., $\alpha \geq \frac{1}{2}$. Since we can figure this out, presumably Bob can too. Even if he is uncertain regarding the value of α , he infers it is at least $\frac{1}{2}$. Hence $\beta \geq \frac{1}{2}$. Since $4 - 5\beta < 2$ if $\beta \geq \frac{1}{2}$, he prefers *Share*. Since we can figure this out, presumably Ann can too. Hence she chooses *Trust*, fully expecting Bob to *Share* (so $\alpha = 1$). Bob figures this out (so that $\beta = 1$), which further reinforces his preference to *Share*. The path (*Trust*, *Share*) is predicted!

The argument depends on belief β being *conditional* on Ann choosing *Trust*. It cannot be recast using GPS’ theory, since $\bar{\mathbf{M}}_i$ contains only initial beliefs, but it can be captured in our framework, since \mathbf{M}_i contains conditional beliefs.

R2: others’ beliefs

There are two independent justifications for letting a player’s utility depend on others’ beliefs. First, this may be an adequate description of how certain social rewards operate. Refer back to Example 3 from the Introduction, where Eva’s preferences over effort depends on Fred’s inferences. It is taken from Dufwenberg and Lundholm [30]. A related example is Bernheim’s [16] model of social conformity. Another example is Caplin and Leahy’s [22] story of an information providing doctor concerned about the belief-dependent anxiety of a patient.¹² These authors develop models where a player’s utility depends on others’ beliefs (although only Caplin and Leahy explicitly refer to psychological games).¹³

The second justification concerns convenience in modeling. Refer back to the discussion of Γ_2 , including the definition of α and β . We modeled Bob’s guilt feelings by letting his psychological payoff depend on β , an updated second order belief. It turns out that there is an equivalent modeling choice. One can assume that Bob’s utility at (ℓ, L) depends directly on α , rather than on β , although Bob is uncertain about the true value of α . He uses probability assessments to weigh the different possibilities. We get Γ_3 in Fig. 3.

¹² For related work see Caplin [19], Caplin and Eliaz [20], and Kőszegi [48] from whom Example 4 of the Introduction is taken.

¹³ The models can be interpreted as psychological games with asymmetric information where the utility of a player depends on the terminal beliefs of another player (cf. Section 6.2).

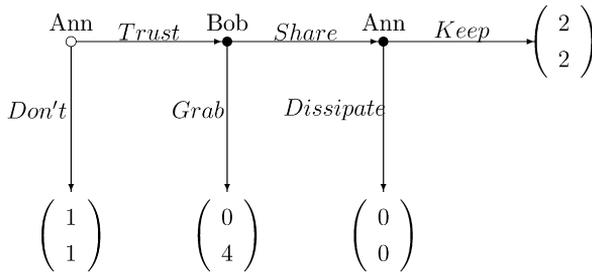


Fig. 4. Modified Trust Game Γ_4 with material payoffs.

After *Trust*, when Bob has to make a choice he compares 2, the payoff of action *Share*, with the conditional expected payoff of action *Grab*, that is $E_2[4 - 5\alpha | Trust] = 4 - 5\beta$; thus, we obtain the same results as with Γ_2 .¹⁴

This illustrates an important point: some belief-dependent motivations can be modeled replacing a conditional own belief of a certain ‘order’ (meaning: how many layers of beliefs about beliefs/choices are involved) with another object involving one degree lower order. This may allow one to work with utilities $u_i : Z \times \prod_{j \neq i} (\mathbf{M}_j \times S_j) \rightarrow \mathbb{R}$, where \mathbf{M}_i is *not* a factor of the domain. This has two advantages. First, it may seem easier to represent preferences with lower order beliefs (like α in Γ_3 rather than β in Γ_2). Second, and most importantly, one is lead to clearly distinguish between the carriers of utility (*i.e.*, elements of $Z \times \prod_{j \neq i} (\mathbf{M}_j \times S_j)$) and how a player deals with uncertainty by making updated probabilistic predictions (described by elements of \mathbf{M}_i). By contrast, when the domain of i ’s utility is $Z \times \mathbf{M}_i \times \prod_{j \neq i} (\mathbf{M}_j \times S_j)$ elements of \mathbf{M}_i serve both purposes.

R3: dependence on plans

Many forms of belief-dependent motivation require preferences to depend on overall strategies, beyond how strategies cause terminal histories when they are carried out.¹⁵ Consider Γ_4 in Fig. 4, a variation of Γ_1 where Ann may ‘dissipate’ some payoff. The payoffs of Γ_4 are material, not necessarily reflecting utilities.

Recall (from the discussion of **R1**) the terminology that ‘Ann is let down’ if the material payoff she gets is less than what she expects. In Γ_4 what she expects to get does not only depend on her beliefs about Bob, but also on how she plans to play. Suppose Ann plans to trust Bob and then keep the surplus. Then her subjectively expected material payoff is 2α , where α is the probability Ann assigns to Bob’s strategy *Share if Trust*. But if she plans to trust Bob and then dissipate the surplus, her expected material payoff is zero independently of α .

¹⁴ We do not suggest that Γ_3 is interesting only in providing a convenient alternative way to analyzing Γ_2 ; the emotion modeled in Γ_3 may make sense in its own right, as a primitive assumption about preferences (akin to Examples 3 and 4 of the Introduction).

¹⁵ This point has been anticipated by Caplin and Leahy [21]. Mariotti [60] also makes a similar point, but not in the context of psychological games.

Again assume that Bob suffers from guilt to the extent that he believes he lets Ann down. A natural way to model this is to let his utility at (*Trust*, *Grab*) be $4 - 5\alpha$ (as in Γ_3) if Ann plans to *Keep*, but 4 if she plans to *Dissipate*.¹⁶

This example illustrates that psychological motivations may exhibit a concern for other players' intentions. Intentions depend on beliefs as well as plans, and the latter dependence goes beyond the endnodes implied by implementing such plans. Therefore, the domain of u_i includes (conditional) beliefs and strategies of other players, on top of terminal histories and own beliefs.

R4: *non-equilibrium analysis*

R1–R3 concern features of players' motivation one may wish to incorporate in a formal framework. The next step is to predict play. We propose a generalization of Kreps and Wilson's [51] sequential equilibrium. We postpone illustrations until we formally introduce the concept in Section 4.

While much of economic theory presumes that players coordinate on an equilibrium, it is not always clear such an assumption is justified. For one thing, people may be quite rational, and confident in others' rationality, even if they fail to coordinate. In conventional game theory, related matters have inspired work on the implications of common belief of rationality; see *e.g.* Bernheim's [15] and Pearce's [65] work on rationalizability. This brings us to **R4**. There is little reason to assume that equilibrium coordination is easier in psychological games than in standard games. In fact, since psychological games often seem more complicated, and since problems of equilibrium multiplicity may be enhanced, assuming equilibrium may be assuming too much *especially* in psychological games.¹⁷

Giving up the equilibrium assumption does not necessarily mean giving up on predictive power. Refer back to the psychological forward induction argument, presented for Γ_2 . Ann and Bob perform deductive reasoning regarding one another's behavior and beliefs, and a clear-cut prediction results despite that no presumption of equilibrium is made. However, the story told was informal, and specific to Γ_2 (or, equivalently, Γ_3). It is natural to wonder about generally applicable formalizations. In Section 5, we develop a framework for analyzing interactive epistemology in psychological games, without postulating equilibrium play. This is a relatively small step because our very definition of psychological game already provides the necessary ingredients. Building on an epistemic theme due to Battigalli and Siniscalchi [11], we extend Pearce's [65] classical notion of (extensive form) rationalizability to psychological games. The concept captures psychological forward induction in simple games like Γ_2 and Γ_3 , and in more complicated games for which long chains of beliefs about beliefs are needed to get sharp predictions.

3. Psychological games

In this section we introduce notation on extensive-forms (3.1), model a universal belief space that accounts for updated beliefs (3.2), and put forth and illustrate our general definition of a psychological game (3.3).

¹⁶ Note also that Ann's anticipation of feeling let down might affect her initial decision. This can be modeled by letting Ann's utility at (*Trust*, *Grab*) be affected by her initial beliefs and her own plan. We pursue this point and its ramifications in Section 6.

¹⁷ For more on this, see Section 6.4.

3.1. Extensive forms with observable actions

We first restrict attention to finite multi-stage games with observable actions, no chance moves, and complete information. These restrictions can be removed, at the cost of additional notational complexity (see Section 6). We assume players move simultaneously at every stage. This is without loss of generality, because the set of feasible actions of a player may depend on actions chosen in previous stages and may be singleton. Simultaneous moves games, perfect information games, and repeated games are special cases (cf. Fudenberg and Tirole [34], §3.3, Osborne and Rubinstein [63], ch. 6). We use the following notation/terminology:

An *extensive form* with observable actions is a tuple $\langle N, H \rangle$ where $N = \{1, \dots, n\}$ is the *player set*, and H is the finite set of feasible *histories*. A history of length ℓ is a sequence $h = (a^1, \dots, a^\ell)$ where each $a^t = (a_1^t, \dots, a_n^t)$ represents the profile of actions chosen at stage t ($1 \leq t \leq \ell$). We assume history h becomes public information as soon as it occurs. The *empty history* (of length 0), denoted h^0 , is an element of H . The set of feasible actions for player i at history h is denoted $A_i(h)$ and may be singleton, meaning that i is not active at h . $A_i(h)$ is empty if and only if h is a *terminal* history. Z denotes the set of terminal histories.

For any given extensive form, we let S_i denote the set of (pure) strategies of player i . A typical strategy is denoted by $s_i = (s_{i,h})_{h \in H \setminus Z}$, where $s_{i,h}$ is the action that would be selected by s_i if history h occurred. Define $S = \prod_{i \in N} S_i$ and $S_{-i} = \prod_{j \neq i} S_j$. The set of i 's strategies that allow history h is denoted $S_i(h)$. Similar notation is used for strategy profiles: $S(h) = \prod_{i \in N} S_i(h)$; $S_{-i}(h) = \prod_{j \in N} S_j(h)$. We let $\zeta(s) \in Z$ denote the terminal history induced by $s = (s_i)_{i \in N}$.

3.2. Conditional beliefs and infinite hierarchies of beliefs

Here we summarize the theory of hierarchies of conditional beliefs due to Battigalli and Siniscalchi [10], which should be consulted for proofs, details, and further references. Consider a decision maker DM who is uncertain about which element in a set X is true. Assume X is a compact Polish space.¹⁸ DM assigns probabilities to events E, F, \dots in the Borel sigma-algebra \mathcal{B} of X according to some (countably additive) probability measure. Let $\Delta(X)$ denote the set of all such probability measures.¹⁹ As events unfold DM updates her beliefs. The actual and/or potential beliefs of DM are described by a conditional probability system (see R enyi [70]). Let $\mathcal{C} \subseteq \mathcal{B}$ denote the collection of potentially observable events (or conditioning events). DM holds probabilistic beliefs conditional on each event $F \in \mathcal{C}$.

Definition 1. A *conditional probability system* (cps) on $(X, \mathcal{B}, \mathcal{C})$ is a function $\mu(\cdot|\cdot) : \mathcal{B} \times \mathcal{C} \rightarrow [0, 1]$ such that for all $E \in \mathcal{B}, F, F' \in \mathcal{C}$

- (1) $\mu(\cdot|F) \in \Delta(X)$,
- (2) $\mu(F|F) = 1$,
- (3) $E \subseteq F' \subseteq F$ implies $\mu(E|F) = \mu(E|F')\mu(F'|F)$.

¹⁸ A topological space X is *Polish* if it admits a compatible metric d such that (X, d) is a complete and separable metric space (see, e.g., [46], p. 13).

¹⁹ Note that (X, \mathcal{B}) is a standard Borel space (see, e.g., [46], Definition 12.5).

We regard the set of cps' on $(X, \mathcal{B}, \mathcal{C})$ as a subset of the topological space $[\Delta(X)]^{\mathcal{C}}$, where $\Delta(X)$ is endowed with the topology of weak convergence of measures and $[\Delta(X)]^{\mathcal{C}}$ is endowed with the product topology.

From now on DM is a player i ; $(X, \mathcal{B}, \mathcal{C})$ is either $X = S_{-i}$ (a finite set) or $X = S_{-i} \times Y$ where Y is a compact Polish space typically representing a set of opponents' beliefs. The Borel sigma-algebra \mathcal{B} is implicitly understood,²⁰ and conditioning events corresponds to histories, i.e., $\mathcal{C} = \{F \subseteq S_{-i} \times Y: F = S_{-i}(h) \times Y, h \in H\}$ (or $\mathcal{C} = \{F \subseteq S_{-i}: F = S_{-i}(h), h \in H\}$ if $X = S_{-i}$). The set of cps' is denoted $\Delta^H(S_{-i} \times Y)$ a subset of $[\Delta(S_{-i} \times Y)]^H$. If conditioning event F corresponds to history h , then we abbreviate as $\mu(\cdot|F) = \mu(\cdot|h)$. Note that our specification of the conditioning events relies on interpreting s_j as an objective description of how j would behave at each decision node. However, we will also interpret s_j as a *plan* in the mind of player j . The implicit assumption underlying our analysis (as well as most papers on interactive epistemology in games) is that each player has correct beliefs, given by his plan of action, about how he would choose at different histories (and there is common certainty of this). This is the reason why, like GPS (and Battigalli and Siniscalchi [11]), we do not explicitly model i 's beliefs about his own behavior. We will come back to this in Section 6.

The following result shows that $\Delta^H(S_{-i} \times Y)$ is a compact Polish space, just like Y .²¹ It is key in our construction of hierarchical beliefs, implying that the domains of higher- and lower-order uncertainty have the same structural properties.

Lemma 2. $\Delta^H(S_{-i})$ is a compact Polish space. Furthermore, if Y is a compact Polish space, also $\Delta^H(S_{-i} \times Y)$ is a compact Polish space.

Hierarchies of cps' are defined recursively as follows:

- $X_{-i}^0 = S_{-i}$ ($i \in N$),
- $X_{-i}^k = X_{-i}^{k-1} \times \prod_{j \neq i} \Delta^H(X_{-j}^{k-1})$ ($i \in N; k = 1, 2, \dots$).

