Carsten S. Nielsen and Alexander Sebald*

Simple Unawareness in Dynamic Psychological Games

DOI 10.1515/bejte-2015-0011

Abstract: Building on Battigalli and Dufwenberg (2009)'s framework of dynamic psychological games and the progress in the modeling of dynamic unawareness by Heifetz, Meier, and Schipper (2013a) we model and analyze the impact of asymmetric awareness in the strategic interaction of players motivated by reciprocity and guilt. Specifically we characterize extensive-form games with psychological payoffs and simple unawareness, define extensive-form rationalizability and, using this, show that unawareness has a pervasive impact on the strategic interaction of psychologically motivated players. Intuitively, unawareness influences players' beliefs concerning, for example, the intentions and expectations of others which in turn impacts their behavior.

Keywords: unawareness, psychological payoffs, extensive-form rationalizability
JEL Classification: C72, C73, D80

1 Introduction

Recent lab and field evidence suggests that people not only care about the monetary consequences of their actions, but that their behavior is also driven by psychological payoffs (for example, Fehr, Kirchsteiger, and Riedl 1993; Charness and Dufwenberg 2006; Falk, Fehr, and Fischbacher 2008; Bellemare, Sebald, and Strobel 2011). Two prominent examples of psychological payoffs in the hitherto existing literature are reciprocity (for example, Rabin 1993; Dufwenberg and Kirchsteiger 2004; Falk and Fischbacher 2006) and guilt aversion (for example, Charness and Dufwenberg 2006; Battigalli and Dufwenberg 2007). Departing from the strictly consequentialist tradition in economics, Geanakoplos, Pearce, and Stacchetti (1989) and Battigalli and Dufwenberg

***Corresponding author: Alexander Sebald,** Department of Economics, University of Copenhagen, Ø ster Farimagsgade 5, Building 26, DK-1353, Copenhagen K, Denmark, E-mail: alexander.sebald@econ.ku.dk
Carsten S. Nielsen, Department of Economics, University of Copenhagen, Ø ster Farimagsgade 5, Building 26, DK-1353, Copenhagen K, Denmark, E-mail: carsten.nielsen@econ.ku.dk

(2009) present general frameworks for analyzing the strategic interaction of players with psychological payoffs: "psychological games". Roughly speaking, psychological games are games in which players' preferences depend upon players' beliefs about the strategies that are being played, players' beliefs about the beliefs of others about the strategies that are being played, and so on.

A widely unspoken assumption that is underlying game-theoretic analyses, and therefore also the analyses of psychological games, is that players are aware of *all* facts characterizing the strategic environment they are in. However, in many real life situations this is not the case. People often have asymmetric awareness levels concerning their own as well as the feasible choices of others although they are part of the same strategic environment. People are frequently surprised in the sense that they become aware of new strategic alternatives by observing actions they had previously been unaware of. In recent years different models of unawareness have been proposed showing the importance of unawareness for individual decision making problems as well as the strategic interaction of players in standard (non-psychological) games (for example, Fagin and Halpern 1988; Dekel, Lipman, and Rustichini 1998; Modica and Rustichini 1999; Halpern 2001; Heifetz, Meier, and Schipper 2006, 2013a, 2013b; Halpern and Rêgo 2006, 2008, 2009; Heifetz, Meier, and Schipper 2008; Li 2009; Grant and Quiggin 2013).

However, it is not only in standard games that unawareness is important. We show in our analysis that unawareness has a profound and distinct impact on the strategic interaction of players in psychological games. To see this consider the following intuitive example: Imagine two friends, *Ann* and *Bob*. Assume it is *Bob*'s birthday, he is planning a party and would be very happy, if *Ann* could come. Unfortunately *Bob*'s birthday coincides with the date of *Ann*'s final exam at university. She can either decide to take the exam the morning after *Bob*'s party or two weeks later at a second date. *Ann* is certain that *Bob* would feel let down, if she were to cancel his party without having a very good excuse. Quite intuitively, although *Ann* would really like to get over her exam as soon as possible, she might anticipate feeling guilty from letting down *Bob* if she canceled his party to take the exam the following morning. As a consequence, *Ann* might choose the second date to avoid letting *Bob* down. In contrast, consider now the following variant of the same example: *Ann* knows that *Bob* is unaware of the second date. In this situation *Ann* might choose to take the exam on the first date and not feel guilty. Since *Bob* is unaware of the second date and the final exam is a good excuse, he does not expect *Ann* to come. *Ann* knows this and, hence, does not feel guilty as *Bob* is not let down. In fact, if she were certain that *Bob* would never become aware of the second date, she probably had an emotional incentive to leave him unaware in order not to

raise his expectations. That is, she had an incentive not to make him aware of the fact that she actually has the time to come to his party, but just wants to get over her exam. Interestingly, if *Ann* were only interested in her own payoff in this strategic situation with unawareness, she would not care whether *Bob* is or will become aware of the second date. She would simply not attend his party irrespective of *Bob*'s awareness. Only her belief-dependent feeling of guilt towards *Bob* creates the strong emotional incentive to leave him unaware.

Bob's unawareness concerning *Ann*'s ability to come to his party and, connectedly, *Ann*'s incentive not to tell him about the second date intuitively highlight the focus of our analysis. We analyze the influence and importance of unawareness concerning feasible paths of play for the strategic interaction of players in psychological games. To simplify the analysis we concentrate on two-player strategic environments with simple unawareness. More specifically, building on Battigalli and Dufwenberg (2009)'s framework of dynamic psychological games and the recent progress in the modeling of (Heifetz, Meier, and Schipper 2006, 2008, 2013a, 2013b), we define a two-player model in which players are motivated by psychological payoffs and one player is potentially unaware of certain feasible paths of play. More specifically, in our two-player setting with simple unawareness, one player is initially aware of all paths of play, whereas the other player is initially unaware of some paths of play. We assume that the aware player is aware of the unaware player's unawareness, but the unaware player is not. We refrain from moves of chance implying that players' information sets are singletons. We restrict ourselves in this way to clearly investigate and highlight the pervasive role of asymmetric awareness on the strategic interaction of players motivated by belief-dependent preferences. Extensions to broader settings are of course feasible, and will definitely allow for the analysis of a lot of interesting applications, but we leave it for future research to explore these directions. Using our framework we provide different examples highlighting the role of unawareness in the strategic interaction of players motivated by reciprocity à la Dufwenberg and Kirchsteiger (2004) and guilt aversion à Battigalli and Dufwenberg (2007).We limit ourselves to two-player environments and simple asymmetric awareness scenarios in order to intuitively introduce our model and clearly uncover the role of unawareness without burdening the analysis with technical issues arising in strategic environments allowing for more players and more complex unawareness.

Our examples demonstrate that the strategic behavior of players motivated by psychological payoffs crucially depends on their awareness concerning the strategic environment they are in, their perception concerning the awareness of others, their perception concerning the perception of others, and so on – a fact that implies both an opportunity as well as a challenge to analyses

empirically investigating the strength and nature of psychological payoffs. On the one hand, in line with experimental evidence suggesting that people are more prone to selfish choices if they believe that others will remain unaware of them (for example, Dana, Cain, and Dawes 2006, Dana, Weber, and Kuang 2007; Broberg, Ellingsen, and Johannesson 2007; Andreoni and Bernheim 2009; Lazear, Malmendier, and Weber 2012),our examples show that redchanging the awareness of players that are motivated by psychological payoffs leads to intuitive and testable predictions distinct from predictions based on consequentialist preferences like selfishness and inequality aversion (Fehr and Schmidt 1999). On the other hand, it poses a challenge for experimental investigations in relatively uncontrolled environments like the field or the Internet. As also seen in our introductory example, not controlling for *Ann*'s perception concerning *Bob*'s awareness might lead to wrong inferences concerning *Ann*'s inclination to feel guilty towards *Bob*. Furthermore, our examples reveal that over and above the actual choices that are made, managing other people's awareness levels has to be understood as an integral and important part of any strategic interaction. By managing other's awareness, we influence the others' expectations and perceptions concerning our intentions, which in turn influences their behavior.

We start out by formulating a model concentrating on two-player extensive-forms with complete information, observable actions and no chance moves. To allow for unawareness we use a standard extensive-form representing the game tree and a subtree thereof, and define extensive-forms with simple unawareness with the help of singleton information sets. These singleton information sets describes at each decision node in the game tree, and copy thereof in the subtree, the frame of mind of a player. Our two-player extensive-forms are in essence a special case of Heifetz, Meier, and Schipper (2013a)'s generalized extensive-forms, and therefore embeddable in their setting.Of course, our extensive-form with unawareness is not typically common knowledge among players, and therefore should be interpreted from the modeler's point of view. In fact, any game that does not explicitly distinguish between the players' description of the strategic environment and the modeler's will fail to capture (see Dekel, Lipman, and Rustichini 1998).

Having defined our class of two-player extensive-forms with unawareness, we formally characterize psychological payoffs in our setting. In synthesis, we define a player's strategies and conditional beliefs about the other player's pure strategies (first-order beliefs), beliefs about the other player's beliefs (second-order beliefs), and so on. The infinite hierarchy of conditional beliefs that we define takes player's awareness, players's perception regarding the other's awareness, and so forth, into account and is used for the general specification

of our psychological payoffs and, hence, the characterization of our class of dynamic psychological games with simple unawareness. As mentioned above, specific types of psychological payoffs that can be embedded in our model are among others reciprocity and guilt aversion.

Dufwenberg and Kirchsteiger (2004), Battigalli and Dufwenberg (2007) and Sebald (2010) propose sequential equilibrium as a solution concept for their psychological games. However, assuming equilibrium play is very demanding in strategic environments involving unawareness. The implicit assumption made when imposing sequential equilibrium on strategic settings with unawareness is that if a player becomes aware of more during the game, he will compute new equilibrium beliefs not rationalizing, for example, why the other player made him aware. Sequential equilibrium only requires a player to reason about the other player's future behavior. For this reason, we impose extensive-form rationalizability (Pearce 1984), which embodies forward induction, as a solution concept for our psychological games with simple unawareness. Extensive-form rationalizability implies, that along each feasible path of play, every active player is always certain that the other player sequential best responds, certain that the other player is certain that he sequential best responds, and so on. If a player finds himself at some information set, where the other player's strategies that could lead to that information set are inconsistent with the players previous certainty in the other player's best response, then the player seeks a best rationalization which could have led to that information set (Battigalli 1997; Battigalli and Siniscalchi 2002). That is, if the player is "surprised" by the other player's unexpected action, and cannot use Bayesian updating, then he forms new beliefs that justify this observed inconsistency. In its simplest form, forward-induction reasoning involves the assumption that, upon observing an unexpected (but undominated) action of the other player, a player maintains the working hypothesis that the latter is a sequential best response. The best rationalization principle captures precisely this type of argument.

After having defined our model, the solution concept and two prominent notions of psychological payoffs, reciprocity and guilt aversion, we describe two examples to highlight the role of unawareness in the interaction of players motivated by reciprocity and guilt aversion. First, we consider a version of the sequential prisoners dilemma also analyzed by Dufwenberg and Kirchsteiger (2004) featuring a reciprocal second mover, *Bob*, who is unaware that the first mover, *Ann*, can defect.[1] Different to Dufwenberg and Kirchsteiger (2004)'s

1 Note that this awareness scenario is similar to the finitely repeated prisoner's dilemma with unawareness analyzed in Feinberg (2004).

analysis assuming full awareness, it is shown that as long as *Bob* is unaware of the fact that *Ann* could have defected, he defects independent of his sensitivity to reciprocity – even when *Ann* chooses to cooperate. The way he perceives *Ann*'s kindness does not only depend on what she does, but also on what *Bob* thinks she could have done given his awareness of the strategic situation. *Ann* anticipates this and defects as long as she cannot cooperate and simultaneously make *Bob* aware of the fact that she could have defected. As a second example, we formally revisit the example of *Ann* not wanting to come to *Bob*'s party because of the exam, and consider how her aversion to guilt affects her choice. If *Ann* knows that *Bob* will not be "let down" if he is unaware of the second exam date, then *Ann* does not feel any guilt towards *Bob* if she chooses not to come to the party. Because *Bob* is unaware of the second exam date he foresees that *Ann* will write the exam on the first date, and thus, does not expect *Ann* to come. Both examples highlight that unawareness in the interaction of players with psychological payoffs leads to very intuitive behavioral predictions distinct from predictions using non-psychological preferences or no unawareness. Furthermore, it becomes evident that managing others' awareness levels is an important and integral part of strategic interactions of players motivated by psychological payoffs.

The organization of the paper is as follows: In Section 2 we introduce dynamic games with simple unawareness. Following this, in Section 3, we define the hierarchies of beliefs and psychological payoffs. Section 4 contains the definition of our solution concept: extensive-form rationalizability. In Section 5 we give two examples of how our model can be applied. Finally, Section 6 concludes.

2 Dynamic Games with Simple Unawareness

A finite extensive-form game with singleton information sets and no chance moves, called the *simple unawareness game*, is played by two players, player i and some other player j, who move one at a time and have possibly different views on the feasible paths of play. One player is initially aware of all paths of play, whereas the other player is initially unaware of some paths of play. The aware player knows that the other player is unaware, while the unaware player is unaware of his own unawareness and thinks that the other player is aware of the same as he. The game specifies material payoffs for each player at each terminal node. These payoffs describe the material consequences of the players' actions, not their preferences. The players' psychological payoffs will be introduced in Section 3.

The simple unawareness game we consider adapts Heifetz, Meier, and Schipper (2013a)'s generalization of the standard extensive-form game (Fudenberg and Tirole 1991, chapter 3.3). By simple unawareness we mean that our game is restricted to unawareness with just two trees – a game tree specifying all physical paths of play and a subtree thereof. Both will be described in detail below. Our game allows for a parsimonious analysis of the applications we consider. More general applications and extensions are certainly very interesting (for example, delusion or awareness of unawareness) but are left for future research in order to not burden our analysis with additional notational complexity.

Game tree. The physical paths of play are given by a finite set T^* of nodes together with a binary relation $\prec$ on T^* that represents precedence. The binary relation $\prec$ must be a partial order, and $(T^*, \prec)$ must form an arborescence: the relation $\prec$ totally orders the predecessors of each member of T^*. The order of play thus constitute a game tree that begins at an initial node with no predecessor and then proceeds along some path from node to an immediate successor, terminating when a node with no successor is reached. The various paths give the various possible orders of play. Let Z denote the set of terminal nodes, and let $X = T^* \setminus Z$ be the set of decision nodes.

Moves in the game tree. To represent the choices available to players at decision nodes, we have a finite set A of actions and a function ψ that labels each non-initial decision node x with the last action taken to reach it. We require that ψ be one-to-one on the set of immediate successors of each decision node x, so that different successors correspond to different actions. Let $A(x)$ denote the set of feasible actions at x. Actions are labeled so that $A(x) \cap A(x') = \emptyset$ for $x \neq x'$. To represent the rules for determining who moves at a decision node, we have a function $\iota : X \to \{i, j\}$ that assigns to each decision node the player whose turn it is.

Subtree. A subtree is defined by a subset of nodes $T \subset T^*$ for which $(T, \prec)$ is also an arborescence – perhaps starting at a different initial node. To ensure that the subtree is associated with well-defined payoffs to the players, we impose that all terminal nodes in the subtree are also in Z. Decision nodes that appear in the subtree, also appear in the game tree. We will need to explicitly differentiate these decision nodes and define them as distinct elements. To this effect, we label by y the copy of the decision node x, whenever the copy of x is part of the subtree T. Let Y be the set of copies of decision nodes in the subtree. The subtree together with the structure introduced below is intended to model the subjective partial view of the player.

Moves in the subtree. To ensure consistency between moves in the game tree and subtree, we impose that feasible actions at the copy y of decision node x are given by a non-empty subset $A(y) \subseteq A(x)$ for which the properties of the

function ψ are not violated. This implies that if the action $a \in A(x)$ leads from x to successor x' in the game tree T^*, then a also leads from the copy y to y' (the copy of x' in the subtree T) whenever both copies appear in the subtree. We also require that the same player moves at decision node x and copy y, so that there is no disagreement about which player that has to move.

To exemplify the implications of the above definitions, consider the the extensive-form underlying the sequential prisoners dilemma also analyzed by Dufwenberg and Kirchsteiger (2004):

Consider first the game tree T^* in Figure 1. At the initial node x^0, *Ann* is active and can corporate (the action C) or defect (the action D). At nodes x^1 and x^2, *Bob* is active and can corporate (the action c) or defect (the action d). Now consider the subtree T. The copy of decision node x^2 is not a part of the subtree, which implies that at the copy y^0 of the initial node x^0, *Ann* only has one feasible action C. *Bob*, on the other hand, can still choose c or d when he is active at copy y^1 of decision node x^1.

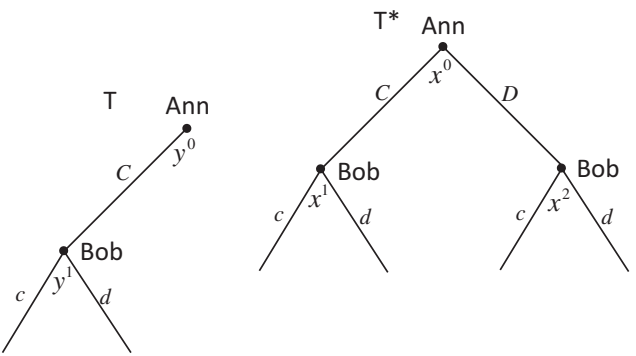

Figure 1: A game tree and a subtree thereof.

Generic decision nodes. A generic decision node n is an element of the disjoint union $N = X \cup Y$. By generic we mean that n can describe both decision nodes $x \in X$ as well as the copies $y \in Y$.

Singleton information sets. The information players have when choosing their actions is the most subtle part of our simple unawareness games. In the simple unawareness game players know all previous moves, but a player's frame of mind may not allow him to be aware of all paths of play. The information possessed by player i is represented by singleton information sets $h_i(n) \in H_i$ for all $n \in N$. The singleton information set $h_i(n)$ defines the frame of mind of player i by identifying the paths of play the player conceives possible at n. Unlike a standard information set, the generic decision node n does not need to be contained in $h_i(n)$. For example,

at decision node x in the game tree T^* it might be that $h_i(x)$ is in the subtree T. When the singleton information set $h_i(n)$ is in the subtree T, player i is unaware of all paths of play not described by the subtree.

In games with psychological payoffs, it is important to represent players' information also at nodes where they are not active. The information structure H_i of player i thus contains, as a subset $I_i \subset H_i$, the information structure of active player i. That is, the subset I_i contains as elements singleton information sets $h_i(n)$ of player $\iota(n) = i$.

Let T' and T'' be two generic trees that each can be either the subtree T or the game tree T^*. Each singleton information set $h_i(n) \in H_i$ has the following static properties that parallel those in Heifetz, Meier, and Schipper (2013a, 59–60) (see also Figure 2)[2]:

U0 Confined awareness: if $n \in T''$ then $h_i(n) \in T'$ with $T' \subseteq T''$.