By repeated applications of Lemma 2, each X_{-i}^k is a cross-product of compact Polish spaces, hence compact Polish itself.²² A cps $\mu_i^k \in \Delta^H(X_{-i}^{k-1})$ is called k -order cps. For $k > 1$, μ_i^k is a joint cps on the opponents' strategies and $(k - 1)$ -order cps'. A *hierarchy of cps'* is a countably infinite sequence of cps' $\mu_i = (\mu_i^1, \mu_i^2, \dots) \in \prod_{k>0} \Delta^H(X_{-i}^{k-1})$. μ_i is *coherent* if the cps' of distinct orders assign the same conditional probabilities to lower-order events:

$$\mu_i^k(\cdot|h) = \text{marg}_{X_{-i}^{k-1}} \mu_i^{k+1}(\cdot|h) \quad (k = 1, 2, \dots; h \in H).$$

It can be shown that a coherent hierarchy μ_i induces a cps v_i on the cross-product of S_{-i} with the sets of hierarchies of cps' of i 's opponents, a compact Polish space.

However, v_i may assign positive probability (conditional on some h) to opponents' incoherence. To rule this out, say that a coherent hierarchy μ_i satisfies belief in coherency of order 1 if the induced cps v_i is such that each $v_i(\cdot|h)$ ($h \in H$) assigns probability one to the opponents'

²⁰ \mathcal{B} obtains from the product of the discrete topology on S_{-i} and the topology of Y .

²¹ This depends on two facts: (1) the collection of conditioning events for player i (corresponding to H) is at most countable (indeed finite), and (2) each conditioning event $S_{-i}(h) \times Y$ (or $S_{-i}(h)$ if $X = S_{-i}$) is both closed and open, as we consider the discrete topology on S_{-i} .

²² The cross-product of countably many compact Polish spaces is also compact Polish.

coherency; μ_i satisfies belief in coherency of order k if it satisfies belief in coherency of order $k - 1$ and the induced cps v_i is such that each $v_i(\cdot|h)$ ($h \in H$) assigns probability one the opponents' coherency of order $k - 1$; μ_i is *collectively coherent* if it satisfies belief in coherency of order k for each positive integer k . The set of collectively coherent hierarchies of player i is a compact Polish space, denoted by \mathbf{M}_i . We let M_i^k denote the set of k -order beliefs consistent with collective coherency, that is, the projection of \mathbf{M}_i on $\Delta^H(X_{-i}^{k-1})$, and let $M_{-i}^k = \prod_{j \neq i} M_j^k$, $\mathbf{M}_{-i} = \prod_{j \neq i} \mathbf{M}_j$, $\mathbf{M} = \prod_{j \in N} \mathbf{M}_j$.

We have now defined all components of the domain of the utility functions. But is this enough for the analysis of strategic reasoning? In order to decide on the best course of action, player i may need to form (conditional) beliefs about the *infinite* hierarchies of (conditional) beliefs of other players, either because they affect his payoff or because his assessment of the behavior and finite-order beliefs of others is derived from assumptions, such as “common belief in rationality,” involving beliefs of infinitely many orders. Does this mean that we need additional layers of beliefs? No. The following result shows that the countably infinite hierarchies of cps' defined above are sufficient for the strategic analysis; \mathbf{M}_i is isomorphic to $\Delta^H(S_{-i} \times \mathbf{M}_{-i})$, so each $\mu_i \in \mathbf{M}_i$ corresponds to a cps on $S_{-i} \times \mathbf{M}_{-i}$:

Lemma 3. *For each $i \in N$ there is a 1-to-1 and onto continuous function*

$$f_i = (f_{i,h})_{h \in H} : \mathbf{M}_i \rightarrow \Delta^H(S_{-i} \times \mathbf{M}_{-i})$$

whose inverse is also continuous. Furthermore, each coordinate function $f_{i,h}$ is such that for all $\mu_i = (\mu_i^1, \mu_i^2, \dots) \in \mathbf{M}_i$, $k \geq 1$,

$$\mu_i^k(\cdot|h) = \text{marg}_{S_{-i} \times M_{-i}^1 \times \dots \times M_{-i}^{k-1}} f_{i,h}(\mu_i).$$

3.3. Psychological games

We are now ready to state our definition of a psychological game:

Definition 4. A *psychological game* based on extensive form $\langle N, H \rangle$ is a structure $\Gamma = \langle N, H, (u_i)_{i \in N} \rangle$ where $u_i : Z \times \mathbf{M} \times S_{-i} \rightarrow \mathbb{R}$ is i 's (measurable and bounded) psychological payoff function.

The numerical examples examined in Section 2 fit this definition: in Γ_2 , u_2 depends on z and $\mu_2^2(\cdot|Trust)$; ²³ in Γ_3 , u_2 depends on z and 1's initial first-order belief, $\mu_1^1(\cdot|h^0)$; finally, the psychological payoff function u_2 proposed to analyze Γ_4 (a game with material payoffs) depends on z , $\mu_1^1(\cdot|h^0)$, and s_1 .

In all these examples, a psychological game is obtained from a *material payoff game* $\langle N, H, (\pi_i : Z \rightarrow \mathbb{R})_{i \in N} \rangle$ according to some formula. We now illustrate a few such derivations, focusing on two-player games.

²³ $\mu_2^2(\cdot|Trust)$ is the conditional second-order belief of player 2 used to compute the expectation β of the probability α initially assigned by 1 to the strategy 'Share if Trust.'

Player j is ‘let down’ if his actual material payoff, denoted $\hat{\pi}_j$, is lower than the payoff he initially expected to get, $E_j[\tilde{\pi}_j]$. This disappointment can be measured by the following expression:²⁴

$$\max\{0, (E_j[\tilde{\pi}_j] - \hat{\pi}_j)\}.$$

Suppose that i likes his material payoff but dislikes disappointing j . A simple way to model such ‘guilt’ motivation is to assume that i wants to maximize the expected value of the following expression:

$$\hat{\pi}_i - \theta_i \max\{0, (E_j[\tilde{\pi}_j] - \hat{\pi}_j)\}.$$

Taking explicitly into account the material payoff game, and that $E_j[\tilde{\pi}_j]$ is a function of j ’s plan and his initial beliefs about i ’s plan, we obtain the following utility function:

$$u_i(z, \mu, s_j) = \pi_i(z) - \theta_i \max\left\{0, \left[\sum_{s'_i} \mu_i^1(s'_i|h^0)\pi_j(\zeta(s_j, s'_i)) - \pi_j(z)\right]\right\}$$

where $\theta_i \geq 0$ is a psychological sensitivity parameter. For the special case of material payoff game Γ_1 , we obtain psychological game Γ_3 by letting $\theta_1 = 0$ and $\theta_2 = \frac{5}{2}$.²⁵

Another motivation is the willingness to give up some material payoff to avoid feeling regret *ex post*. The regret of i can be captured by the distance between the actual material payoff and the maximal expected payoff that could have been obtained ‘with the benefit of hindsight,’ *i.e.*, using the terminal beliefs. Formally, i ’s regret when he gets material payoff $\hat{\pi}_i$ and has terminal belief $\tau_i \in \Delta(S_j)$ equals

$$\max_{s'_i} \sum_{s'_j} \tau_i(s'_j)\pi_i(\zeta(s_i, s'_j)) - \hat{\pi}_i.$$

Taking into account that the actual payoff obtains at a terminal history z , which is observed by player i with first-order cps μ_i^1 , we obtain the utility function

$$u_i(z, \mu, s_j) = \pi_i(z) - \theta_i \left[\max_{s'_i} \sum_{s'_j} \mu_i^1(s'_j|z)\pi_i(\zeta(s_i, s'_j)) - \pi_i(z) \right]$$

where $\theta_i \geq 0$ is a psychological sensitivity parameter.²⁶ This shows that it may be natural to let utility depend on what players believe at the end of the game (for other examples of this kind see Subsection 6.2).

Note that while we allow for a general functional form with very complex arguments, $u_i(z, \mu, s_j)$, the utility functions used in specific applications can be relatively simple.

²⁴ Bell [13] and Loomes and Sugden [57] present decision-theoretic models of belief-dependent disappointment. What we have here is a game-theoretic extension. In this section we use it as input to formulate guilt (see next paragraph). Concern for own disappointment may naturally depend on one’s own plan, which is excluded under Definition 4 but considered again in Section 6.3.

²⁵ The modeling choices here reflect the discussion of **R2** in Section 2 and are in line with Battigalli and Dufwenberg’s [8] notion of simple guilt to whom we refer for more guilt examples (the mathematical formulation in that paper is slightly more complex, but leads to a utility with the same best response correspondence).

²⁶ Bell [12] and Loomes and Sugden [56] develop theories of regret, in which a decision maker’s experienced utility depends on the post-choice revelation of a state-of-nature. Our formulation preserves that spirit, but extends it to belief-dependent motivation. This is natural in a strategic setting, where players cannot perfectly observe *ex post* the state of the world, which includes what another player *would* have chosen.

So far we have illustrated three forms of belief-dependent motivation which all relied on (different forms of) first-order beliefs (own initial, others' initial, own and/or others' terminal). Of course, Definition 4 allows also for higher-order belief-dependence. One example is Battigalli and Dufwenberg's [8] notion of "guilt from blame," which involves dependence on the co-player's third-order beliefs.²⁷ The rest of this section leads up to another class of examples, which involve infinite-order beliefs.

As we noted in Section 2, letting beliefs and strategies of others in u_i may yield simpler functional forms and clarify the distinction between prediction and psychological motivation. Yet, as reflected by Γ_2 and Γ_3 it is always possible to obtain an equivalent functional form where a player's utility depends only on the terminal history and his own conditional beliefs. Given $u_i(z, \mu, s_{-i})$ let

$$\hat{u}_i(z, \mu_i) := E_{\mu_i}[u_i|z]$$

where $E_{\mu_i}[u_i|z]$ is the conditional expectation operator given cps $f_i(\mu_i) \in \Delta^H(S_{-i} \times M_{-i})$ (see Lemma 3). In words, $\hat{u}_i(z, \mu_i)$ is how much i values z after he has observed it. It can be shown that \hat{u}_i and u_i yield the same (sequential) best responses.²⁸

The functional form \hat{u}_i has infinite hierarchy μ_i as an argument, but this is just because we want an abstract and general expression. If u_i depends only on beliefs of order k , then \hat{u}_i depends only on beliefs of order $k + 1$. An example that necessarily involves infinite-order belief-dependence arises if belief-dependent motivation appears together with 'interactive altruism.' Suppose that the utility of i is given by two terms as follows

$$\begin{aligned} u_i &= \phi_i + \theta u_j \quad (0 < \theta < 1), \\ \phi_i &= \phi_i(z, \mu^1, s_j) \end{aligned}$$

(ϕ_i could be a guilt or regret component). Then we can compute an explicit form for the expected utility conditional on z by repeated substitution, where the term that depends on $(k + 2)$ -order beliefs has exponentially decreasing weight θ^k :

$$\begin{aligned} \hat{u}_i(z, \mu_i) &= E_{\mu_i}[\phi_i + \theta u_j|z] \\ &= E_{\mu_i^2}[\phi_i|z] + \theta E_{\mu_i}[u_j|z] \\ &= E_{\mu_i^2}[\phi_i|z] + \theta E_{\mu_i^3}[E_{\mu_j^2}[\phi_j|z]|z] + \theta^2 E_{\mu_i}[E_{\mu_j}[u_j|z]|z] \\ &= \dots \end{aligned}$$

($E_{\mu_i^k}[\cdot|z]$ is the conditional expectation operator given cps μ_i^k).

4. Equilibrium analysis

Kreps and Wilson's sequential equilibrium has become a benchmark for the analysis of standard games. We extend this concept to the class of psychological games defined in Section 3. (The restriction to multi-stage game forms with complete information simplifies but is not essential; cf. Section 6.) We next define and interpret mixed strategies and assessments (4.1), give the main definition and provide an existence theorem (4.2), and consider examples (4.3).

²⁷ To save space we refer to the original source for precise definitions and examples. The intuition is that i experiences guilt (judging by his terminal beliefs) to the extent that j 's terminal beliefs indicate that i intended to disappoint j .

²⁸ See the first part of Appendix A and Lemma 17.

4.1. Randomized strategies and consistent assessments

The equilibrium concept we develop refers to randomized choices. However, in our interpretation, we exclude actual randomization. Rather, we interpret a randomized choice of a given player i as the common first-order belief of i 's opponents about i (cf. Aumann and Brandenburger [4]). This is akin to the following characterization of a Nash equilibrium in a standard simultaneous-move game: $(\sigma_1, \dots, \sigma_n) \in \Delta(A_1) \times \dots \times \Delta(A_n)$ is an equilibrium if, for each i , each action in the support of σ_i is a best reply to σ_{-i} .

We focus on behavior strategies $\sigma_i = (\sigma_i(\cdot|h))_{h \in H \setminus Z} \in \prod_{h \in H \setminus Z} \Delta(A_i(h))$, interpreting σ_i as an array of common conditional first-order beliefs held by i 's opponents. This interpretation is part of the notion of ‘consistency’ of profiles of strategies and hierarchical beliefs defined below.

Kreps and Wilson [51] argue that an appropriate definition of equilibrium in extensive form games must refer to ‘assessments’: profiles of (behavior) strategies *and* conditional (first-order) beliefs. They proceed in two steps: first a ‘consistency’ condition for assessments is put forward, and then sequential equilibrium is defined as a consistent assessment satisfying sequential rationality. It turns out that consistency captures the assumptions that each player regards his opponents’ choices at different histories as stochastically independent, and any two players have the same (prior and conditional) beliefs about any third player. In our framework an *assessment* is a profile $(\sigma, \mu) = (\sigma_i, \mu_i)_{i \in N}$ where σ is a behavioral strategy profile and $\mu \in \mathbf{M}$. We extend the definition of consistency by adding a requirement concerning the higher-order beliefs that need to be specified in psychological games.