U1 Generalized reflexivity: for $T \subset T^*$, $x \in T^*$, and $h_i(x) \in T$, if T contains a copy y of x then $y \in h_i(x)$.

U2 Introspection: if $n' \in h_i(n)$ then $h_i(n') = h_i(n)$.

U3 The subtree preserve awareness: for $T \subset T^*$, $x \in T^*$, and $x \in h_i(x)$, if T contains a copy y of x then $y \in h_i(y)$.

Properties U0-U3 are all static properties, as a dynamic property we impose perfect recall. We require that if generic decision node $n'' \in h_i(n')$ and if n is a predecessor of n', then there is a predecessor $\hat{n}$ of n'' in the same tree as n'' for which $h_i(\hat{n}) = h_i(n)$, and that the action taken at n along the path to n' is the same as the action taken at $\hat{n}$ along the path to n''. Intuitively, the generic decision nodes n' and n'' are distinguished by information player i does not have, so he cannot have had it when he was at information set $h_i(n)$; n' and n'' must be consistent with the same action at $h_i(n)$, since the player remembers his action there. Perfect recall together with confined awareness implies that the player cannot become unaware along a path of play. Suppose that $y' \in h_i(n')$, $h_i(n') \in T$ and generic decision node n is a predecessor of n'. For y' there is – by perfect recall – a copy $y \in T$ that is a predecessor of y', such that $h_i(y) = h_i(n)$. By confined awareness it must be that $h_i(n) \in T$. Thus, player i could not have been aware of more at n.

In games with psychological payoffs, where players somehow care about the beliefs or intentions of others, it is also important to be explicit about what a player is aware of and what he – given this awareness – thinks the other player

2 The numbers correspond to the respective property in Heifetz, Meier, and Schipper (2006, 83). Heifetz, Meier, and Schipper (2013a, 59–60) has two additional properties: "U4 Subtrees preserve ignorance" and "U5 Subtrees preserve knowledge", which in our simple setting are obsolete.

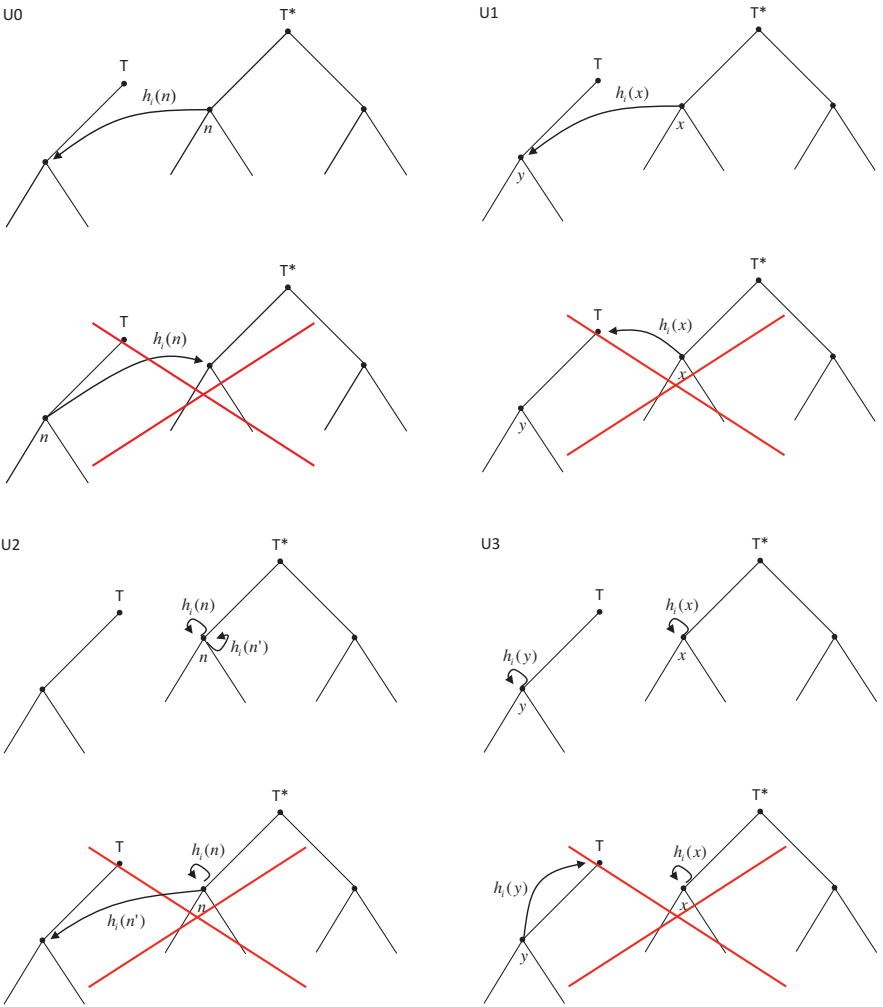

Figure 2: Examples that agree and disagree with properties U0-U3.

is aware of. To be explicit about this, we use composite singleton information sets $h_{ji}(n) = h_j \circ h_i(n)$.[3] Finally, each player's singleton information set is assumed a primitive of the game. When there is no need to be explicit it thus

3 Composite information sets can be more involved, for example, $h_i \circ h_j \circ h_i(n)$ which is player i's perception at node n of player j's perception of his frame of mind. However, in our setting with simple unawareness such involved iterative frames of mind will be redundant (due to property "U2 Introspection").

makes sense to write h_i instead of $h_i(n)$, use the notation $\iota(h_i)$ and $A(h_i)$ instead of $\iota(n)$ and $A(n)$, and write h_{ji} instead of $h_{ji}(n)$.

Figure 3 displays the extensive-form underlying the sequential prisoners dilemma considered in Figure 1 now with singleton information sets that are consistent with properties U0-U3 and perfect recall added. In the figure, singleton information sets are shown as arrows. The "solid arrows" indicate *Ann*'s singleton information sets, while the "broken arrows" indicate *Bob*'s singleton information sets. For the sake of simplicity, we omit *Bob*'s redundant singleton information sets at y^0 and y^1. *Ann* is aware when active. When choosing an action at singleton information set $h_A(x^0)$, she considers the physical paths of play. At $h_A(y^0)$ the action C should be interpreted as the action she would have taken had she only been aware of feasible path of play in the subtree. Moreover, she thinks that *Bob* is unaware since $h_{BA}(x^0) \in T$. When *Bob* is active at $h_B(x^1)$, he is unaware of the physical path where *Ann* could have chosen D. He thinks that *Ann* is also unaware since $h_{AB}(x^1) \in T$. If *Ann* chooses D, then *Bob* is aware and active at $h_B(x^2)$. At this singleton information set *Bob* will be surprised. He realizes that had he observed *Ann* choosing C, then he would not have suspected that she could have chosen anything other than that action.

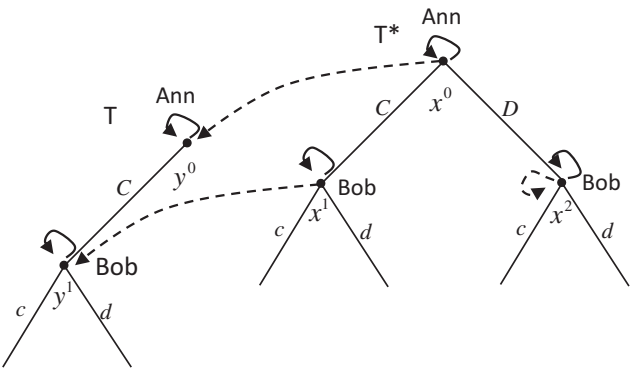

Figure 3: A game tree and a subtree thereof with singleton information sets.

Pure strategies. Let $A_i = \bigcup_{h_i \in H_i} A_i(h_i)$ be the set of all actions for player i. A pure strategy for aware player i is a map $s_i : I_i \to A_i$, with $s_i(h_i) \in A(h_i)$ for $h_i \in I_i$. A pure strategy for player i thus specifies an action at each of the singleton information sets at which the player is active. Player i's set of pure strategies, S_i, is simply the space of all such s_i. Since each of these pure strategies is a map

from singleton information sets to some action, we can write S_i as the Cartesian product of the action sets at each h_i:

$$S_i = \prod_{h_i \in I_i} A(h_i).$$

Because the aware player i is aware of all of his singleton information sets, his set of pure strategies is equal to S_i.

Remember, in our simple unawareness game only the game tree T^* represents the physical paths of play. The subtree T represents the restricted subjective view of the feasible paths in the mind of an unaware player, or the view of the feasible paths that an aware player assign to the unaware player, and so on. Moreover, as the game evolves a player may become aware of paths of which he was unaware earlier. A strategy can thus not in the simple unawareness game be conceived as an ex ante plan of actions. Instead, it should be interpret as a list of answers to the question "what would player i do if h_i were the singleton information set he considers as possible?"

For example, in Figure 3 we can identify aware *Ann*'s strategies at the singleton information set $h_A(x^0)$ with the actions C and D – the actions she actually takes. Aware *Ann*'s set of strategies is thus $S_A = \{CC, DC\}$, where the first index refers to the action taken at $h_A(x^0)$ and the second index to the action taken at $h_A(y^0)$. Following *Ann*'s action D, at singleton information set $h_B(x^2)$, *Bob* is aware, not only of the fact that he can take actions c and d, but also that he could have taken the actions c and d at $h_B(x^1)$ following *Ann*'s action C. We can thus identify *Bob*'s strategies by $S_B = \{cc, cd, dc, dd\}$.

Denote by $S = S_i \times S_j$ the set of strategy pairs. The path caused by a strategy pair $s \in S$ yields a terminal node denoted $\boldsymbol{z}(s) \in Z$. Strategy s_i reaches singleton information set h_i if the path induced by $s \in S$ reaches h_i. Otherwise, we say that h_i is excluded by the strategy s_i. The set of player i's strategies that reaches h_i is denoted $S_i(h_i)$. We will sometimes also write $S_j(h_i)$, meaning player j's strategies that reach h_i.

Unaware player i is unaware of singleton information sets in the game tree. For strategy $s_i \in S_i$, we denote by s_i^T the strategy of unaware player i induced by s_i. Strategy s_i induces strategy s_i^T if $s_i(h_i) = s_i^T(h_i)$ for every $h_i \in T$. For example, in Figure 3 the induced strategy of *Ann*'s strategy $s_i = (CC)$ is $s_A^T = (C)$, while the induced strategy of *Bob*'s strategy $s_B^T = (cd)$ is $s_B^T = (c)$. Let S_i^T be the set of player i's induced strategies, and $S_i^T(h_i)$ denote the set of strategies that reaches $h_i \in T$. The path caused by the induced strategy pairs yields a terminal node $\boldsymbol{z}(s^T) \in Z$. If $R_i \subseteq S_i$ is some set of strategies of player i, denote by R_i^T the set of strategies induced by R_i in the subtree T. In the rest of the paper we will use the notation

$s_i^{\lambda(h_i)} \in S_i^{\lambda(h_i)}$, with index function $\lambda(h_i) = \{\emptyset\}$ if $h_i \in T^\star$ and $\lambda(h_i) = \{T\}$ if $h_i \in T$, to define player i's strategies at a singleton information set h_i.

It is important to understand that *Bob* at $h_B(x^2)$ (after he has become aware) is not deluded to think that the strategic interaction at $h_B(x^1)$ is described by paths of play in the subtree, nor does he think that *Ann* was ever unaware. Rather, *Bob* interprets *Ann*'s induced strategy $s_A^T = \{C\}$ as describing the action *Ann* would have taken, had she decide to keep him unaware. For any solution concept we need to analyse what, for example, *Ann* thinks of *Bob*'s actions in the subtree, such actions are determined by *Bob*'s induced strategy. This is why *Ann*'s strategy also determines her action at $h_A(y^0)$ in the subtree, even though she will never actually only consider this action as possible given that she is aware of all feasible paths of play.

Material payoffs. Above we defined the strategic interaction in an extensive form with simple unawareness. To obtain a dynamic game with simple unawareness we add a specification of the players' material payoffs assigned to the terminal nodes. Because each terminal node $z \in Z$ completely determines a path through the game tree, we can assign to player i material payoffs using functions $\pi_i : Z \to \mathbb{R}$.

3 Beliefs and Psychological Payoffs

Extensive games with simple unawareness, as described above, assumed that payoffs depend only on induced paths of play. This in not sufficient for describing the motivations and choices of players who care about, for example, guilt aversion and reciprocity. Psychological games, on the other hand, allow payoffs to depend directly on beliefs (about beliefs), and via such beliefs capture, for example, emotions like reciprocity and guilt.

Beliefs. Conditional on each singleton information set $h_i \in H_i$, player i holds an updated, or revised, belief $\alpha_i(\cdot|h_i) \in \Delta(S_j^{\lambda(h_i)}(h_i))$;

$$\alpha_i = (\alpha_i(\cdot|h_i))_{h_i \in H_i} \in \prod_{h_i \in H_i} \Delta\left(S_j^{\lambda(h_i)}(h_i)\right)$$

is the system of first-order beliefs of player i.

For example, at $h_B(x^0)$ in Figure 3, *Bob* is certain that *Ann*'s strategy is $s_A^T = \{C\}$. If *Bob* subsequently finds himself at $h_B(x^2)$, then his belief will change, so that he now is certain that *Ann*'s strategy is $s_A = \{CD\}$. *Bob* thus becomes aware that *Ann* could have kept him unaware by choosing strategy $s_A = \{CC\}$, but instead chose to make him aware by choosing action D.

At h_i player i also holds a second-order belief $\beta_i(h_{ji})$ about the first-order belief system α_j of player j, a third-order belief about the second-order beliefs, and so on. For the purpose of this paper, we may assume that higher-order beliefs are degenerate point beliefs. Thus, with a slight abuse of notation we identify $\beta_i(h_{ji})$ with a particular first-order belief system α_j. A similar notational convention applies to other higher-order beliefs. Let the sequence $\mu_i = (\alpha_i, \beta_i, \ldots)$ denote player i's hierarchy of beliefs, and M_i the (compact) set of such hierarchies. In our applications we consider beliefs at most of the second order.

At $h_A(x^0)$ in Figure 3, *Ann* holds a first-order belief $\alpha_A(\cdot|h_A(x^0)) \in \Delta(\{cc, cd, dc, dd\})$. She thinks that *Bob* is unaware, and her belief about *Bob*'s first-order belief about her strategy must reflect this. *Ann* thus needs to consider *Bob*'s frame of mind. That is why her second-order belief $\beta_A(h_{BA}(x^0))$ is conditioned on the composite singleton information set, implying that she is certain that *Bob* is certain that her strategy is $s_A^T = \{C\}$.

Players should not change their beliefs unless the play reaches a singleton information set which falsifies it. We therefore assume that player i's hierarchy of beliefs M_i are consistent such that: there is at least one strategy of player j in the support of $\alpha_i(\cdot|h_i)$ at some h_i, and that beliefs must satisfy Bayes' rule and common knowledge of Bayes' rule whenever possible. Consistency of the updating system requires that $\alpha_i(\cdot|h_i)$, $\beta_i(h_{ji})$, and so on, at h_i are consistent with h_i being reached and that no beliefs are abandoned unless falsified. Thus, when *Bob* in Figure 3 finds himself at $h_B(x^2)$, he must change his beliefs such that they are consistent with being aware.

Psychological payoffs. Section 2 defines dynamic games with simple unawareness. To obtain a dynamic psychological game with simple unawareness we extent payoffs to include beliefs. Specifically, we assign to player i psychological payoffs using functions $u_i : Z \times M_i \to \mathbb{R}$. The psychological payoffs will be obtained from the material payoff functions $\pi_i : Z \to \mathbb{R}$. We exemplify this definition by considering two prominent functional forms capturing simple guilt aversion and reciprocity, respectively.

Simple guilt aversion. Simple guilt aversion as in Battigalli and Dufwenberg (2007) implies that player i judges the initial expectations of player j concerning his material payoff and feels guilty whenever he does not live up to these expectations. Given his strategy $s_j^{\lambda(h_j(x^0))}$ and the first-order belief system α_j, player j forms an initial expectation about his material payoffs π_j:

$$\mathbb{E}_{s_j^{\lambda(h_j(x^0))}, \alpha_j}[\pi_j|h_j(x^0)] = \sum_{s_i^{\lambda(h_j(x^0))} \in S_i^{\lambda(h_j(x^0))}} \alpha_j\left(s_i^{\lambda(h_j(x^0))}|h_j(x^0)\right)\pi_j\left(z\left(s_j^{\lambda(h_j(x^0))}, s_i^{\lambda(h_j(x^0))}\right)\right).$$

For a given terminal node z the function

$$D_j\left(z, s_j^{\lambda(h_j(x^0))}, \alpha_j\right) = \max\left\{0, \mathbb{E}_{s_j^{\lambda(h_j(x^0))}, \alpha_j}[\pi_j|h_j(x^0)] - \pi_j(z)\right\}$$

measures how much player j is "let down". Player i does not know player j's strategy and first-order beliefs, but holds a belief about these. Denote player i's belief about player j's "let down" by $D_{ij}(z, s_j^{\lambda(h_{ji}(x^0))}, \beta_i)$, where β_i is player i's second-order belief system. Given this, player i is motivated by simply guilt aversion if his psychological payoffs are represented by:

$$u_i(z, \mu_i) = \pi_i(z) - \theta_{ij} D_{ij}\left(z, s_j^{\lambda(h_{ji}(x^0))}, \beta_i\right) \tag{1}$$

where θ_{ij} reflect player i's sensitivity to guilt. Guilt averse player i senses a psychological cost connected to his feeling of guilt in case he does not live up to his belief about player j's initial expectation and takes this into account when deciding on his optimal behavior. Players' feelings of guilt thus depend on their awareness, and their beliefs about the awareness of the other player.