Let $\Pr_{\sigma_j}(\cdot|\hat{h}) \in \Delta(S_j(\hat{h}))$ denote the probability measure over j 's strategies conditional on \hat{h} derived from behavior strategy σ_j under the assumption of independence across histories:

$$\forall s_j \in S_j(\hat{h}), \quad \Pr_{\sigma_j}(s_j|\hat{h}) := \prod_{h \in H \setminus Z: h \not\prec \hat{h}} \sigma_j(s_j, h|h)$$

($h \not\prec \hat{h}$ means that h is not a predecessor, or prefix, of \hat{h}).²⁹

Definition 5. A profile of first-order cps’ $\mu^1 = (\mu_i^1)_{i \in N}$ is *derived* from a behavioral strategy profile $\sigma = (\sigma_i)_{i \in N}$ if for all $i \in N$, $s_{-i} \in S_{-i}$, $\hat{h} \in H$,

$$\mu_i^1(s_{-i}|\hat{h}) = \prod_{j \neq i} \Pr_{\sigma_j}(s_j|\hat{h}). \tag{1}$$

Clearly, if μ^1 is derived from σ then for any three players i, j, k , the beliefs of i and j about k coincide:

$$\forall \hat{h} \in H, \quad \text{marg}_{S_k} \mu_i^1(\cdot|\hat{h}) = \Pr_{\sigma_k}(\cdot|\hat{h}) = \text{marg}_{S_k} \mu_j^1(\cdot|\hat{h}).$$

We are now ready for the main definition of this subsection.

Definition 6. Assessment (σ, μ) is *consistent* if

- (a) μ^1 is derived from σ ,

²⁹ Cf. Kuhn [53]: $\Pr_{\sigma_j}(\cdot|h) \in \Delta(S_j(h))$ is only one out of the many probability measures that induce the same outcome probabilities starting from h , for all $s_{-j} \in S_{-j}(h)$. But note that realization-equivalent beliefs may yield different psychological utilities!

(b) higher-order beliefs in μ assign probability 1 to the lower-order beliefs:

$$\forall i \in N, \forall k > 1, \forall h \in H, \quad \mu_i^k(\cdot|h) = \mu_i^{k-1}(\cdot|h) \times \delta_{\mu_{-i}^{k-1}}$$

where \times denotes the product of measures and δ_x is the Dirac measure assigning probability 1 to singleton $\{x\}$.

A justification of the (strong) condition (b) comes from (i) the classical interpretation of equilibrium beliefs as the end-product of a transparent reasoning process by intelligent players, and (ii) the trembling hand assumption underlying the sequential equilibrium concept.³⁰ (i) implies that any two players must share the same initial first-order beliefs about any other player, and every player comes to a correct conclusion about the (hierarchical) beliefs of his opponents because he is able to replicate their reasoning. (ii) implies that unexpected moves are explained as mistakes, not as the result of unexpected beliefs, therefore players never change their beliefs about the conditional beliefs that the opponents would hold at each h . Of course, by observing the actual play-path each player infers the current actual beliefs of his opponents, but interesting forms of learning about others' beliefs are ruled out. We offer further comments on consistency and sequential equilibrium in Section 6.4.

We also note that (b) is analogous to a condition used by GPS to define psychological Nash equilibrium, requiring that players hold common, correct beliefs about each others' beliefs. Indeed, (b) is equivalent to the requirement that, for each player i and each history h , the conditional belief on $S_{-i} \times \mathbf{M}_{-i}$ induced by hierarchy μ_i assigns probability one to μ_{-i} .³¹

4.2. Sequential equilibrium assessments

We now move to the section's main definition: a consistent assessment is a sequential equilibrium if it satisfies sequential rationality. Formally, fix a hierarchy of cps' μ_i a (non-terminal) history h and a strategy s_i consistent with h . The expectation of u_i conditional on h , given s_i and μ_i is

$$E_{s_i, \mu_i}[u_i|h] := \int_{S_{-i} \times \mathbf{M}_{-i}} u_i(\zeta(s_i, s_{-i}), \mu_i, \mu_{-i}, s_{-i}) f_{i,h}(\mu_i)(ds_{-i}, d\mu_{-i}). \tag{2}$$

Definition 7. An assessment (σ, μ) is a *sequential equilibrium (SE)* if it is consistent and for all $i \in N, h \in H \setminus Z, s_i^* \in S_i(h)$,

$$\Pr_{\sigma_i}(s_i^*|h) > 0 \quad \Rightarrow \quad s_i^* \in \arg \max_{s_i \in S_i(h)} E_{s_i, \mu_i}[u_i|h]. \tag{3}$$

Note that, by consistency, σ_i represents the first-order beliefs of i 's opponents about i , and furthermore there is common certainty of the true belief profile μ at every history; therefore the sequential rationality condition (3) can equivalently be written as

$$\forall j \neq i, \quad \text{supp marg}_{S_i} \mu_j^1(\cdot|h) \subseteq \arg \max_{s_i \in S_i(h)} \sum_{s_{-i} \in S_{-i}(h)} \mu_i^1(s_{-i}|h) u_i(\zeta(s_i, s_{-i}), \mu, s_{-i}). \tag{4}$$

³⁰ Kreps and Wilson [51] show that (in standard games) sequential equilibrium is generically equivalent to trembling hand perfect equilibrium.

³¹ That is, (b) holds iff $\forall i \in N, \forall h \in H, f_{i,h}(\mu_i)(S_{-i} \times \{\mu_{-i}\}) = 1$.

This clarifies that SE is a notion of equilibrium in beliefs. Indeed we could have given an equivalent definition of SE with no reference to behavioral strategies.

We can also take the point of view of an ‘agent’ (i, h) of player i , in charge of the move at history h , who seeks to maximize i ’s conditional expected utility given the consistent assessment (σ, μ) . The expected utility of i conditional on h and $a_i \in A_i(h)$ given (σ, μ) can be expressed as

$$E_{\sigma, \mu}[u_i|h, a_i] := \sum_{s_{-i} \in S_{-i}(h)} \prod_{j \neq i} \Pr_{\sigma_j}(s_j|h) \sum_{s_i \in S_i(h, a_i)} \Pr_{\sigma_i}(s_i|h, a_i) u_i(\zeta(s), \mu, s_{-i}), \tag{5}$$

where $\Pr_{\sigma_i}(s_i|h, a_i) := \prod_{h' \in H \setminus Z: h' \not\prec h} \sigma_i(s_i, h'|h')$ ($h' \not\prec h$ means that h' is not h or a predecessor of h). This specification presumes that (i, h) assesses the probabilities of actions by other agents of player i in the same way as each player $j \neq i$; that is using the behavioral strategy σ_j .³² It can be shown that a version of the One-Shot-Deviation (or unimprovability) principle holds in our framework³³:

Proposition 8. *A consistent assessment (σ, μ) satisfies (3) and hence is an SE if and only if for all $i \in N, h \in H \setminus Z$,*

$$\text{supp}(\sigma_i(\cdot|h)) \subseteq \arg \max_{a_i \in A_i(h)} E_{\sigma, \mu}[u_i|h, a_i]. \tag{6}$$

Proof. Available on request. \square

We obtain the following existence theorem:

Theorem 9. *If the psychological payoff functions are continuous, there exists at least one sequential equilibrium assessment.*

The proof relies on a ‘trembling hand’ argument (cf. Selten [74]). We only provide a sketch here (but details are available on request). Consider ε -perturbed games where there is strictly positive minimal probability of choosing any action at any history, i.e. $\varepsilon = (\varepsilon_{i,h}(a_i, h))_{a_i \in A_i(h), i \in N, h \in H}$ is a strictly positive vector such that $\sum_{a_i \in A_i(h)} \varepsilon(a_i, h) < 1$ for each h . For each behavior strategy profile σ , let $\mu = \beta(\sigma)$ denote the unique profile of hierarchies of cps’ such that (σ, μ) is consistent.³⁴ Define an (agent-form, psychological) ε -equilibrium as an ε -constrained behavior strategy profile σ_ε such that for each h and each i , a pure action a_i that does not maximize the expectation of u_i (given h and $\beta(\sigma_\varepsilon)$) is assigned the minimal probability $\varepsilon(a_i, h) > 0$. It can be shown by standard compactness-continuity arguments that each ε -perturbed game has an ε -equilibrium (cf. the existence proof for psychological Nash equilibria in GPS). Fix a sequence $\varepsilon^k \rightarrow 0$ and a corresponding sequence σ^k of ε^k -equilibrium assessments. By compactness, σ^k has an accumulation point σ^* . By upper-hemicontinuity of the local

³² Suppose that u_i depends only on terminal histories and beliefs, not on s_{-i} . Then we obtain the more familiar formula

$$E_{\sigma, \mu}[u_i|h, a_i] = \sum_z \Pr_\sigma(z|h, a_i) u_i(z, \mu),$$

where $\Pr_\sigma(z|h, a_i)$ is the probability of terminal history z conditional on (h, a_i) determined by σ .

³³ See Section 6.3 for further discussion.

³⁴ $(\mu_i^j)_{i \in N}$ is derived from σ via formula (1) and the infinite hierarchies of cps’ are obtained by assuming correct higher order beliefs.

best response correspondences, for each (i, h) , $\sigma_i^*(\cdot|h)$ assigns positive probability only to actions that are best responses to $(\sigma^*, \beta(\sigma^*))$ at h . So, by Proposition 8 $(\sigma^*, \beta(\sigma^*))$ is a sequential equilibrium assessment.

We next show that the SE concept generalizes subgame perfect equilibrium for standard games with observable actions (recall: sequential and subgame perfect equilibrium coincide in games with observable actions). This is a corollary of a more general result for games where psychological utilities depend only on terminal nodes and beliefs: $u_i : Z \times \mathbf{M} \rightarrow \mathbb{R}$. For any such game $\Gamma = \langle N, H, (u_i)_{i \in N} \rangle$ and any profile of hierarchies of cps' $\mu = (\mu_i)_{i \in N}$, we can obtain a *standard* game $\Gamma^\mu = \langle N, H, (v_i^\mu)_{i \in N} \rangle$ with payoff functions $v_i^\mu(z) = u_i(z, \mu)$.

Remark 10. Suppose $\forall i \in N, u_i : Z \times \mathbf{M} \rightarrow \mathbb{R}$. Then an assessment (σ, μ) is a SE if and only if it is consistent and σ is a subgame perfect (hence sequential) equilibrium of the standard game Γ^μ .

4.3. Examples

We illustrate the SE concept with three examples, first a simultaneous move game illustrating how we can reproduce the essence of a leading example of GPS, then two versions of the Trust Game connecting back to some of the key notions previously highlighted in Section 2.

4.3.1. Equilibrium beliefs in the Bravery Game

The Bravery Game is a numerical example used in GPS (p. 66) to show that a psychological game may have multiple, isolated mixed strategy equilibria even if there is only one active player, which is impossible in standard games. We consider a modified version to illustrate our definition of equilibrium in beliefs. Let $A_1 = \{\text{Wait}\}$, $A_2 = \{\text{bold}, \text{timid}\}$. Player 1 (Ann) is inactive so we can ignore her payoffs, but her beliefs matter. Player 2 (Bob) is concerned about what Ann thinks about him. Acting boldly is dangerous, but worthwhile if Ann expects Bob to act boldly. GPS model the situation with a payoff function of the form $\bar{u}_2 : A \times \bar{\mathbf{M}}_2 \rightarrow \mathbb{R}$. Specifically, let $\alpha := \mu_1^1(\text{bold}|h^0)$ denote Ann's first-order belief about Bob (a random variable from Bob's point of view), and let $\beta := \int \alpha \mu_2^2(d\mu_1^1)$ denote (a feature of) the Bob's second-order beliefs. GPS' payoff function is

$$\bar{u}_2(a_2, \bar{\mu}_2) = \begin{cases} 2 - \beta, & \text{if } a_2 = \text{bold}, \\ 3(1 - \beta), & \text{if } a_2 = \text{timid}. \end{cases}$$

We modify \bar{u}_2 , considering instead $u_2 : A \times \mathbf{M} \rightarrow \mathbb{R}$ defined by

$$u_2(a_2, \mu) = \begin{cases} 2 - \alpha, & \text{if } a_2 = \text{bold}, \\ 3(1 - \alpha), & \text{if } a_2 = \text{timid}. \end{cases}$$

Clearly, the expectation of u_2 given a_2 and Bob's second-order belief β is \bar{u}_2 . There are three equilibria, with $\beta = \alpha = 1$, $\beta = \alpha = 0$ and $\beta = \alpha = \frac{1}{2}$.³⁵

4.3.2. Trust game with guilt aversion

Consider Γ_3 in Fig. 3 (or equivalently Γ_2). Recall that α (a function of μ_1^1) is the probability Ann assigns to strategy 'Share if Trust' at the beginning of the game, and $\beta = \int \alpha \mu_2^2(d\mu_1^1|Trust)$

³⁵ These are essentially the same equilibria as those obtained by GPS. But GPS allow for actual randomization; thus the first-order beliefs of Bob are degenerate on the equilibrium (mixed) strategy of Ann, and higher-order beliefs of each player are degenerate on the equilibrium lower-order beliefs of the other player.

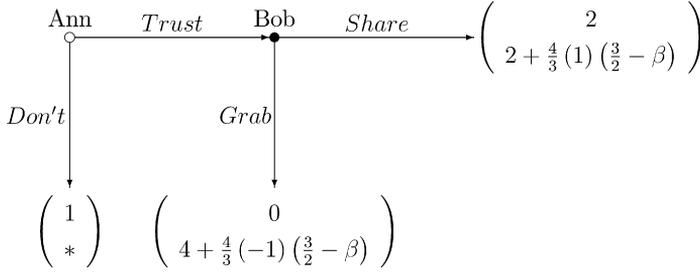


Fig. 5. Trust Game Γ_5 with reciprocity payoffs.

is the relevant feature of the conditional second-order beliefs of Bob. We let $\tau = \mu_2^1(Trust|h^0)$ denote Bob’s initial first-order belief. An assessment is summarized by (τ, α, β) , where (τ, α) corresponds to a behavior strategy profile. The indifference condition for Bob is $\beta = \frac{2}{5}$, the indifference condition for Ann is $\alpha = \frac{1}{2}$; consistency yields $\alpha = \beta$. The game has three SEs: $\tau = \alpha = \beta = 1$ (trust), $\tau = \alpha = \beta = 0$ (no trust), and $\tau = 0, \alpha = \beta = \frac{2}{5}$ (insufficient trust). Note that only the first equilibrium is consistent with forward induction reasoning (as described in Section 2, and further elaborated on in Subsection 5.1 below).