Reciprocity. Different from guilt aversion, reciprocity as in Dufwenberg and Kirchsteiger (2004) assumes that players judge the kindness of others. Whenever player i judges player j to be kind, he reciprocates by being kind himself. Whenever player i judges player j to be unkind, he acts unkindly in return. More formally, given a singleton information set h_j, a strategy $s_j^{\lambda(h_j)}$ that reaches h_j and the first-order belief system α_j, player j forms an expectation about player i's material payoff π_i:

$$\mathbb{E}_{s_j^{\lambda(h_j)}, \alpha_j}[\pi_i|h_j] = \sum_{s_i^{\lambda(h_j)} \in S_i^{\lambda(h_j)}} \alpha_j\left(s_i^{\lambda(h_j)}|h_j\right)\pi_i\left(\boldsymbol{z}\left(s_i^{\lambda(h_j)}, s_j^{\lambda(h_j)}\right)\right)$$

Player j's kindness towards player i is described by how much player j expects to give player i relative to some equitable payoff $\pi_i^e(h_j)$ at singleton information set h_j:

$$K_j(h_j) = \mathbb{E}_{s_j^{\lambda(h_j)}, \alpha_j}[\pi_i|h_j] - \pi_i^e(h_j).$$

The equitable payoff is the threshold or neutral payoff above (below) which player j treats player i kindly (unkindly). In other words, if $K_j(\cdot) > 0$, then player j treats player i kindly. Conversely, if $K_j(\cdot) < 0$, then player j treats player i unkindly. Let the equitable payoff be

$$\pi_i^e(h_j) = \frac{1}{2} \times \left[\max_{s_i^{\lambda(h_j)} \in S_i^{\lambda(h_j)}} \mathbb{E}_{s_j^{\lambda(h_j)}, \alpha_j}[\pi_i|h_j] + \min_{s_i^{\lambda(h_j)} \in S_i^{\lambda(h_j)}} \mathbb{E}_{s_j^{\lambda(h_j)}, \alpha_j}[\pi_i|h_j]\right].$$

The equitable payoff is the average player j is able to give to player i in material terms based on his awareness and first-order belief. As in the case of guilt aversion, player i does not know strategy s_j and first-order beliefs α_j, but holds a first- and second-order belief about them. Denote player i's judgement of player j's kindness at information set h_i by $K_{ji}(h_{ji}) = \mathbb{E}_{s_j^{\lambda(h_{ji})},\beta_i}[\pi_i|h_{ji}] - \pi_i^e(h_{ji})$, where β_i is player i's second-order belief system. We say player i is motivated by reciprocity if he has belief-dependent preferences represented by a psychological payoff function of the form:

$$u_i(z, \mu_i|h_{ji}) = \pi_i(z) + Y_i \times K_{ji}(h_{ji}) \times \pi_j(z), \tag{2}$$

where $Y_i > 0$ is player i's sensitivity to reciprocity. Whenever player i perceives player j to be kind, player i is motivated to also maximize player j's material payoff. In case player i judges player j to be unkind, player i is motivated to reduce player j's material payoff. This definition of reciprocity implies that if player i is unaware of paths of play, he judges the kindness of player j based on the paths that he is aware of.

4 Extensive-Form Rationalizability

We adapt to the present framework the extensive-form rationalizability concept of Pearce (1984). To do so, we first extend the notion of sequential rationality and then redefine the best-rationalization principle as outlined by Battigalli (1997) and Battigalli and Siniscalchi (2002).

Sequential rationality. Our basic behavioral assumption is that player i chooses and carries out a strategy s_i that reaches singleton information set h_i and is optimal given his hierarchy of beliefs μ_i, conditional upon any singleton information set consistent with s_i. It is thus not required that a strategy specifies behavior at singleton information sets that cannot be reached by s_i.

Fix a singleton information set h_i. Player i's expectation of the psychological payoff u_i, given s_i and μ_i is

$$\mathbb{E}_{s_i,\mu_i}[u_i|h_i] = \sum_{s_j^{\lambda(h_i)} \in S_j^{\lambda(h_i)}} \alpha_i\left(s_j^{\lambda(h_i)}|h_i\right) \times u_i\left(\boldsymbol{z}\left(s_i^{\lambda(h_i)}, s_j^{\lambda(h_i)}\right), \mu_i\right).$$

Given a hierarchy of beliefs $\mu_i \in M_i$, strategy s_i is a sequential best response to μ_i if for all $h_i \in I_i$:

$$s_i \in \arg\max_{s_i \in S_i(h_i)} \mathbb{E}_{s_i,\mu_i}[u_i|h_i].$$

For any hierarchy of beliefs μ_i, let $BR_i(\mu_i)$ denote the set of strategies s_i that are sequential best responses to μ_i. The set of best responses thus consists of strategies s_i of player i that, for a given μ_i, are undominated at every singleton information set $h_i \in H_i$, given his awareness at h_i.

Clearly, BR_i is nonempty valued. The first-order belief $\alpha_i(\cdot|h_i)$ is continuous. Since u_i is also continuous, we have that $\alpha_i(\cdot|h_i) \times u_i(\cdot,\cdot)$ is continuous, which implies that BR_i is an upper hemicontinuous correspondence.

To see how the extension of sequential rationality works, consider *Bob* in Figure 3. *Bob*'s strategy $s_B = \{cd\}$, for example, is sequential rational if it is undominated by his other strategies $\hat{s}_B \in \{cc, dc, dd\}$ at every $h_B \in H_B$. *Bob* is unaware when he is active at singleton information set $h_B(x^1)$. Here his strategy $s_B = \{cd\}$ is undominated if his induced strategy $s_B^T = \{c\}$ gives him an expected psychological payoff that is equal or higher than his induced strategy $\hat{s}_B^T = \{d\}$, given *Ann*'s induced strategy $s_A^T = \{C\}$ that reach $h_B(x^1)$. If *Bob* is active at $h_B(x^2)$, then he has become aware. Here his strategy $s_B = \{cd\}$ is undominated if it gives him an expected psychological payoff that is equal or higher than his strategies $\hat{s}_B \in \{cc, dc, dd\}$, given *Ann*'s strategy $s_A = \{CD\}$ that reach $h_B(x^2)$. Thus, *Bob*'s strategy $s_B = \{cd\}$ is sequential rational if, and only if, it is undominated at both $h_B(x^1)$ and $h_B(x^2)$.

Best-rationalization principle. Our redefinition of the best-rationalization principle requires that players' beliefs conditional upon observing singleton information set h_i be consistent with the highest degree of strategic sophistication of the other player. In the following we clarify what we mean by strategic sophistication in terms of dynamic psychological games with unawareness. Consider the following extensive-form rationalization procedure for player i:

$$M_i[1] = M_i,$$

$$R_i[1] = \left\{ \begin{array}{ll} s_i \in S_i & \text{such that there exists a hierarchy of beliefs } \mu_i \in M_i[1] \\ & \text{for which } s_i \in BR(\mu_i) \end{array} \right\},$$

$$\vdots$$

$$M_i[k] = \left\{ \begin{array}{ll} \mu_i \in M_i[k-1] & \text{such that for all singleton information sets } h_i \in H_i, \text{if} \\ & R_j[k-1](h_i) \neq \emptyset, \text{then } \alpha_i(R_j^{\lambda(h_i)}[k-1]|h_i) = 1 \end{array} \right\},$$

$$R_i[k] = \left\{ \begin{array}{ll} s_i \in S_i & \text{such that there exists a hierarchy of beliefs } \mu_i \in M_i[k] \\ & \text{for which } s_i \in BR(\mu_i) \end{array} \right\}.$$

Let $R_i[\infty] = \bigcap_{k \geq 0} R_i[k]$. Player i's strategies in $R_i[\infty]$ are said to be extensive-form (correlated) rationalizable in a dynamic psychological game with simple unawareness.

Intuitively our definition of extensive-form rationalizability starts at the level of strategic thinking of player i, whose step 1 rationalizable strategies are sequential best responses to some nonrestricted hierarchy of beliefs. Strategies that are not sequential best responses to any nonrestricted hierarchy of beliefs are not step 1 rationalizable.

Next, player i restricts his hierarchies of beliefs to those for which he, at each singleton information set h_i, is certain (given his awareness) of those step 1 rationalizable strategies of player j that reach h_i. Player i then chooses step 2 rationalizable strategies that are sequential best responses to these restricted hierarchies of beliefs. Strategies that are not sequential best responses to these restricted hierarchies of beliefs, on player j's step 1 rationalizable strategies, are not step 2 rationalizable.

Furthermore, player i must restrict his restricted hierarchies of beliefs to those for which he, at each singleton information set h_i, is certain (given his awareness) of those step 2 rationalizable strategies of player j that reach h_i. Player i then chooses step 3 rationalizable strategies that are sequential best responses to these twice restricted hierarchies of beliefs. Strategies that are not sequential best responses to these twice restricted hierarchies of beliefs, on player j's step 2 rationalizable strategies, are not step 3 rationalizable, and so on.

The sequence $\{R_i[k] : k \geq 0\}$ is a weakly decreasing, that is, $R_i[k+1] \subseteq R_i[k]$ for all k. Since R_i is finite the sequence converges in countably many steps. The limit is given by the first integer K such that $R_i[K] = R_i[K+1]$. Consider the limit set $R_i[K]$ of player i. The sequence $R_j[0], R_j[1], \ldots, R_j[K-1]$ represents a hierarchy of increasingly strong hypotheses of player i about the behavior of player j. When player i implements a strategy $s_i \in R_i[K]$, he always optimize accordingly. At the beginning of the game, it is the common knowledge that all players update and behave in this way.

The set M_k (for $k > 1$) implies that along each feasible path of play, at a singleton information set an active player is certain that the other player sequential best responds, certain that the other player is certain he sequential best responds, and so on. If a player finds himself at a succeeding singleton information set, where the other player's strategies that could lead to that singleton information set are inconsistent with the player's previous certainty in the other player's best response, then the player seeks a best rationalization that could have led to that singleton information set. That is, if the player is "surprised" by the other player's unexpected action, and cannot use Bayesian updating, then he forms new beliefs that justify this observed inconsistency. In its simplest form, forward-induction reasoning involves the assumption that, upon observing an unexpected (but undominated) action of the other player,

a player maintains the working hypothesis that the latter is a sequential rational. The best-rationalization principle captures precisely this type of argument.

Forward-induction reasoning thus implies that at $h_B(x^0)$ and onwards, unaware *Bob* is certain that *Ann*'s sequential best response is $s_A^T = \{C\}$. However, if singleton information set $h_B(x^2)$ is reached and *Bob* becomes aware, then he is certain given his newly found awareness that *Ann*'s action $s_A = \{CD\}$ is a sequential best response to some hierarchy of beliefs of hers. At $h_B(x^2)$, *Bob* has no choice but to revert to being certain that *Ann* would not choose the strategy $s_A = \{CC\}$ rationally, and excludes *Ann*'s hierarchies of beliefs for which $s_A = \{CC\}$ is a sequential best response.

5 Two Examples

In the following we present two examples to highlight the impact and importance of simple unawareness in dynamic psychological games. In particular, our examples demonstrate that the strategic behavior of players motivated by psychological payoffs can be effected by unawareness even in situations where it would not have mattered had the player been selfish. In these situations the aware player often has a subtle incentive to manage the awareness of the unaware player by disclosing paths of play if possible. In the first example, we analyze a sequential prisoners dilemma featuring unawareness and reciprocity. Second, we re-visit our introductory Ann/Bob-example featuring guilt aversion. A full description of the strategic interaction with all possible awareness levels is beyond the scope of this paper. Therefore, we limit the analysis to specific awareness scenarios.

Example 1 (reciprocity). Consider the the extensive-form underlying the sequential prisonors dilemma with unawareness depicted in Figure 3 now with *Ann*'s and *Bob*'s material payoffs added.

Remember, at the initial node x^0, *Ann* is active and can corporate (the action *C*) or defect (the action *D*). At nodes x^1 and x^2, *Bob* is active and can corporate (the action *c*) or defect (the action *d*). In the strategic setting depicted in Figure 4, *Ann* is initially aware of everything, whereas *Bob* is initially unaware. *Ann* knows this, and knows that he only becomes aware of everything if she chooses *D*. Also, *Ann* knows that *Bob* perceives her to be unaware. The "solid arrows" indicate *Ann*'s singleton information sets, while the "broken arrows" indicate *Bob*'s singleton information sets. For the sake of simplicity, we omit *Bob*'s redundant singleton information sets at y^0 and y^1.

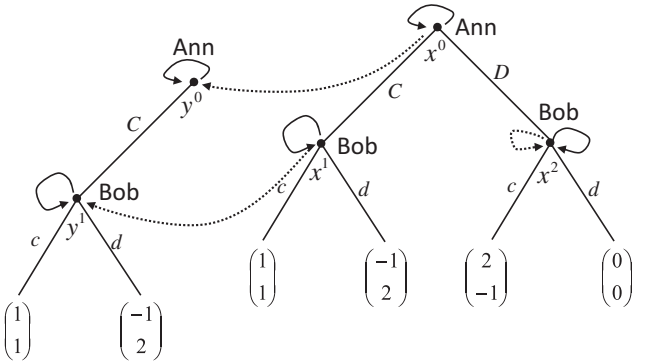

Figure 4: Sequential prisoners dilemma game with unawareness.

We assume that *Bob* is motivated by reciprocity and *Ann* is selfish. *Bob*'s psychological payoff is thus described by eq. [2], while *Ann* only cares about her material payoff $\pi_A(z)$. *Ann*'s optimal strategy depends on her first-order belief about *Bob* strategy (conditional on her behavior). Her strategy $s_A = \{CC\}$ is a sequential best response as long as her expected material payoff from strategy $s_A = \{CC\}$ exceeds her expected material payoff from strategy $s_A = \{CD\}$. Her expected material payoff from strategy $s_A = \{CC\}$ is:

$$\mathbb{E}_{\{CC\},\alpha_A}[\pi_A|h_A(x^0)] = \alpha_A(\{c\cdot\}|h_A(x^0)) \cdot 1 + (1 - \alpha_A(\{c\cdot\}|h_A(x^0))) \cdot -1,$$

where $\alpha_A(\{c\cdot\}|h_A(x^0)) := \sum_{s_B \in \{cc,\, cd\}} \alpha_A(s_B|h_A(x^0))$ is a shorthand notation for *Ann*'s first-order belief about the strategies of *Bob* that select the action *c* at the singelton information set $h_B(x^1)$. *Ann*'s expected material payoff from following strategy $s_A = \{CD\}$ is:

$$\mathbb{E}_{\{CD\},\alpha_A}[\pi_A|h_A(x^0)] = \alpha_A(\{\cdot c\}|h_A(x^0)) \cdot 2,$$

where $\alpha_A(\{\cdot c\}|h_A(x^0)) := \sum_{s_B \in \{cc,\, dc\}} \alpha_A(s_B|h_A(x^0))$ is an akin shorthand notation. The step 1 rationalizable strategies of *Ann* are thus

$$R_A[1] = \{s_A : \alpha_A(\{c\cdot\}|h_A(x^0)) - \alpha_A(\{\cdot c\}|h_A(x^0)) \geq \frac{1}{2} \Rightarrow s_A = \{CC\},$$

$$\text{otherwise} \Rightarrow s_A = \{CD\}\}.$$

Bob, on the other hand, is initially passive and only becomes active at singleton information sets $h_B(x^1)$ and $h_B(x^2)$. Remember that at singleton information set $h_B(x^1)$, *Bob* is unaware and holds a first-order belief $\alpha_B(s_A^T|h_B(x^1))$ and a second-order point belief $\beta_B(h_B(x^1))$ concerning *Ann*'s first-order belief. Given his

unawareness at $h_B(x^1)$, *Bob* thinks that *Ann*'s only action is *C*. Independent of his second-order belief and his sensitivity to reciprocity Y_B he thus judges *Ann* as neither kind nor unkind (because $K_{BAB}(h_B(x^1)) = 0$). Consequently, whenever *Bob* finds himself at $h_B(x^1)$ he will simply maximize his own material payoff by choosing *d*.

At singleton information set $h_B(x^2)$, *Bob* is aware of everything and knows that he could have earned a material payoff of 2 had *Ann* initially chosen to keep him unaware by choosing *C*. *Bob*'s judgment of *Ann*'s intention towards him at $h_B(x^2)$ is thus

$$\mathbb{E}_{\{CD\},\beta_B}[\pi_B|h_B(x^2)] = \beta_B(\{\cdot c\}|h_B(x^2)) \cdot -1,$$

where $\beta_B(\{\cdot c\}|h_B(x^2)) := \sum_{s_B \in \{cc,\, dc\}} \beta_B(h_B(x^2))$ is a shorthand notation for *Bob*'s second-order belief about the likelihood with which *Ann* believes he chooses *c* at information set $h_B(x^2)$. Clearly, $2 > \mathbb{E}_{\{CD\},\beta_B}[\pi_B|h_B(x^2)]$ independent of second-order belief $\beta_B(h_B(x^2))$. Hence, *Bob* judges *Ann* as unkind when finding himself at $h_B(x^2)$. The step 1 rationalizable strategies of *Bob* are thus

$$R_B[1] = \{s_B : \text{for all } \beta_B, Y_B \Rightarrow s_B = \{dd\}\}.$$

Although *Bob* is motivated by reciprocity as defined in eq. [2], his behavior in our sequential prisoners dilemma with unawareness is independent of his (second-order) beliefs and independent of his sensitivity to reciprocity Y_B.

Bob's set of step 1 rationalizable strategies is a singleton set $R_B[1] = \{dd\}$. *Ann* is thus at all her singelton information sets certain that *Bob* follows strategy $s_B = \{dd\}$. That is, *Ann*'s hierarchy of beliefs $\mu_A \in M_A[2]$ are all such that $\alpha_A(\{dd\}|\cdot) = 1$. Being certain that *Bob* chooses *d* no matter what, *Ann*'s sequential best response strategy must also select *D* as an action (since $0 > -1$). The step 2 rationalizable strategies of *Ann* are thus

$$R_A[2] = \{s_A : \text{for all } \mu_A \in M_A[2] \Rightarrow s_A = \{CD\}\}.$$

Ann anticipates that given his awareness, *Bob* judges her as unkind and chooses *d* independent of what she does. Consequently, since *Ann* is only interested in her own material payoff, she chooses *D* herself to get a material payoff of 0 instead of -1.

To study the impact of unawareness, we compare this scenario to the rationalizable solution of the sequential prisoners dilemma with reciprocity and full awareness (see Figure 5).

Now *Ann* chooses to *C* in the initial singleton information set $h_A(x^0)$ as long as she believes sufficiently strongly that *Bob* will choose *c*. Her expected payoff

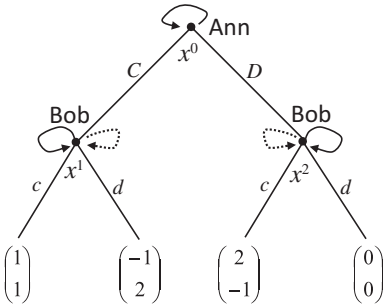

Figure 5: Sequential prisoners dilemma with full awareness.

from choosing either $s_A = \{CC\}$ or $s_A = \{CD\}$ at $h_A(x^0)$ is the same as before, and her step 1 rationalizable strategies are again:

$$R_A[1] = \{s_A : \alpha_A(\{c\cdot\}|h_A(x^0)) - \alpha_A(\{\cdot c\}|h_A(x^0)) \geq \frac{1}{2} \Rightarrow s_A = \{CC\},$$
$$\text{otherwise} \Rightarrow s_A = \{CD\}\}.$$

Bob's optimal behavior at $h_B(x^2)$ remains the same as before. That is, *Bob* chooses d out of material and reciprocal reasons. However, *Bob*'s optimal behavior at $h_B(x^1)$ now depends on his sensitivity to reciprocity Y_B. Let *Bob*'s (second-order) belief about *Ann*'s belief concerning the likelihood with which he chooses c at $h_B(x^1)$, be $\beta_B(\{c\cdot\}|h_B(x^1)) := \sum_{s_B \in \{cc,\, cd\}} \beta_B(h_B(x^1))$. The step 1 rationalizable strategies of *Bob* are thus:

$$R_B[1] = \{s_B : \beta_B(\{c\cdot\}|h_B(x^1)) \leq 2 - \frac{1}{Y_B} \Rightarrow s_B = \{cd\}, \quad (3)$$
$$\text{otherwise} \Rightarrow s_B = \{dd\}\}.$$

The lower *Bob*'s second-order belief is, the kinder he perceives *Ann*'s strategy $s_A = \{CC\}$ (because $K_{BAB} = 1 - \frac{1}{2} \cdot \beta_B(\{c\cdot\}|h_B(x^1))$), which provides him with a payoff which is higher than if she had chosen strategy $s_A = \{CD\}$. At $h_B(x^1)$, *Bob* never actually thinks *Ann* is unkind. The question at this augmented history simply is whether he thinks she is kind enough, given his sensitivity to reciprocity Y_B, such that he prefers to reciprocate her kindness.