4.3.3. Trust game with reciprocity

Our framework is adequate for modeling reciprocity in extensive games. To support this claim, we show how the essence of Dufwenberg and Kirchsteiger’s theory can be captured in Γ_1 : Let α, β and τ be defined as in the previous example. The key tenets of the theory concern player i ’s kindness to player j (K_{ij}), and i ’s belief in j ’s kindness to i (\hat{K}_{iji}). At each history, player i maximizes utility defined by the sum of material payoffs (as in Γ_1) and reciprocity payoffs equal to $\theta_i \times K_{ij} \times \hat{K}_{iji}$, where θ_i is a constant measuring i ’s sensitivity to reciprocity. Assume that Ann’s and Bob’s sensitivities are $\theta_1 = 0$ and $\theta_2 = \frac{4}{3}$. One can show that K_{ij} and \hat{K}_{iji} can be reproduced in our framework and notation; in particular we need the following:

- Bob’s kindness following *Trust* = -1 for *Grab* and = 1 for *Share*,
- Bob’s belief in Ann’s kindness following *Trust* = $\frac{3}{2} - \beta$.

Γ_5 in Fig. 5 displays the relevant utilities as conceived by the players when they move (Bob is not active at h^0 , so we put no utility for him following *Don’t*):³⁶

Applying Definition 7, Γ_5 has a unique SE with $\tau = 1, \alpha = \beta = \frac{3}{4}$. No ‘pure’ SE exists,³⁷ just like in Dufwenberg and Kirchsteiger’s theory (cf. 6.3 below).

5. Interactive epistemology

We argued in Section 2 that alternatives to equilibrium analysis are worth exploring. The definition of \mathbf{M}_i provides us with all the ingredients to analyze strategic reasoning by means of

³⁶ As with Γ_2 vs. Γ_3 , we can replace Bob’s conditional second-order belief β with Ann’s initial first-order belief α in Bob’s payoffs, and get analogous conclusions.

³⁷ In any SE we have $\alpha = \beta$. With $\theta_2 = \frac{4}{3}$, the indifference condition for Bob yields $\beta = \frac{3}{4}$. If $\alpha = \beta < \frac{3}{4}$ then \hat{K}_{212} shoots up, so Bob prefers *Share* to *Grab*, which in SE would imply $\alpha = \beta = 1, \dots$ a contradiction. If $\alpha = \beta > \frac{3}{4}$ then \hat{K}_{212} goes down, so Bob prefers *Grab* to *Share*, implying $\alpha = \beta = 0, \dots$ another contradiction.

interactive epistemology, *i.e.*, assumptions about players’ rationality and what they believe about each other at any node. We show how to express such assumptions in the language of events and belief operators (5.1), and then analyze a notion of extensive form rationalizability (5.2).

5.1. States of the world, events, and belief operators

A state of the world specifies, for each player i and history h , what i would do and believe if h were reached. Note the subjunctive conditional: game-theoretic analysis does not only concern the actual path of actions and beliefs, but also considers how players *would* react (in terms of choices and beliefs) to histories that do not actually occur at the true state. The state of a player is therefore given by his strategy and his hierarchy of cps’, (s_i, μ_i) (as explained in Subsection 3.2 we interpret s_i both as an objective description of i ’s contingent behavior and as i ’s plan of action). The set of states for player i is denoted by $\Omega_i = S_i \times \mathbf{M}_i$, and the set of *states of the world* is $\Omega = \prod_{i=1}^n \Omega_i$. We let $\Omega_{-i} = \prod_{j \neq i} \Omega_j$ denote the set of possible states of i ’s opponents. With a slight abuse of notation we often write $\Omega = \Omega_i \times \Omega_{-i}$ with typical element $\omega = (\omega_i, \omega_{-i})$.

An *event* is a (Borel) subset $E \subseteq \Omega$; its complement is denoted $\neg E = \Omega \setminus E$. An event about i is any set of states $E = E_i \times \Omega_{-i}$, where E_i is a Borel subset of Ω_i . We let \mathcal{E}_i denote the family of events about i . Events about the opponents of i are similarly defined; the collection of such events is denoted \mathcal{E}_{-i} .

We often use brackets to denote specific events. In particular, for any function (random variable) $\mathbf{x} : \Omega \rightarrow X$ and value $x^* \in X$, we use the notation $[\mathbf{x} = x^*] = \{\omega : \mathbf{x}(\omega) = x^*\}$. When \mathbf{x} is understood, we simply write $[x^*]$. For example, $[s_i^*] = \{(s_i, \mu_i, \omega_{-i}) : s_i = s_i^*\} \in \mathcal{E}_i$ is the event “ i plays s_i^* ”; here it is understood that \mathbf{x} is the projection function $\mathbf{x}(s_i, \mu_i, \omega_{-i}) = s_i$. Similarly, $[h] = \prod_{i \in N} S_i(h) \times \mathbf{M}_i$ is the event that history h occurs.

Recall that we follow both GPS and Battigalli and Siniscalchi [11] in disregarding players’ beliefs about themselves. At state $\omega = (s_i, \mu_i, \omega_{-i})$, player i would believe event $E = \Omega_i \times E_{-i} \in \mathcal{E}_{-i}$ conditional on history h with probability $f_{i,h}(\mu_i)(E_{-i})$ (cf. Section 3.2). Thus $\{(s_i, \mu_i, \omega_{-i}) : f_{i,h}(\mu_i)(E_{-i}) = 1\}$ is the event “ i would believe E conditional on h .” E may concern the beliefs of i ’s opponents.

We use belief operators to represent events about interactive beliefs: a *belief operator* for player i is a mapping with domain \mathcal{E}_{-i} and range \mathcal{E}_i . For any given history $h \in H$, the *h-conditional belief operator* for i is defined as follows:

$$\forall E = \Omega_i \times E_{-i} \in \mathcal{E}_{-i}, \quad B_{i,h}(E) = \{(s_i, \mu_i, \omega_{-i}) : f_{i,h}(\mu_i)(E_{-i}) = 1\}.$$

h may be counterfactual at ω , because strategies played at ω may not induce h ; in this case “ i would believe E conditional on h ” is a counterfactual statement about i ’s beliefs at ω . Clearly, $B_{i,h}(E) \in \mathcal{E}_i$.³⁸ $B_{i,h}(\cdot)$ satisfies monotonicity [$E \subseteq F$ implies $B_{i,h}(E) \subseteq B_{i,h}(F)$] and conjunction [$B_{i,h}(E \cap F) = B_{i,h}(E) \cap B_{i,h}(F)$]. Furthermore $B_{i,h}(E) = B_{i,h}(E \cap [h])$ because i always believes what he observes.

The basic event we are interested in is players’ rationality. We focus on a condition that does not distinguish between realization-equivalent strategies. Say that i is *rational* at state $(s_i^*, \mu_i, \omega_{-i})$ iff s_i^* maximizes i ’s conditional expected utility $E_{s_i, \mu_i}[u_i | h]$ (defined in (2)) con-

³⁸ For any Borel set $\Omega_i \times E_{-i}$, $B_{i,h}(\Omega_i \times E_{-i})$ is also a Borel set because the h -coordinate belief function $f_{i,h}$ is continuous (see Lemma 3).

ditional on each history h allowed by s_i^* . More formally, let $H_i(s_i^*) = \{h \in H \setminus Z: s_i^* \in S_i(h)\}$ denote the set of non-terminal histories allowed by s_i^* ; we require $s_i \in r(\mu_i)$ where

$$r_i(\mu_i) := \left\{ s_i^*: \forall h \in H(s_i^*), s_i^* \in \arg \max_{s_i \in S_i(h)} E_{s_i, \mu_i}[u_i|h] \right\}. \tag{7}$$

The event “player i is rational” is $R_i = \{(s_i, \mu_i, \omega_{-i}): s_i \in r_i(\mu_i)\}$. It can be shown that $r_i(\mu_i)$ can be obtained via a backward induction algorithm and that R_i is a well-defined nonempty event (cf. proof of Lemma 14 in Appendix A).

To illustrate how these concepts can be used, we re-examine two psychological versions of the Trust Game. As regards notation, we have to distinguish the event “Bob shares,” which in the extensive form implies that “Ann trusts Bob,” from the event “Bob would share if Ann trusted Bob” which is a subjunctive conditional, logically independent on whether Ann trusts Bob or not. Similar considerations hold for the action *Grab*. We use bold letters to denote subjunctive conditionals (which in this case correspond to strategies of Bob), as in [Share] and [Grab].

Consider the Trust Game with guilt aversion Γ_3 in Fig. 3. The game can be solved by forward induction reasoning: it is rational for Ann to trust Bob only if she assigns at least 50% probability to strategy **Share**, i.e. only if $\alpha \geq \frac{1}{2}$, where $\alpha: \mathbf{M}_1 \rightarrow \mathbb{R}$ is the random variable defined by $\alpha(\mu_1) = \mu_1^1(\mathbf{Share}|h^0)$.³⁹ If Bob believes in Ann’s rationality when he has to move (even if he is ‘surprised’), he infers from Ann’s action *Trust* that $\alpha \geq \frac{1}{2}$. Therefore $\beta \geq \frac{1}{2}$, where $\beta: \mathbf{M}_2 \rightarrow \mathbb{R}$ is the random variable defined by $\beta(\mu_2) = \int \alpha(\mu_1) f_{2, Trust}(\mu_2)(d\mu_1)$. His rational response is to share. If Ann anticipates Bob’s reasoning she trusts him.

The formal counterpart of this argument is as follows (the events listed are nonempty; we rely on the monotonicity of the belief operators):

$$\begin{aligned} R_1 &= \left\{ (s_1, \mu_1, \omega_2): \alpha(\mu_1) > \frac{1}{2} \Rightarrow s_1 = Trust, \alpha(\mu_1) < \frac{1}{2} \Rightarrow s_1 = Don't \right\}, \\ R_2 &= \left\{ (\omega_1, s_2, \mu_2): \beta(\mu_2) > \frac{2}{5} \Rightarrow s_2 = \mathbf{Share}, \beta(\mu_2) < \frac{2}{5} \Rightarrow s_2 = \mathbf{Grab} \right\}, \\ R_1 \cap [Trust] &\subseteq \left[\alpha \geq \frac{1}{2} \right], \\ B_{2, Trust}(R_1) &= B_{2, Trust}(R_1 \cap [Trust]) \subseteq B_{2, Trust}\left(\left[\alpha \geq \frac{1}{2} \right]\right) \subseteq \left[\beta \geq \frac{1}{2} \right], \\ R_2 \cap B_{2, Trust}(R_1) &\subseteq R_2 \cap \left[\beta \geq \frac{1}{2} \right] \subseteq [\mathbf{Share}], \\ R_1 \cap B_{1, h^0}(R_2 \cap B_{2, Trust}(R_1)) &\subseteq R_1 \cap [\alpha = 1] \subseteq [Trust]. \end{aligned}$$

Now consider the Trust Game with Reciprocity Γ_5 in Fig. 5. Without an equilibrium supposition, one is at loss for predictive power: if $\theta_2 = \frac{4}{3}$, Bob’s best response depends on whether β is below or above $(\frac{3}{2} - \frac{1}{\theta_2}) = \frac{3}{4}$. This cannot be resolved by forward induction reasoning, which yields (as explained above) $\beta \geq \frac{1}{2}$.

However, if one uses other values of θ_2 one can draw clear conclusions merely using backward induction: if $\theta_2 < \frac{2}{3}$, Bob’s best response (given *Trust*) is *Grab* independently of β ,

³⁹ In some formulas, we have to make explicit the dependence of random variable α on the state of the world. The same holds for random variable β .

thus $R_2 \subseteq [\mathbf{Grab}]$ and $R_1 \cap B_{1,h^0}(R_2) \subseteq [Don't]$; on the other hand, if $\theta_2 > 2$, $R_2 \subseteq [\mathbf{Share}]$ and $R_1 \cap B_{1,h^0}(R_2) \subseteq [Trust]$. Furthermore, a subtle issue arises when $\frac{2}{3} < \theta_2 < 1$: backward induction cannot pin down Bob's best response, which is *Grab* only if $\beta \geq (\frac{3}{2} - \frac{1}{\theta_2})$, while forward induction yields $\beta \geq \frac{1}{2}$. This puts an *upper bound* on how kind Bob believes Ann is.⁴⁰ With $\frac{2}{3} < \theta_2 < 1$ the best response is *Grab* (formally, $R_2 \cap B_{2,Trust}(R_1) \subseteq [\mathbf{Grab}]$, $B_{1,h^0}(R_2 \cap B_{1,Trust}(R_1)) \subseteq [\alpha = 0]$ and $R_1 \cap B_{1,h^0}(R_2 \cap B_{1,Trust}(R_1)) \subseteq [Don't]$).

One can show that the SE prediction implies $0 < \alpha = \beta = (\frac{3}{2} - \frac{1}{\theta_2}) < \frac{1}{2}$, $\tau = 0$. Thus, SE and forward induction reasoning yield the same path, but very different predictions about how Bob would revise his beliefs off that path.⁴¹

5.2. Rationalizability

The basic rationalizability concept for standard games is equivalent to iterated strict dominance and is motivated by the assumption that players are rational and there is common belief in rationality. Several modifications of rationalizability have been proposed, to handle sequential rationality and to reflect alternative epistemological assumptions.⁴² In psychological games payoffs are affected by hierarchical beliefs, so rationalizability has to be defined as a property of a whole state of the world rather than of strategies. One could, e.g., stipulate that a state $\omega = (s_i, \mu_i)_{i \in N}$ is rationalizable if at ω players are rational and there is common belief in rationality at the beginning of the game.⁴³

One could go on to examine an array of modifications. That is not our goal. Rather, we wish to indicate that the class of psychological games we have defined is amenable to interactive epistemology analysis in principle, and to illustrate the potential cutting power of such an approach. We provide tools to perform a particular forward-induction analysis. Following Battigalli and Siniscalchi [11], we first define a ‘strong belief operator’ SB_i as follows: $SB_i(\emptyset) = \emptyset$ and

$$\forall E \in \mathcal{E}_{-i} \setminus \{\emptyset\}, \quad SB_i(E) = \bigcap_{[h] \cap E \neq \emptyset} B_{i,h}(E).$$

In words, $SB_i(E)$ is the event “player i would believe E conditional on every history that does not contradict E ”⁴⁴; e.g., $SB_i([s_j])$ is the event “player i would believe player j plays s_j at each history h allowed by s_j .”