If *Bob*'s sensitivity to reciprocity is low ($Y_B \leq \frac{1}{2}$), such that he for sure chooses d if *Ann* chooses C, then *Ann* will choose strategy $s_A = \{CD\}$ as this provides her with a higher expected material payoff. Conversely, if *Bob*'s sensitivity to reciprocity is high ($Y_B \geq 1$), *Ann* is certain that *Bob* chooses c if she chooses C. Given this, she chooses strategy $s_A = \{CC\}$ as this provides her with a higher expected material payoff. Notice, if *Bob* is sensitive enough to *Ann*'s kindness, then he

chooses c at $h_B(x^1)$ independent of his second-order belief. Given this *Ann* also chooses C, something she would not do were she sure that *Bob* would be unaware. Based on $R_B[1]$, *Ann*'s beliefs are

$$M_A[2] = \{\mu_A : \text{for all } Y_B \geq 1 \Rightarrow \alpha_A(\{c\cdot\}|h_A(x^0)) = 1, \text{ for all } Y_B < \frac{1}{2} \Rightarrow \alpha_A(\{c\cdot\}|h_A(x^0)) = 0\}.$$

At step 2 reasoning, *Ann* is certain that a very sensitive *Bob* will always choose c if she also chooses C, whereas a very insensitive *Bob* will always choose d no matter what. For intermediate sensitivity levels, $\frac{1}{2} \leq Y_B < 1$, her beliefs are $M_A[2] = M_A[1]$. *Ann*'s step 2 rationalizable strategies are thus:

$$R_A[2] = \{s_A : \text{for } Y_B \geq 1 \Rightarrow s_A = \{CC\}, \text{ for } \frac{1}{2} \leq Y_B < 1, \gamma_A(\{c\cdot\}|\boldsymbol{h}_A(n^0)) \leq 2 - \frac{1}{Y_B} \Rightarrow s_A = \{CC\}, \text{ otherwise } \Rightarrow s_A = \{CD\}\}.$$

where γ_A denotes *Ann*'s (third-order) point belief about *Bob*'s second-order belief. Based on $R_A[2]$, *Bob*'s step 3 beliefs are

$$M_B[3] = \{\mu_B : \text{for } Y_B \geq 1 \Rightarrow \beta_B(\{c\cdot\}|h_b(x^1)) = 1, \text{ for } Y_B < \frac{1}{2} \Rightarrow \beta_B(\{c\cdot\}|h_B(x^1)) = 0, \text{ otherwise } \Rightarrow \delta_B(\{c\cdot\}|h_B(x^1)) \leq 2 - \frac{1}{Y_B}, \beta_B(\{c\cdot\}|h_B(x^1)) = 1\}.$$

where δ_B denotes *Bob*'s (fourth-order) belief about *Ann*'s third-order beliefs. $M_B[3]$ implies that *Bob* believes that *Ann* expects him to choose $s_B = \{cd\}$ whenever he finds himself at singleton information set $h_B(x^1)$. *Bob*'s step 3 rationalizable strategies are thus:

$$R_B[3] = \{s_B : \text{for } Y_B \geq 1 \Rightarrow s_B = \{cd\}, \text{ otherwise } Y_B < 1 \Rightarrow s_B = \{dd\}\}.$$

If *Bob* is sensitive enough to reciprocity ($Y_B \geq 1$), he always chooses c if *Ann* chooses C, and chooses d if *Ann* chooses D. If he is insensitive enough to reciprocity ($Y_B < \frac{1}{2}$), then he always chooses d no matter what *Ann* does. However, if his sensitivity to reciprocity is $\frac{1}{2} \leq Y_B < 1$ and he finds himself at $h_B(x^1)$, then he is certain that *Ann* believes that he will choose c (because $\beta_B(\{c\cdot\}|h_B(x^1)) = 1$). That is, $1 > 2 - \frac{1}{Y_B}$ and he defects although he believes that *Ann* believes that he will choose c. Finally, *Ann*'s step 4 rationalizable strategies are:

$$R_A[4] = \{s_A : \text{for all } Y_B \geq 1 \Rightarrow s_A = \{CC\}, \text{ for all } Y_B < 1 \Rightarrow s_B = \{CD\}\}.$$

With full awareness, if *Bob* is sufficiently sensitive to reciprocity ($Y_B \geq 1$), then *Ann* chooses *C* since this induces *Bob* to choose *c*. This stands in contrast to the result of the first awareness scenario in which *Bob* was unaware of *Ann*'s possibility to defect even after her choice *C*. This example highlights that, although reciprocity is only based on first- and second-order beliefs, the recursive nature of extensive-form rationalizability requires the specification of higher (potentially infinite) orders of beliefs.

In synthesis, although *Bob*'s sensitivity to reciprocity might be very high, his behavior in the first scenario stands in contrast to the result in the second senario. With full awareness *Bob*'s behavior following *Ann*'s action *C* depends on *Bob*'s sensitivity to reciprocity. For sufficiently high levels of sensitivity *Bob* reciprocates by choosing *c*. With unawareness *Bob*'s behavior is independent of his sensitivity to reciprocity. *Bob* simply defects as he perceives *Ann*'s action as unkind no matter what she does. As a consequence, also *Ann*'s behavior is qualified – she defects as well. Interestingly, *Bob* behaves *as if* he is selfish, although he is not. It is only his subjective perception concerning the strategic environment which drives his behavior.

It is at the intersection of these two scenarios that the implications of unawareness for the behavior of people motivated by psychological payoffs become most visible. If *Bob* were only interested in his own material payoff, his behavior in the two scenarios would be the same. Most importantly, being only interested in his material payoff *Bob* would choose *d* following *Ann*'s decision to choose *C* independent of whether there are different awareness levels (as in the first scenario) or not (as in the second scenario). It is only his psychological payoff which explains the above-described difference in behavior. *Ann* thus has a strong incentive to make *Bob* aware and choose *C*. As long as *Bob* is unaware that *Ann* could have defected, he will not perceive her corporation as kind. However, were *Ann* able to make *Bob* aware, then she (as well as *Bob*) would have been better off. This shows, similar to our example in the introduction, how the opportunity to disclose paths of play is important when players are motivated by psychological payoffs.

Example 2 (guilt aversion): Consider again the example from the introduction featuring *Bob*'s birthday party and *Ann*'s final exam at university. Assume *Bob* would like to organize a party and would be happy if *Ann* could come. *Bob* gets a material payoff of 2 from not organizing the party (denoted *NP*), he gets a material payoff of 3 from organizing the party (denoted *P*) if *Ann* can join in, and a material payoff of 1 if she cannot make it. Thus, he would rather not organize

the party if *Ann* cannot make it. Furthermore, assume that *Bob* is unaware of the second date, and hence, he is certain that she cannot make it.

Suppose that *Ann* knows what *Bob* is aware of, thus she knows that he is certain that she has an exam the day after his birthday party and cannot come, but in addition she also knows that she could actually write the final exam 2 months later at a second date. Denote her decision to write the exam on the first date (the day after the party) by *No* 1 and the second date by *No* 2. She can only attend his party if she decides to write the exam on date *No* 2, if either *Bob* does not organize a party (*NP*) or *Ann* chooses *No* 1, she writes the exam on the first date. Assume *Ann*'s material payoff from not going to the party and writing the exam on the second date is 4, whereas her material payoff from going to the party and writing the exam on the second date is 2. Figure 6 depicts this simple unawareness game. Assume *Ann* is motivated by simple guilt aversion as described by eq. [1] with $\theta > 2$, and that *Bob* is only interested in his own material payoff. The "solid arrows" now indicate *Bob*'s singleton information sets, while the "broken arrows" indicate *Ann*'s singleton information sets. Again we omit redundant singleton information sets.

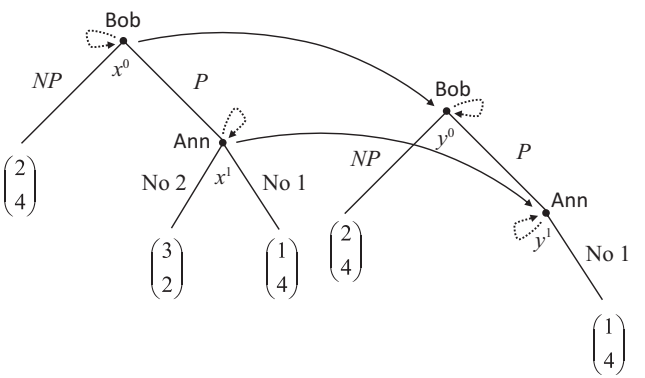

Figure 6: Example 2 with guilt aversion.

Before looking at this unawareness game consider – as a benchmark – a situation of full awareness. With full awareness the rationalizable solution is that *Bob* chooses *P* and *Ann* chooses *No 2*. The reason is the following: *Bob* only chooses *P* if it implies an expected material payoff which is higher than 2. That is,

$$R_B[1] = \{s_B : \alpha_B(No\ 2 | h_B(x^0)) \geq \frac{1}{2} \Rightarrow s_B = \{P\},$$
$$\text{otherwise} \Rightarrow s_B = \{NP\}\}.$$

If *Ann* finds herself at singleton information set $h_A(x^1)$ she has to believe that *Bob*'s belief regarding the likelihood that she chooses *No* 2 is $\alpha_B(No\ 2|h_B(x^1)) \geq \frac{1}{2}$. That is, *Ann*'s second-order belief is $\beta_A(No\ 2|h_A(x^1)) \geq \frac{1}{2}$. *Ann* will choose to postpone the exam and attend the party (*No* 2) as long as this gives her a higher expected psychological payoff. Since $\theta > 2$ it holds that:

$$R_A[1] = \{s_A : \beta_A(No\ 2|h_A(x^1)) \geq \frac{1}{2} \Rightarrow s_A = \{No\ 2\}\}.$$

Since *Bob* can be certain that *Ann* will choose *No* 2 if he chooses *P* this is the only rationalizable outcome.

Consider now the case with unawareness on *Bob*'s side. *Bob* would only choose to organize a party if he sufficiently believed that *Ann* could come. Because *Bob* is unaware of the second exam date, he is certain that he will get a material payoff of 1 if he organizes the party and 2 if he does not. *Bob*'s only rationalizable strategy is to not organize the party (that is, choose *NP*). *Ann* knows that unaware *Bob* is not "let down" if she chooses *No* 1 at $h_A(x^1)$ following *Bob*'s decision to organize the party, that is $\alpha_B(No\ 2|h_B(x^0)) = 0$ which implies $\beta_A(No\ 2|h_A(x^1)) = 0$ so that $D_{AB} = 0$. *Ann* is therefore certain that she will get a psychological payoff of 4 if she takes the exam on the date *No.* 1. If *Ann* takes the exam on the second date then she is certain that she will not feel guilt, however she will only get a psychological payoff of 2. *Ann*'s only rationalizable strategy is to take the exam on the first date (that is, choose *No* 1).

Like the previous example featuring reciprocity, also this example with guilt aversion demonstrates the impact of unawareness on players motivated by psychological payoffs and highlights the subtle incentive to disclose paths of play to the unaware player. In this game *Ann* has no incentive to make *Bob* aware of the second date as long as she is sure that he will never become aware of it because this leaves her with a payoff of 4 instead of 2. Interestingly, were *Ann* only interested in his own material payoff, she would not be concerned about *Bob* being or not being aware of her action *No* 2. *Ann* would in any case choose to write the exam on date *No* 1 irrespective of *Bob*'s awareness.

6 Conclusion

We have analyzed the influence and importance of simple unawareness concerning feasible paths of play for the strategic interaction of players in dynamic psychological games, and defined a two-player model in which players are motivated by psychological payoffs and simple unawareness of certain feasible

paths of play. Using this model we provide different examples highlighting the role of unawareness in the strategic interaction of players motivated by reciprocity à la Dufwenberg and Kirchsteiger (2004) and guilt aversion à la Battigalli and Dufwenberg (2007). Our examples demonstrate that the strategic behavior of players motivated by belief-dependent preferences crucially depends on their awareness concerning the strategic environment they are in, their perception concerning the awareness of others, their perception concerning the perception of others etc. In other words, unawareness has a profound and intuitive impact on the strategic interaction of players motivated by psychological payoffs – a fact that creates both an opportunity as well as a challenge to empirically investigations analyzing the strength and nature of psychological payoffs. Concentrating on two-player environments and simple awareness scenarios obviously puts limits to the strategic situations that can be analyzed with our model. Nevertheless our simple model has allowed us to uncover intriguing effects. More general strategic environments with more complex awareness scenarios are left for future research.

Acknowledgements: We would like to thank the editor, Burkhard Schipper, and an anonymous referee for their very helpful comments. We are also grateful to Geir Asheim, Pierpaolo Battigalli, Mie la Cour Sonne, Martin Dufwenberg, Aviad Heifetz, Georg Kirchsteiger, Peter Norman Sø rensen, Lars Peter Ø sterdal and participants at the EDGE Jamboree in Dublin, the DGPE Workshop in Copenhagen, the Econometric Society's European Winter Meeting in Rome, the Royal Economic Society's Postgraduate Meeting in London, the Seminar in Microeconomics in Lund, the CNEE Workshop in Copenhagen and the CES seminar in Leuven for helpful comments and suggestions. Both authors gratefully acknowledge the financial support from the Danish Council for Independent Research in Social Sciences (Grant ID: DFF-4003-00032). All errors are our own.

References

Andreoni, J., and B. Bernheim. 2009. "Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica* 77 (5):1607–36.

Battigalli, P. 1997. "On Rationalizability in Extensive Games." *Journal of Economic Theory* 74 (1):40–61.

Battigalli, P., and M. Dufwenberg. 2007. "Guilt in Games." *The American Economic Review, Papers and Proceedings* 97 (2):170–6.

Battigalli, P., and M. Dufwenberg. 2009. "Dynamic Psychological Games." *Journal of Economic Theory* 144:1–35.

Battigalli, P., and M. Siniscalchi. 2002. "Strong Belief and Forward Induction Reasoning." *Journal of Economic Theory* 106 (2):356–91.

Bellemare, C., A. Sebald, and M. Strobel. 2011. "Measuring the Willingness to Pay to Avoid Guilt: Estimation Using Equilibrium and Stated Belief Models." *Journal of Applied Econometrics* 26 (3):437–53.

Broberg, T., T. Ellingsen, and M. Johannesson. 2007. "Is Generosity Involuntary?." *Economics Letters* 94 (1):32–7.

Charness, G., and M. Dufwenberg. 2006. "Promises and Partnership." *Econometrica* 74 (6):1579–601.

Dana, J., D. Cain, and R. Dawes. 2006. "What You Don't Know Won't Hurt Me: Costly (but Quiet) Exit in Dictator Games." *Organizational Behavior and Human Decision Processes* 100 (2):193–201.

Dana, J., R. Weber, and J. Kuang. 2007. "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness." *Economic Theory* 33 (1):67–80.

Dekel, E., B. Lipman, and A. Rustichini. 1998. "Standard State-Space Models Preclude Unawareness." *Econometrica* 66:159–73.

Dufwenberg, M., and G. Kirchsteiger. 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior* 47 (2):268–98.

Fagin, R., and J. Halpern. 1988. "Beliefs, Awareness, and Limited Reasoning." *Artificial Intelligence* 34:39–76.

Falk, A., E. Fehr, and U. Fischbacher. 2008. "Testing Theories of Fairness–Intentions Matter." *Games and Economic Behavior* 62 (1):287–303.

Falk, A., and U. Fischbacher. 2006. "A Theory of Reciprocity." *Games and Economic Behavior* 54 (2):293–315.

Fehr, E., G. Kirchsteiger, and A. Riedl. 1993. "Does Fairness Prevent Market Clearing? an Experimental Investigation." *The Quarterly Journal of Economics* 108 (2):437–59.

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

Feinberg, Y. 2004. "Subjective Reasoning Games with Unawareness." Stanford University Research Paper (1875).

Fudenberg, D., and J. Tirole. 1991. *Game Theory*. Cambridge: MIT Press.

Geanakoplos, J., D. Pearce, and E. Stacchetti. 1989. "Psychological Games and Sequential Rationality." *Games and Economic Behavior* 1 (1):60–79.

Grant, S., and J. Quiggin. 2013. "Inductive Reasoning About Unawareness." *Economic Theory* 54 (3):717–55.

Halpern, J. 2001. "Alternative Semantics for Unawareness." *Games and Economic Behavior* 37:321–39.

Halpern, J., and L. Rêgo. 2006. "Extensive Games with Possibly Unaware Players." In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, ACM, pp. 744–751.

Halpern, J., and L. Rêgo. 2008. "Interactive Unawareness Revisited." *Games and Economic Behavior* 62:323–262.

Halpern, J., and L. Rêgo. 2009. "Reasoning About Knowledge of Unawareness." *Games and Economic Behavior* 67 (2):503–25.

Heifetz, A., M. Meier, and B. Schipper. 2006. "Interactive Unawareness." *Journal of Economic Theory* 130:78–94.

Heifetz, A., M. Meier, and B. Schipper. 2008. "A Cannonical Model for Interactive Unawareness." *Games and Economic Behavior* 62:304–24.

Heifetz, A., M. Meier, and B. Schipper. 2013a. "Dynamic Unawareness and Rationalizable Behavior." *Games and Economic Behavior* 81:50–68.

Heifetz, A., M. Meier, and B. Schipper. 2013b. "Unawareness, Beliefs, and Speculative Trade." *Games and Economic Behavior* 77:100–21.

Lazear, E., U. Malmendier, and R. Weber. 2012. "Sorting in Experiments with Application to Social Preferences." *American Economic Journal: Applied Economics* 4 (1):136–63.

Li, J. 2009. "Information Structures with Unawareness." *Journal of Economic Theory* 144 (3):977–93.

Modica, S., and A. Rustichini. 1999. "Unawareness and Partitional Information Structures." *Games and Economic Behavior* 27:265–98.

Pearce, D. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52 (4):1029–50.

Rabin, M. 1993. "Incorporating Fairness Into Game Theory and Economics." *The American Economic Review* 83 (5):1281–302.

Sebald, A. 2010. "Attribution and Reciprocity." *Games and Economic Behavior* 68 (1):339–52.