We are interested in events of the form $SB_i(R_{-i} \cap E)$, where $R_{-i} = \bigcap_{j \neq i} R_j$ is the event that i 's opponents are rational and E is either Ω or some event concerning beliefs, and we consider assumptions like “everybody strongly believes that the opponents are rational.” To write this concisely, we define a *mutual strong belief operator*. Let \mathcal{E} denote the collection of events of the form $E = \bigcap_{i \in N} E_i$ ($E_i \in \mathcal{E}_i$). For example, $R = \bigcap_{i \in N} R_i \in \mathcal{E}$. For each $E = \bigcap_{i \in N} E_i \in \mathcal{E}$, the event “mutual strong belief in E ” is $SB(E) = \bigcap_{i \in N} SB_i(\bigcap_{j \neq i} E_j)$. Note that $SB(E) \in \mathcal{E}$.

⁴⁰ The higher β , the more Bob believes that Ann's choice to trust him is self-interested.

⁴¹ This can happen in standard games too, but for different reasons and with more complex extensive forms (see, e.g., Reny, 1992, [69]). Furthermore, we will show in Section 6 that in some psychological games SE and forward induction yield different paths!

⁴² See, e.g., Battigalli and Bonanno [6], Asheim [2], and references therein.

⁴³ Here is an exact definition: for every event $E = \bigcap_{i \in N} E_i$, $E_i \in \mathcal{E}_i$, let $B(E) = \bigcap_{i \in N} B_{i,h^0}(\bigcap_{j \neq i} E_j)$. Require that $\omega \in R \cap (\bigcap_{k \geq 1} B^k(E))$, where $B^k(E) = B^{k-1}(E)$.

⁴⁴ $SB_i(\cdot)$ is not a monotone operator, and satisfies only a weak form of conjunctiveness $[SB_i(E) \cap SB_i(F) \subseteq SB_i(E \cap F)]$. For more on this, see [11].

We explore the consequences of the following assumptions:

- (0) each player is rational [$= R$],
- (1) mutual strong belief in (0) [$= SB(R)$],
- (2) mutual strong belief in (0) & (1) [$= SB(R \cap SB(R))$],
- (3) mutual strong belief in (0), (1) & (2) [$= SB(R \cap SB(R \cap SB(R)))$],

and so on... Such assumptions are more easily expressed with formulas if we introduce an auxiliary ‘correct strong belief’ operator:

$$\forall E \in \mathcal{E}, \quad CSB(E) = E \cap SB(E).$$

The conjunction of assumptions (0)–(k) corresponds to the event $CSB^k(R)$, where for any $E \in \mathcal{E}$, $CSB^0(E) = E$ and $CSB^k(E) = CSB(CSB^{k-1}(E))$.⁴⁵ Rationalizability is defined by considering the limit as $k \rightarrow \infty$:

Definition 11. A state of the world ω is *rationalizable* if $\omega \in \bigcap_{k \geq 0} CSB^k(R)$.

Battigalli and Siniscalchi [11] show that the strategies consistent with event $CSB^k(R)$ in standard games are those surviving the first $k + 1$ steps of Pearce’s [65] extensive-form rationalizability procedure. This explains the terminology of Definition 11. To illustrate the concept, we note that it captures the forward induction solution of the Trust Game with guilt aversion (either Γ_2 or Γ_3). However, that conclusion requires only two layers of mutual correct strong belief: the solution obtains at all states $\omega \in CSB^2(R)$.

To illustrate the full power of Definition 11, we analyze a *Generalized Trust Game with guilt aversion*, reminiscent of Ben-Porath and Dekel’s [14] money-burning game: Ann can either (evenly) distribute the total surplus of \$2 (action D), or reinvest it in one out of L projects. Project $\ell = 1, \dots, L$ yields $2(1 + \frac{\ell}{L})$. Bob controls surplus distribution. He can either *Grab* or (evenly) *Share*. $Trust_\ell$ denotes the action of investing in project ℓ , and $Share_\ell$ denotes the conditional choice of sharing if Ann invests in project ℓ . Let $\alpha_\ell(\mu_1) = \mu_1^1(Share_\ell|h^0)$ and $\beta_\ell(\mu_2) = \int \alpha_\ell(\mu_1^1)\mu_2^2(d\alpha_\ell(\mu_1^1)|Trust_\ell)$. Assume that Ann’s utility is her material payoff, whereas Bob is averse to guilt. Applying the guilt formula of Subsection 3.3, the players’ utilities are given by

$$\begin{aligned} u_i(D) &= 1, \quad i = 1, 2, \\ u_i(Trust_\ell, Share) &= \left(1 + \frac{\ell}{L}\right), \quad i = 1, 2, \\ u_1(Trust_\ell, Grab) &= 0, \\ u_2(Trust_\ell, Grab) &= 2\left(1 + \frac{\ell}{L}\right) - \theta_2\alpha_\ell\left(1 + \frac{\ell}{L}\right), \end{aligned}$$

Bob (strictly) prefers to share the yield of project ℓ if and only if $\theta_2\beta_\ell > 1$. For $L = 1$ and $\theta_2 = \frac{5}{2}$ we obtain Γ_3 , and the forward induction argument works if and only if $\theta_2 > 2$. When $L > 1$ rationalizability yields the efficient sharing outcome also for much lower values of θ_2 :

⁴⁵ For example, (0) & (1) & (2) is $R \cap SB(R) \cap SB(R \cap SB(R)) = CSB^2(R)$.

Proposition 12. *In the Generalized Trust Game with guilt aversion, if $\theta_2 > 1 + \frac{1}{L}$ then, for every rationalizable state (s_1, μ_1, s_2, μ_2) , $s_1 = \text{Trust}_L$, $s_2 = (\text{Share}_\ell)_{\ell=1}^L$, $\alpha_\ell(\mu_1) = \beta_\ell(\mu_2) = 1$ ($\ell = 1, \dots, L$).*

Proof. Available on request. \square

Our extension of Pearce’s solution concept is well behaved:

Theorem 13. *If psychological utilities are continuous the set $\bigcap_{k \geq 0} \text{CSB}^k(R)$ of rationalizable states is nonempty and compact.*

Proof. By definition

$$\text{CSB}^{k+1}(R) = \text{CSB}(\text{CSB}^k(R)) = \text{CSB}^k(R) \cap \text{SB}(\text{CSB}^k(R)) \subseteq \text{CSB}^k(R).$$

We prove by induction that each element $\text{CSB}^k(R) = \bigcap_{\ell=0}^k \text{CSB}^\ell(R)$ of the nested sequence $\{\text{CSB}^k(R)\}_{k \geq 0}$ is closed and nonempty. Lemma 3 implies Ω is compact; thus, the closed subset $\bigcap_{k \geq 0} \text{CSB}^k(R)$ is compact. Furthermore, the finite intersection property of compact spaces implies $\bigcap_{k \geq 0} \text{CSB}^k(R) \neq \emptyset$.

The argument relies on three preliminary results, proved in Appendix A:

Lemma 14. *Correspondence $r_i : \mathbf{M}_i \rightarrow S_i$ is nonempty valued. If u_i is continuous, r_i has a closed graph and R_i is a nonempty closed set.*

Lemma 15. *For every closed event $E \in \mathcal{E}$, $\text{SB}(E)$ is closed.*

Lemma 16. *Let $\{E^\ell\}_{\ell=0}^{\ell=k}$ be a decreasing sequence of nonempty events in \mathcal{E} ($\emptyset \neq E^k \subseteq E^{k-1} \subseteq \dots \subseteq E^0$), then $\bigcap_{\ell=0}^{\ell=k} \text{SB}(E^\ell)$ is also nonempty.*

For notational convenience let $\text{CSB}^{-1}(E) = \Omega$. We prove by induction that, for each $k \geq 0$, $\text{CSB}^k(R)$ is nonempty closed and can be expressed as

$$\text{CSB}^k(R) = R \cap \left(\bigcap_{\ell=-1}^{k-1} \text{SB}(\text{CSB}^\ell(R)) \right).$$

Basis step. The statement is true for $k = 0$ because by Lemma 14 $\text{CSB}^0(R) = R$ is nonempty closed, and R can be expressed as

$$R = R \cap \Omega = R \cap \text{CSB}^{-1}(R).$$

Inductive step. Suppose the statement is true for each $\ell = 0, \dots, k$, then

$$\begin{aligned} \text{CSB}^{k+1}(R) &= \text{CSB}(\text{CSB}^k(R)) = \text{CSB}^k(R) \cap \text{SB}(\text{CSB}^k(R)) \\ &= R \cap \left(\bigcap_{\ell=-1}^{k-1} \text{SB}(\text{CSB}^\ell(R)) \right) \cap \text{SB}(\text{CSB}^k(R)) \\ &= R \cap \left(\bigcap_{\ell=-1}^k \text{SB}(\text{CSB}^\ell(R)) \right). \end{aligned}$$

By the inductive hypothesis each $CSB^\ell(R)$ is nonempty and closed ($\ell = 0, \dots, k$). By Lemma 15 also $SB(CSB^\ell(R))$ is closed ($\ell = 0, \dots, k$). R is also closed (Lemma 14). Hence $CSB^{k+1}(R)$ is closed. $\{CSB^\ell(R)\}_{\ell=0}^k$ is a decreasing sequence of nonempty events in \mathcal{E} . Therefore Lemma 16 implies that $\bigcap_{\ell=-1}^k SB(CSB^\ell(R)) \neq \emptyset$. Pick any $\omega = (s_i, \mu_i)_{i \in N} \in \bigcap_{\ell=-1}^k SB(CSB^\ell(R))$. Since the latter is just an event about beliefs, modifying the strategies in ω we obtain another state in the same event. By definition of R , $\prod_{i \in N} r_i(\mu_i) \times \{\mu_i\} \subseteq R$. By Lemma 14, $r_i(\mu_i) \neq \emptyset$. We get

$$\emptyset \neq \prod_{i \in N} r_i(\mu_i) \times \{\mu_i\} \subseteq R \cap \left(\bigcap_{\ell=-1}^k SB(CSB^\ell(R)) \right).$$

Hence $CSB^{k+1}(R) \neq \emptyset$. This proves the inductive step, and the theorem. \square

6. Discussion and extensions

We compare our approach with that of GPS (6.1), consider imperfect information, chance moves, and asymmetric information (6.2), own-plan dependence, dynamic inconsistency, and multi-self utility (6.3), and finally discuss assumptions, solution concepts, and avenues for further research (6).

6.1. Comparison with GPS

In Section 2 we presented our framework as a generalization of GPS. This is not literally true. The reason is twofold. First, GPS allow for imperfect information and chance moves. As we show below, these complications can be included in our framework. Second, GPS allow for actual randomization whereas we exclude it. *Prima facie*, this difference may seem immaterial. GPS assume players maximize expected (psychological) utility given beliefs, and in their framework there is no incentive to randomize. It might seem that the only role played by randomization is to guarantee the existence of equilibrium, a result we obtain by looking at equilibrium in beliefs. However, unlike standard games, in psychological games there may be a difference (to a player’s utility) between a belief assigning probability one to a randomized choice that, say, picks a or b with probability $\frac{1}{2}$, and the belief that assigns probability $\frac{1}{2}$ to each of a and b . These beliefs are equivalent if psychological utility functions satisfy a linearity property, which is not satisfied in all applications (see *e.g.* Sebald [72]). Note, however, that one can deal with this by adding to the game explicit moves whereby a player chooses a lottery over elementary actions (which is what Sebald does).

Now look at the version of GPS that is a special case of our framework: psychological games with utilities of the form $u_i : Z \times \bar{\mathbf{M}}_i \rightarrow \mathbb{R}$, where $\bar{\mathbf{M}}_i$ is the space of infinite hierarchies of *initial* beliefs of i , and first-order beliefs are probability measures over pure strategies of i ’s opponents. How much is lost by restricting the analysis to such games? We have argued that many interesting phenomena such as sequential reciprocity, psychological forward induction, and regret cannot be analyzed. However, we can prove a partial equivalence result. Suppose the initial beliefs of others enter the utility, that is, $u_i : Z \times \prod_{j \in N} \bar{\mathbf{M}}_j \rightarrow \mathbb{R}$. Then there is a psychological game with utilities $\bar{u}_i : Z \times \bar{\mathbf{M}}_i \rightarrow \mathbb{R}$ that has the same sequential equilibrium assessments as the former game.⁴⁶ This does not mean that in this class of games conditional higher-order beliefs are

⁴⁶ The intuition is relatively simple: each initial belief hierarchy $\bar{\mu}_i$ induces a probability measure $\bar{f}_i(\bar{\mu}_i) \in \Delta(S_{-i} \times \bar{\mathbf{M}}_{-i})$ which can be used to compute an expectation $\bar{u}_i(z, \bar{\mu}_i)$ of $u_i(z, \bar{\mu}_i, \cdot)$. Since in a consistent assessment there

irrelevant. First, the equivalence result only concerns sequential equilibria, and we argued that the non-equilibrium analysis of psychological games is important (as exemplified by psychological forward induction). Second, our very definition of sequential equilibrium makes essential use of conditional beliefs. Having such beliefs in the framework is conceptually useful as it helps understanding the epistemic assumptions underlying the SE concept.

GPS define a ‘psychological subgame perfect equilibrium’ concept (p-SPE) in two steps. First, they define a ‘psychological Nash equilibrium’ as a situation in which each player’s (randomized) strategy is an *ex ante* best reply to his hierarchy of beliefs and beliefs of all orders are correct. Then, they define a p-SPE as a psychological equilibrium profile $(\sigma_i, \bar{\mu}_i)_{i \in N}$ such that $(\sigma_i)_{i \in N}$ is a subgame perfect equilibrium of the standard game with utility functions $\bar{u}_i(\cdot, \bar{\mu}_i) : Z \rightarrow \mathbb{R}$ ($i \in N$). Remark 10 shows that our SE coincides with p-SPE for games with utility functions *à la* GPS. This result can be extended to games with imperfect information (replacing p-SPE with a similarly defined ‘psychological sequential equilibrium’ concept).

GPS also define a ‘psychological trembling-hand perfect equilibrium’ (p-TPE): $(\sigma_i, \bar{\mu}_i)_{i \in N}$ is a p-TPE if it is a psychological equilibrium and $(\sigma_i)_{i \in N}$ is a trembling-hand perfect equilibrium of the standard game with utility functions $\bar{u}_i(\cdot, \bar{\mu}_i) : Z \rightarrow \mathbb{R}$ ($i \in N$). They show that some games with continuous utility functions have no p-TPE. Yet, our argument to prove existence of SE (see the sketch in Subsection 4.2) shows that continuous games always have equilibria that can be obtained as the limit of ε -equilibria with trembles. These two results are mutually consistent because, even for the special case of games *à la* GPS, they concern different notions of trembling-hand perfection. GPS perturb the strategies, but not the beliefs; we perturb both, as equilibrium beliefs are determined by the strategies *via* the consistency condition.⁴⁷

6.2. Imperfectly observable actions, chance moves, and asymmetric information

We chose to focus on games with observable actions and no chance moves for the sake of simplicity. But our concepts and results carry over to the more general case of games where past actions need not be perfectly observed and chance may play a role (as in GPS). Let $N = \{c, 1, \dots, n\}$ where index c denotes the chance player, and let \mathbf{H}_i be the partition of the finite set of histories H into information sets of player i ($i \neq c$). Assume perfect recall holds. Then the set of strategy profiles consistent with any information set \mathbf{h} of player i must have the form $S(\mathbf{h}) = S_i(\mathbf{h}) \times S_{-i}(\mathbf{h})$ [where $S_i(\mathbf{h}) := \bigcup_{h \in \mathbf{h}} S_i(h)$, $S_{-i}(\mathbf{h}) := \bigcup_{h \in \mathbf{h}} S_{-i}(h)$]. Consider, for the first-order beliefs of player i , the collection of conditioning events $\{F_i : F_i = S_{-i}(\mathbf{h}), \mathbf{h} \in \mathbf{H}_i\}$. Let X_{-i}^{k-1} be the space of $(k-1)$ -order uncertainty for player i ; we obtain the set of k -order cps’ $\Delta^{\mathbf{H}_i}(X_{-i}^{k-1})$, and the k -order uncertainty space $X_{-i}^k = X_{-i}^{k-1} \times \prod_{j \neq i, 0} \Delta^{\mathbf{H}_j}(X_{-j}^{k-1})$. The resulting set of infinite hierarchies of cps’ \mathbf{M}_i is homeomorphic to $\Delta^{\mathbf{H}_i}(S_{-i} \times \mathbf{M}_{-i})$ *via* a belief mapping $f_i = (f_{i,\mathbf{h}})_{\mathbf{h} \in \mathbf{H}_i}$.

Note that \mathbf{H}_i specifies i ’s information at each node/history, including those where i is inactive and in particular including the terminal nodes. This would be redundant in standard games, but is

is ‘common knowledge’ of the hierarchical beliefs, no observation will make the players change their mind about the initial beliefs of the opponents, hence for any consistent assessment u_i and \bar{u}_i have the same set of maximizing actions at each history. (If the game has only one stage u_i and \bar{u}_i are fully equivalent, *i.e.*, they have the same best response correspondences.)

⁴⁷ Kolpin [47] makes a related point.

crucial to model some belief-dependent motivations, such as regret or blame avoidance, whereby players' conditional beliefs matter even if they are inactive.⁴⁸

Let $\sigma_c = \sigma_c(\cdot|h) \in \prod_{h \in H \setminus Z} \Delta^o(A_c(h))$ denote the strictly positive 'objective' probabilities of chance moves. It is routine to define closed and compact subsets of hierarchies $\hat{\mathbf{M}}_i$ ($i = 1, \dots, n$) reflecting the assumption that each players's beliefs about the chance player c are determined by σ_c and there is common certainty of this. The analysis of rationalizability is easily adapted to this environment. The definition of sequential equilibrium requires some care.

An assessment $(\sigma, \mu) = (\sigma_i, \mu_i)_{i=1}^n$ is *consistent* if there is a sequence of strictly positive behavioral strategy profiles $\sigma^k \rightarrow \sigma$ such that for all $i = 1, \dots, n$, $\mathbf{h} \in \mathbf{H}_i$, $s_{-i} \in S_{-i}(\mathbf{h})$,

$$\mu_i^1(s_{-i}|\mathbf{h}) = \lim_{k \rightarrow \infty} \frac{\Pr_{\sigma_c}(s_c) \prod_{j \neq i, c} \Pr_{\sigma_j^k}(s_j)}{\sum_{s'_{-i} \in S_{-i}(h)} \Pr_{\sigma_c}(s'_c) \prod_{j \neq i, c} \Pr_{\sigma_j^k}(s'_j)}$$

(since σ_c is strictly positive, the denominator is positive),⁴⁹ and furthermore for all $\ell > 1$, μ_i^ℓ assigns probability 1 to $\mu_{-i}^{\ell-1}$ (cf. Definition 6). (σ, μ) is a *sequential equilibrium* if it is consistent and for all $i = 1, \dots, n$, $\mathbf{h} \in \mathbf{H}_i$, $s_i^* \in S_i(\mathbf{h})$,

$$\Pr_{\sigma_i}(s_i^*|\mathbf{h}) > 0 \Rightarrow s_i^* \in \arg \max_{s_i \in S_i(\mathbf{h})} E_{s_i, \mu_i}[u_i|\mathbf{h}],$$

where $E_{s_i, \mu_i}[u_i|\mathbf{h}]$ is given by the obvious modification of formula (2). A straightforward adaptation of the trembling-hand argument used to prove Theorem 1 shows that if the utility functions are continuous a sequential equilibrium exists.

This extended framework also allows to analyze situations with *incomplete information*, modeling them as games with asymmetric information about an initial chance move: at the empty history h^0 the only active player is c (chance), $A_c(h^0) = \Theta$, where $\Theta \subseteq \Theta_0 \times \Theta_1 \times \dots \times \Theta_n$, each player i observes only coordinate θ_i of $\theta = (\theta_0, \theta_1, \dots, \theta_n)$; θ may affect payoffs, or actions sets, or the probabilities of future chance moves.⁵⁰ This extension is important for applications. Unless one models interaction within a family or amongst friends, it is probably not realistic to assume that players know one another's psychological propensities. For example, in the analysis of I_2 (or I_3) we assumed that Ann knows that Bob's sensitivity to guilt is $\theta_2 = \frac{5}{2}$, which may be a stretch.⁵¹ Another example comes up in Caplin and Leahy's [22] model of doctor-patient interaction: the patient's well-being depends on his anxiety and on whether or not he likes early resolution of uncertainty; the concerned doctor wants to help the patient but is uninformed of the nature of patient's preferences regarding resolution of uncertainty. Yet another reason to allow for incomplete information is that a player may care about the ex post beliefs of others about some of his characteristics, which are not common knowledge, as in the models of Bernheim

⁴⁸ For example, the terminal information partition is crucial in Battigalli and Dufwenberg's [8] model of guilt from blame and Tadelis' [77] analysis of shame in a Trust Game.

⁴⁹ Kreps and Wilson [51] have a similar condition that refers to conditional probabilities of histories/nodes.

⁵⁰ We do not present a more direct formalization of incomplete information psychological games for reasons of space. A word of caution is in order. When incomplete information is represented as imperfect, asymmetric information about an initial chance move one introduces fictitious *ex ante* beliefs (the initial beliefs in the imperfect information extensive game). This does not affect the equilibrium analysis of standard games, but it is known to affect the rationalizable outcomes (see Battigalli *et al.* [7] for a discussion), and it may as well affect the analysis when players have belief-dependent preferences.

⁵¹ Ample evidence in psychology suggests emotional sensitivities differ among people. See Krohne [52] for a general discussion, and Tangney [78] on guilt specifically.

and Dufwenberg and Lundholm and several other economic models where agents are assumed to have an intrinsic concern about their reputation.⁵²

One final reflection relating to imperfect and asymmetric information before we proceed. Standard games of incomplete information can be regarded as a special case of our framework. We have already clarified in Section 2.1 how psychological games are different from games of incomplete information. But one may wonder whether our framework is “too general” in the following sense: *Is it possible that the same qualitative behavior explained by psychological game models can be explained using standard games of incomplete information?* The answer is No, although it is possible in some cases (see, for example, Levine [54], and Gul and Pesendorfer [40] who put forward models of reciprocity featuring incomplete information about “behavioral types”).⁵³ The best way to see this is to remember that psychological preferences may be affected by players’ beliefs at information sets where they are *inactive*, such as their terminal information sets. Changing this information typically alters predicted behavior. This comparative statics effect (which has been demonstrated experimentally by Tadelis [77]) is impossible in standard games, because behavior predicted by sequential equilibrium or rationalizability in such games depends only on the information of players when they are active.⁵⁴

6.3. Own-plan dependence, dynamic (in)consistency, and multi-self utility

We have argued that it is natural to let a player’s utility depend directly on other players’ plans, but have so far assumed it does not depend directly on his own plan. This allowed us to apply standard dynamic programming techniques and prove Theorem 13. We next show that allowing for own-plan dependence is natural and gives rise to an interesting form of dynamic inconsistency.

Consider the following version of the Trust Game: Ann can either trust Bob or opt out. If she opts out no surplus is created; if she trusts Bob the total surplus is \$4 and Bob can either grab \$3 or let Ann allocate the surplus (action *Leave*). If Bob let Ann choose the allocation, she can either split the surplus or reward Bob, giving him the \$3 he could have grabbed.

Now assume that if Ann gets less money than she expected she feels disappointed and that the anticipation of this negative feeling affects her decisions. This is captured by the following utility function:

$$u_1(z, \mu, s_1) = \pi_1(z) - \theta \left[\sum_{s_2} \mu_1^1(s_2|h^0) \pi_1(\zeta(s_1, s_2)) - \pi_1(z) \right]$$

where θ is a sensitivity parameter and the term in brackets in the Ann’s disappointment at z . Γ_6 in Fig. 6 builds on u_1 , and thereby turns out to exhibit own-plan dependent utility for Ann. We let $u_2(z, \mu, s) = \pi_2(z)$.

The utility assigned by Ann to terminal history (*Trust, Grab*) depends on her initial belief $\alpha = \mu_1^1(\mathbf{Leave}|h^0)$, and on how she plans to behave if Bob leaves for her to allocate the surplus. The own-plan dependence arises because Ann cannot feel disappointed if she plans to reward Bob

⁵² Models with an intrinsic concern for reputation are sometimes presented as a reduced forms of repeated interaction models. See, for example, Morris [62] and Ottaviani and Sorensen [64].

⁵³ Gul and Pesendorfer offer a recursive construction of a canonical space of behavioral types.

⁵⁴ Battigalli and Guaitoli [9] put forward a notion of conjectural (or self-confirming) equilibrium that crucially depends on the terminal information structure. But such an equilibrium concept does not apply to one-shot interactions, as it is meant to capture, with a “static” definition, stable patterns of behavior in situations of recurrent interaction.

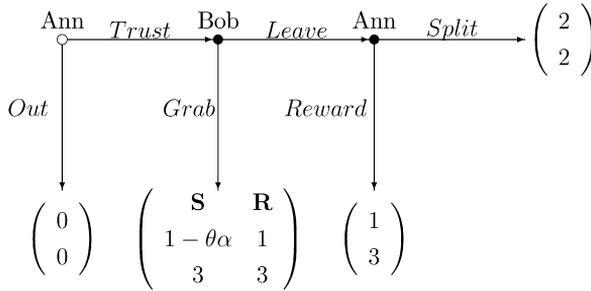


Fig. 6. Psychological game Γ_6 : aversion to disappointment.

(column **R**), since in this case the resulting material payoff is 1 regardless of what Bob chooses. For Ann to be disappointed at terminal history $(Trust, Grab)$ requires that she plans to split and that $\alpha > 0$. In this case she (initially) expects a material payoff which is larger than her material payoff at history $(Trust, Grab)$ $[(1 - \alpha) \times 1 + \alpha \times 2 > 1]$. The disappointment yields utility $(1 - \theta\alpha)$ (column **S**). The expected utility of plan $(Trust, Split)$ is thus $(1 - \alpha) \times (1 - \theta\alpha) + \alpha \times 2$, which could be lower than 1 (in fact, even lower than 0) if θ is large enough. In this case, the *ex ante* expected utility maximizing plan is $(Trust, Reward)$ (which prevents disappointment and yields 1). However, $(Trust, Reward)$ is not dynamically consistent because the best choice after history $(Trust, Leave)$ is to split. As a result there is no strategy that maximizes Ann’s expected utility at the beginning of the game and also at history $(Trust, Leave)$. Ann cannot commit to rewarding Bob, so she should initially maximize under the constraint that she would split in the endgame; the resulting outcome is *Out*. Such ‘consistent planning’ (Strotz [76]) implies that Ann’s strategy be immune to one-shot deviations. Therefore, the example shows that the One-Shot-Deviation principle (cf. Proposition 8) fails with own-plan dependence. A similar kind of dynamic inconsistency arises in Caplin and Leahy’s [21] theory of psychological expected utility, in work by Kőszegi and Rabin [49,50] on reference dependent preferences, and in Mariotti’s [60] abstract analysis of generalized extensive form games. These papers can be shown to be consistent with our extended framework allowing own-plan dependence.

We now present an example where even the multi-self equilibrium fails to exist unless we allow for uncertainty about one’s own strategy.

In game Γ_7 of Fig. 7 a pure strategy of Bob is immune to one-shot deviations if and only if it corresponds to a pure Nash equilibrium of the following game between two selves of Bob:

Bob _ℓ \ Bob _r	L''	R''
L'	3, 0	1, 1
R'	2, 3	2, 1

This companion game has no pure equilibrium. The unique mixed equilibrium is $(\frac{2}{3}L' + \frac{1}{3}R', \frac{1}{2}L'' + \frac{1}{2}R'')$. Working backwards, Ann’s best reply is ℓ . This is the unique multi-self SE of Γ_7 .

This example suggests that in order to make rationality possible in psychological games with own strategy dependence we have to allow a player to be uncertain about her own strategy. This can be done, at the cost of additional complexity, within a richer framework where player i ’s first-order beliefs are defined over S rather than S_{-i} (cf. Battigalli and Siniscalchi [10]).

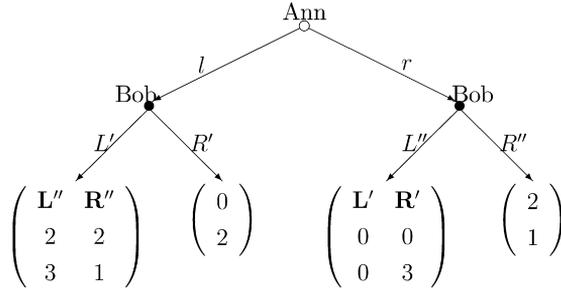


Fig. 7. Psychological game Γ_7 .

Besides own-strategy dependence, a more direct way to allow for dynamic inconsistencies is to adopt a multi-self approach and model a player’s preferences with an array of ‘local’ utility functions $(u_{i,h} : Z \times \mathbf{M} \times S \rightarrow \mathbb{R})_{h \in H \setminus Z}$. The sequential equilibrium analysis of Section 4 applies to this extended framework almost *verbatim*.

This formulation is relevant to reciprocity theory. We have already seen how our basic framework could reproduce Dufwenberg and Kirchsteiger’s theory in an example (Γ_5). However, to handle general games one needs a multi-selves approach, and it is then possible to (essentially) reformulate Dufwenberg and Kirchsteiger’s model (details are available on request).

6.4. Discussion of solution concepts

We have explained in Section 4 that SE relies on the assumption that players never change their beliefs about the (conditional) beliefs of their opponents. This assumption is questionable in standard games, and becomes even more questionable in psychological games where players’ perceptions of opponents’ intentions (“What is my co-player trying to achieve?”) are key.

One way to assess the opponent’s intention is to draw inferences about his unobservable beliefs from his observable actions *via* forward induction (FI) reasoning. It is well known that FI reasoning yields beliefs inconsistent with backward induction reasoning and, more generally, with sequential equilibrium (see, *e.g.*, Reny [69]). Yet in (generic) standard games at least one sequential equilibrium outcome is consistent with FI reasoning.⁵⁵ The following simultaneous-move example shows that this need not be true in games with belief-dependent preferences.

Let $A_1 = \{U, M, D\}$, $A_2 = \{\ell, r\}$. Player 1 (row) only cares about his material payoff, whereas the preferences of 2 (column) also depend on his terminal second-order beliefs. Let $\alpha_1 := \mu_1^1(\ell|h^0)$, $\beta_2(a_1) := E_{\mu_2}[\alpha_1|a_1]$. The game is described by the following table.

	ℓ	r
U	$3, 4[\beta_2(U) - \beta_2(D)]$	$0, 1$
M	$2, 4[\beta_2(U) - \beta_2(D)]$	$2, 1$
D	$0, 4[\beta_2(U) - \beta_2(D)]$	$3, 1$

⁵⁵ For example, in generic finite games with complete and perfect information the backward induction terminal node is the same as the terminal node implied by extensive form rationalizability and the iterated deletion of weakly dominated strategies, even though the implied conditional beliefs about continuation strategies may well be very different, as shown by Reny [69]. For finite games with imperfect information see Govindan and Wilson [38].

In a *sequential equilibrium* $\beta_2(U) - \beta_2(D) = 0$, because all conditional second order beliefs assign probability one to the true value of α_1 . Therefore the SE outcome must be (D, r) .

Forward induction implies that 2 would believe whenever possible that 1 is rational (see Section 5). Note that U (D) is a best reply to α_1 iff $\alpha_1 \geq \frac{2}{3}$ ($\alpha_1 \leq \frac{1}{3}$). Therefore FI implies that $\beta_2(U) - \beta_2(D) \geq \frac{1}{3}$. Hence $4[\beta_2(U) - \beta_2(D)] > 1$, player 2 chooses ℓ , and anticipating this 1 chooses U . The FI outcome is (U, ℓ) .

This example suggests that it is reasonable to relax the consistency condition of SE to allow players to change their higher order beliefs at least when they observe unexpected moves by the opponents. Other conditions on updated beliefs (such as strong belief in rationality) may be imposed to obtain sharp predictions.

Another reason to relax consistency is the following. Suppose that, as suggested at the end of Section 6.3, we explicitly model a player's plan as a (possibly false or uncertain) belief about his own contingent behavior. Then a player's intentions are fully determined by his beliefs about himself and others, and the consistency condition of SE implies that (in games with complete information) players never change their beliefs about the co-players' intentions, no matter what they do. This trivializes the analysis of extensive form psychological games. For example, Bob can respond unkindly to an unexpectedly generous action by Ann simply because he initially believed, and continues to believe, that Ann meant to be unkind to him.

Finally, we note that it would be interesting to approach psychological games from a learning point-of-view. Here is the main issue to be aware of if this task is addressed: It is often argued that players learn to play some equilibrium because through recurrent play they come to hold correct beliefs about the opponents' actions (see, e.g., Fudenberg and Levine [33] and references therein). This may not be enough for psychological games; since payoffs depend on hierarchical beliefs, players would have to be able to learn others' beliefs, but unlike actions beliefs are typically not observable ex post. This suggests that a different notion of equilibrium is worth exploring, whereby beliefs are required to be consistent with observed outcomes, as in the conjectural/self-confirming equilibrium concept.⁵⁶ The interactive epistemology analysis of self-confirming equilibrium in (standard) signaling games by Battigalli and Siniscalchi [11] may provide a useful starting point.

Acknowledgments

We thank Geir Asheim, Oliver Board, Adam Brandenburger, Andrew Caplin, Amanda Friedenberg, Drew Fudenberg, Georg Kirchsteiger, Botond Koszegi, Jing Li, David Pearce, Ludovic Renou, Klaus Ritzberger, Joel Sobel, and participants in several seminars for helpful discussions. For their kind hospitality, we thank the Economics Departments of the Stern School of Business at NYU, the European University Institute (Battigalli), and Göteborg University (Dufwenberg). We gratefully acknowledge financial support from Bocconi University and MIUR (Battigalli) and from the NSF (Dufwenberg).

Appendix A

We start with some preliminaries about rationality and backward induction on belief-induced decision trees, and then prove Lemmata 14, 15 and 16. For any fixed hierarchy of cps' μ_i , we ob-

⁵⁶ See, for example, Battigalli and Guaitoli [9] and Fudenberg and Levine [32].

tain a well-defined decision tree that can be solved by backward induction: define value functions $V_{\mu_i} : H \rightarrow \mathbb{R}$ and $\bar{V}_{\mu_i} : (H \setminus Z) \times A_i \rightarrow \mathbb{R}$ as follows

- For terminal histories $z \in Z$, let

$$V_{\mu_i}(z) := \int_{S_{-i} \times \mathbf{M}_{-i}} u_i(z, \mu_i, \mu_{-i}, s_{-i}) f_{i,z}(\mu_i)(ds_{-i}, d\mu_{-i}).$$

- Assuming that $V_{\mu_i}(h, a)$ has been defined for the immediate successors (h, a) of history h , let

$$\bar{V}_{\mu_i}(h, a_i) := \sum_{a_{-i} \in A_{-i}(h)} \mu_i^1(S_{-i}(h, a_{-i})|h) V_{\mu_i}(h, (a_i, a_{-i}));$$

for each $a_i \in A_i(h)$; then $V_{\mu_i}(h)$ is defined as

$$V_{\mu_i}(h) := \max_{a_i \in A_i(h)} \bar{V}_{\mu_i}(h, a_i).$$

Recall that, for any strategy $s_i \in S_i$, $H_i(s_i) = \{h \in H \setminus Z : s_i \in S_i(h)\}$ denotes the set of histories allowed by s_i . The proof of the following result is available by request:

Lemma 17. *The sequential best reply correspondence $r_i : M_i \rightarrow S_i$ can be characterized as follows*

$$r_i(\mu_i) = \left\{ s_i : \forall h \in H(s_i), s_{i,h} \in \arg \max_{a_i \in A_i(h)} \bar{V}_{\mu_i}(h, a_i) \right\}.$$

Proof of Lemma 14. By Lemma 17 $r_i(\mu_i) = \{s_i : \forall h \in H(s_i), s_{i,h} \in \arg \max_{a_i \in A_i(h)} \bar{V}_{\mu_i}(h, a_i)\}$. Clearly, the RHS is nonempty. Therefore $r_i(\cdot)$ is nonempty-valued and R_i is nonempty. The belief function f_i is continuous (Lemma 3). If u_i is also continuous, then $E_{s_i, \mu_i}[u_i|h]$ is continuous (in μ_i), which implies that R_i is closed. \square

Proof of Lemma 15. We must show that for every closed event $E \in \mathcal{E}_{-i}$, $\text{SB}_i(E)$ is closed. $\text{SB}_i(\emptyset) = \emptyset$, a closed set, by definition. Suppose that $E = \Omega_i \times E_{-i}$ where E_{-i} is nonempty and closed. Recall that $\text{SB}_i(E) = \bigcap_{h: [h] \cap E \neq \emptyset} \text{B}_{i,h}(E)$. For each h ,

$$\text{B}_{i,h}(E) = S_i \times f_{i,h}^{-1}(\Delta(E_{-i} \cap (S_{-i}(h) \times \mathbf{M}_{-i}))) \times \Omega_{-i},$$

where for any measurable space X and any $F \subseteq X$ we let $\Delta(F)$ denote the set of probability measures on X that assign probability one to F . Note that if F is closed, $\Delta(F)$ is also closed. The coordinate function $f_{i,h} : M_i \rightarrow \Delta(\Omega_{-i})$ is continuous and M_{-i} is closed (Lemma 3); hence $E_{-i} \cap (S_{-i}(h) \times M_{-i})$, $\Delta(E_{-i} \cap (S_{-i}(h) \times M_{-i}))$ and $f_{i,h}^{-1}(\Delta(E_{-i} \cap (S_{-i}(h) \times \mathbf{M}_{-i})))$ are closed. It follows that $\text{B}_{i,h}(E)$ ($h \in H$) and $\text{SB}_i(E)$ are closed. \square

Proof of Lemma 16. Let $\{E^\ell\}_{\ell=0}^{\ell=k}$ be a decreasing sequence of nonempty events in \mathcal{E} ($\emptyset \neq E^k \subseteq E^{k-1} \subseteq \dots \subseteq E^0$); we show that $\bigcap_{\ell=0}^{\ell=k} \text{SB}(E^\ell)$ is also nonempty. For each ℓ and i , $E^\ell \in \mathcal{E}$ can be written $E^\ell = E_i^\ell \times E_{-i}^\ell$, where $E_{-i}^\ell \subseteq \Omega_{-i}$, and by definition of $\text{SB}(\cdot)$

$$\bigcap_{\ell=0}^{\ell=k} \text{SB}(E^\ell) = \bigcap_{i \in N} \bigcap_{\ell=0}^{\ell=k} \text{SB}_i(\Omega_i \times E_{-i}^\ell).$$

Therefore we must show that $\bigcap_{\ell=0}^{\ell=k} \text{SB}_i(\Omega_i \times E_{-i}^\ell) \neq \emptyset$ ($i \in N$). Let $\Delta^H(\Omega_{-i}; E_{-i}^\ell)$ denote the set of cps' $\mu \in \Delta^H(\Omega_{-i})$ such that $\mu(E_{-i}^\ell|h) = 1$ for each h such that $E_{-i}^\ell \cap (S_{-i}(h) \times M_{-i}) \neq \emptyset$. Note that

$$\bigcap_{\ell=0}^{\ell=k} \text{SB}_i(\Omega_i \times E_{-i}^\ell) = S_i \times f_i^{-1} \left(\bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell) \right) \times \Omega_{-i}.$$

We show below that $\bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell) \neq \emptyset$. Since f_i is onto (Lemma 3), it follows that $f_i^{-1}(\bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell)) \neq \emptyset$. Hence $\bigcap_{\ell=0}^{\ell=k} \text{SB}_i(\Omega_i \times E_{-i}^\ell) \neq \emptyset$.

We show that $\bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell) \neq \emptyset$ with a recursive construction. Say that h is ‘reached’ by probability measure $\nu \in \Delta(\Omega_{-i})$ if $\nu(S_{-i}(h) \times M_{-i}) > 0$. Note that if h is reached by ν , every predecessor of h is also reached by ν . Say that $\mu(\cdot|h)$ is ‘derived’ from ν , where ν reaches h , if for every Borel set $F_{-i} \subseteq \Omega_{-i}$

$$\mu(F_{-i}|h) = \frac{\nu(F_{-i} \cap (S_{-i}(h) \times M_{-i}))}{\nu(S_{-i}(h) \times M_{-i})}.$$

Pick any probability measure ν in the (nonempty) set $\Delta(E_{-i}^k)$. For each h reached by ν let $\mu(\cdot|h)$ be derived from ν . Thus, $\mu(\cdot|h)$ has been defined for a nonempty set of histories closed w.r.t. precedence (that is, if h is in the set every predecessor of h is in the set), the set is nonempty because it contains the initial history h^0 . Now suppose that $\mu(\cdot|h)$ has been defined for some set of histories \hat{H} closed w.r.t. precedence. If $\hat{H} \neq H$, for each $h \in H \setminus \hat{H}$ such that the immediate predecessor of h belongs to \hat{H} , pick a probability measure ν_h in the set $\Delta(E_{-i}^{\ell(h)} \cap (S_{-i}(h) \times M_{-i}))$, where $\ell(h)$ is the highest index $\ell \in \{-1, 0, \dots, k\}$ such that $E_{-i}^\ell \cap (S_{-i}(h) \times M_{-i}) \neq \emptyset$, and by convention we let $E_{-i}^{-1} = \Omega_{-i}$. Let $\mu(\cdot|h')$ be derived from ν_h whenever h' weakly follows h and is reached by ν_h . Now $\mu(\cdot|h)$ is defined for a set of histories \hat{H}' closed under the precedence relation and strictly larger than \hat{H} . Proceed in this way until the whole H is covered. We claim that the resulting vector of probability measures $(\mu(\cdot|h))_{h \in H}$ is a cps $\mu \in \bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell)$.

To see that $(\mu(\cdot|h))_{h \in H}$ is a cps we only have to check that the ‘chain rule’ (3) in Definition 1 holds. Suppose that h precedes h' . To write formulas more transparently, let $C = S_{-i}(h) \times M_{-i}$, $C'_{-i} = S_{-i}(h') \times M_{-i}$, $\mu(\cdot|h) = \mu(\cdot|C)$, $\mu(\cdot|h') = \mu(\cdot|C'_{-i})$. Since h precedes h' , $S_{-i}(h') \subseteq S_{-i}(h)$, hence $C'_{-i} \subseteq C$. If h' is not reached by $\mu(\cdot|C)$ then (3) holds trivially as $0 = 0$. If h' is reached by $\mu(\cdot|C)$, then $\mu(\cdot|C)$ and $\mu(\cdot|C'_{-i})$ are both derived from the same measure – say $\nu \in \Delta(\Omega_{-i})$ – reaching both h and h' ; thus, for every Borel set $F_{-i} \subseteq C'_{-i}$

$$\mu(F_{-i}|C) = \frac{\nu(F_{-i})}{\nu(C)} = \frac{\nu(F_{-i})}{\nu(C'_{-i})} \frac{\nu(C'_{-i})}{\nu(C)} = \mu(F_{-i}|C'_{-i})\mu(C'_{-i}|C).$$

To see that $\mu \in \bigcap_{\ell=0}^{\ell=k} \Delta^H(\Omega_{-i}; E_{-i}^\ell)$, note that by construction $\mu(E^{\ell(h)}|h) = 1$ for all $h \in H$. Suppose that, for any index $\ell \in \{0, \dots, k\}$ and any $h \in H$, $E_{-i}^\ell \cap (S_{-i}(h) \times M_{-i}) \neq \emptyset$. Then $\ell(h) \geq \ell$ and $\mu(E^\ell|h) \geq \mu(E^{\ell(h)}|h) = 1$; hence $\mu(E^\ell|h) = 1$ as desired. \square

References

[1] W. Albers, R. Pope, R. Selten, B. Vogt, Experimental evidence for attractions to chance, Ger. Econ. Rev. 1 (2000) 113–130.
 [2] G.B. Asheim, The Consistent Preferences Approach to Deductive Reasoning in Games, Springer-Verlag, Berlin, 2006.

- [3] G. Attanasi, R. Nagel, A survey of psychological games: Theoretical findings and experimental evidence, in: A. Innocenti, P. Sbriglia (Eds.), *Games, Rationality and Behaviour. Essays on Behavioural Game Theory and Experiments*, Palgrave MacMillan, Houndmills, 2007, pp. 204–232.
- [4] R.J. Aumann, A. Brandenburger, Epistemic conditions for Nash equilibrium, *Econometrica* 63 (1995) 1161–1180.
- [5] M. Bacharach, G. Guerra, D.J. Zizzo, The self-fulfilling property of trust: An experimental study, *Theory Dec.* 63 (2007) 349–388.
- [6] P. Battigalli, G. Bonanno, Recent results on belief, knowledge and the epistemic foundations of game theory, *Res. Econ.* 53 (1999) 149–226.
- [7] P. Battigalli, A. Di Tillio, E. Grillo, A. Penta, *Interactive epistemology and solution concepts for games with asymmetric information*, Mimeo, Bocconi University, 2007.
- [8] P. Battigalli, M. Dufwenberg, Guilt in games, *Amer. Econ. Rev.* 97 (2007) 170–176, *Papers & Proceedings*.
- [9] P. Battigalli, D. Guaitoli, Conjectural equilibria and rationalizability in a game with incomplete information, in: P. Battigalli, A. Montesano, F. Panunzi (Eds.), *Decisions, Games and Markets*, Kluwer Academic Publishers, Dordrecht, 1997, pp. 97–124.
- [10] P. Battigalli, M. Siniscalchi, Hierarchies of conditional beliefs and interactive epistemology in dynamic games, *J. Econ. Theory* 88 (1999) 188–230.
- [11] P. Battigalli, M. Siniscalchi, Strong Belief and Forward Induction Reasoning, *J. Econ. Theory* 106 (2002) 356–391.
- [12] D.E. Bell, Regret in decision making under uncertainty, *Oper. Res.* 30 (1982) 961–981.
- [13] D.E. Bell, Disappointment in decision making under uncertainty, *Oper. Res.* 33 (1985) 1–27.
- [14] E. Ben-El-Mechaieq, E. Dekel, Signaling future actions and the potential for sacrifice, *J. Econ. Theory* 57 (1992) 36–51.
- [15] D. Bernheim, Rationalizable strategic behavior, *Econometrica* 52 (1984) 1007–1028.
- [16] D. Bernheim, A theory of conformity, *J. Polit. Economy* 102 (1994) 841–877.
- [17] J. Bouckaert, G. Dhaene, Sequential reciprocity in two-player, two-stages games: An experimental analysis, Mimeo, University of Antwerp and Catholic University of Leuven, 2007.
- [18] A. Brandenburger, E. Dekel, Hierarchies of beliefs and common knowledge, *J. Econ. Theory* 59 (1993) 189–198.
- [19] A. Caplin, Fear as a policy instrument: Economic and psychological perspectives on intertemporal choice, in: G. Loewenstein, D. Read, R.F. Baumeister (Eds.), *Time and Decision*, Russell Sage Foundation, New York, NY, 2003.
- [20] A. Caplin, K. Eliaz, AIDS policy and psychology: A mechanism design approach, *RAND J. Econ.* 34 (2003) 631–646.
- [21] A. Caplin, J. Leahy, Psychological expected utility theory and anticipatory feelings, *Quart. J. Econ.* 116 (2001) 55–79.
- [22] A. Caplin, J. Leahy, The supply of information by a concerned expert, *Econ. J.* 114 (2004) 487–505.
- [23] G. Charness, M. Dufwenberg, Promises and partnership, *Econometrica* 74 (2006) 1579–1601.
- [24] M. Dufwenberg, Marital investment, time consistency and emotions, *J. Econ. Behav. Organ.* 48 (2002) 57–69.
- [25] M. Dufwenberg, Time-consistent Wedlock with endogenous trust, *Doctoral Dissertation*, Uppsala University, 1995.
- [26] M. Dufwenberg, S. Gaechter, H. Hennig-Schmidt, The framing of games and the psychology of play, Mimeo, University of Arizona, University of Nottingham, and University of Bonn, 2007.
- [27] M. Dufwenberg, U. Gneezy, Measuring beliefs in an experimental lost wallet game, *Games Econ. Behav.* 30 (2000) 163–182.
- [28] M. Dufwenberg, G. Kirchsteiger, Reciprocity and wage undercutting, *Europ. Econ. Rev.* 44 (2000) 1069–1078.
- [29] M. Dufwenberg, G. Kirchsteiger, A theory of sequential reciprocity, *Games Econ. Behav.* 47 (2004) 268–298.
- [30] M. Dufwenberg, M. Lundholm, Social norms and moral hazard, *Econ. J.* 111 (2001) 506–525.
- [31] A. Falk, U. Fischbacher, A theory of reciprocity, *Games Econ. Behav.* 54 (2006) 293–315.
- [32] D. Fudenberg, D. Levine, Self-confirming equilibrium, *Econometrica* 61 (1993) 523–545.
- [33] D. Fudenberg, D. Levine, *The Theory of Learning in Games*, MIT Press, Cambridge, MA, 1998.
- [34] D. Fudenberg, J. Tirole, *Game Theory*, MIT Press, Cambridge, MA, 1991.
- [35] J. Geanakoplos, The Hangman paradox and the Newcomb's paradox as psychological games, *Cowles Foundation Discussion Paper No. 1128*, 1996.
- [36] J. Geanakoplos, D. Pearce, E. Stacchetti, Psychological games and sequential rationality, *Games Econ. Behav.* 1 (1989) 60–79.
- [37] I. Gilboa, D. Schmeidler, Information dependent games: Can common sense be common knowledge? *Econ. Letters* 27 (1988) 215–221.
- [38] S. Govindan, R. Wilson, On forward induction, Mimeo, Department of Economics, University of Iowa, and Stanford Business School, Stanford University, 2008.
- [39] G. Guerra, D.J. Zizzo, Trust responsiveness and beliefs, *J. Econ. Behav. Organ.* 55 (2004) 25–30.

- [40] F. Gul, W. Pesendorfer, The canonical type space for interdependent preferences, Mimeo, Princeton University, 2007.
- [41] J. Harsanyi, Games of incomplete information played by Bayesian players. Parts I, II, III, *Manage. Sci.* 14 (1967–1968) 159–182, 320–334, 486–502.
- [42] P.H. Huang, H.-M. Wu, More order without more law: A theory of social norms and organizational cultures, *J. Law, Econ., Organ.* 12 (1994) 390–406.
- [43] S. Huck, D. Kübler, Social pressure, uncertainty, and cooperation, *Econ. Governance* 1 (2000) 199–212.
- [44] E. Karni, Utility theory with probability dependent outcome valuation, *J. Econ. Theory* 57 (1992) 111–124.
- [45] E. Karni, E.E. Schlee, Utility theory with probability dependent outcome valuation: Extensions and applications, *J. Risk Uncertainty* 10 (1995) 127–142.
- [46] A. Kechris, *Classical Descriptive Set Theory*, Springer-Verlag, Berlin, 1995.
- [47] V. Kolpin, Equilibrium refinements in psychological games, *Games Econ. Behav.* 4 (1992) 218–231.
- [48] B. Köszegi, Emotional agency, *Quart. J. Econ.* 12 (2006) 121–156.
- [49] B. Köszegi, M. Rabin, A model of reference-dependent preferences, *Quart. J. Econ.* 121 (2006) 1133–1166.
- [50] B. Köszegi, M. Rabin, Reference-dependent risk attitudes, *Amer. Econ. Rev.* 97 (2007) 1047–1073.
- [51] D. Kreps, R. Wilson, Sequential equilibrium, *Econometrica* 50 (1982) 863–894.
- [52] H.W. Krohne, Individual differences in emotional reactions and coping, in: R.J. Davidson, K.R. Scherer, H.H. Goldsmith (Eds.), *Handbook of Affective Sciences*, Oxford University Press, Oxford, 2003, pp. 698–725.
- [53] H.W. Kuhn, Extensive games and the problem of information, in: H.W. Kuhn, A.W. Tucker (Eds.), *Contributions to the Theory of Games II*, Princeton University Press, Princeton, NJ, 1953, pp. 193–216.
- [54] D.K. Levine, Modeling altruism and spitefulness in experiments, *Rev. Econ. Dynam.* 1 (1998) 593–622.
- [55] J. Li, The power of convention: A theory of social preferences, *J. Econ. Behav. Organ.* 65 (2008) 489–505.
- [56] G. Loomes, R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty, *Econ. J.* 92 (1982) 805–824.
- [57] G. Loomes, R. Sugden, Disappointment and dynamic consistency in choice under uncertainty, *Rev. Econ. Stud.* 53 (1986) 271–282.
- [58] M. Machina, ‘Rational’ decision making versus ‘rational’ decision modeling, *J. Math. Psychology* 24 (1981) 163–175.
- [59] M. Machina, Dynamic consistency and non-expected utility models of choice under uncertainty, *J. Econ. Lit.* 24 (1989) 163–175.
- [60] M. Mariotti, Strategic extensive form games: Definitions and comments on sequential rationality, with an application to rational procrastination, Mimeo, Queen Mary, University of London, 2003.
- [61] J.-F. Mertens, S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory* 14 (1985) 1–29.
- [62] S. Morris, Political correctness, *J. Polit. Economy* 109 (2001) 231–265.
- [63] M. Osborne, A. Rubinstein, *A Course in Game Theory*, MIT Press, Cambridge, MA, 1994.
- [64] M. Ottaviani, P. Sorensen, Reputational cheap talk, *RAND J. Econ.* 37 (2006) 155–175.
- [65] D. Pearce, Rationalizable strategic behavior and the problem of perfection, *Econometrica* 52 (1984) 1029–1050.
- [66] R. Pope, Biases from omitted risk effects in standard gamble utilities, *J. Health Econ.* 23 (2004) 695–735.
- [67] M. Rabin, Incorporating fairness into game theory and economics, *Amer. Econ. Rev.* 83 (1993) 1281–1302.
- [68] M. Rabin, Psychology and economics, *J. Econ. Lit.* 83 (1998) 11–46.
- [69] P. Reny, Backward induction, normal form perfection and explicable equilibria, *Econometrica* 60 (1992) 626–649.
- [70] A. Renyi, On a new axiomatic theory of probability, *Acta Math. Acad. Sci. Hungaricae* 6 (1955) 285–335.
- [71] B.J. Ruffle, Gift giving with emotions, *J. Econ. Behav. Organ.* 39 (1999) 399–420.
- [72] A. Sebald, Procedural concerns and reciprocity, ECARES Discussion Paper 2007/62, Université Libre de Bruxelles, 2007.
- [73] U. Segal, J. Sobel, Tit for tat: Foundations for preferences for reciprocity in strategic settings, *J. Econ. Theory* 136 (2007) 197–216.
- [74] R. Selten, Re-examination of the perfectness concept for equilibrium points in extensive games, *Int. J. Game Theory* 4 (1975) 25–55.
- [75] J. Sobel, Interdependent preferences and reciprocity, *J. Econ. Lit.* 43 (2005) 392–436.
- [76] R.H. Strotz, Myopia and inconsistency in dynamic utility maximization, *Rev. Econ. Stud.* 23 (1956) 165–180.
- [77] S. Tadelis, The power of shame and the rationality of trust, Mimeo, UC Berkeley, 2007.
- [78] J.P. Tangney, Recent advances in the empirical study of shame and guilt, *Amer. Behav. Sci.* 38 (1995) 1132–1145.