

Meier, Martin; Schipper, Burkhard C.

Working Paper

Bayesian games with unawareness and unawareness perfection

Working Paper, No. 12-9

Provided in Cooperation with:

University of California Davis, Department of Economics

Suggested Citation: Meier, Martin; Schipper, Burkhard C. (2012) : Bayesian games with unawareness and unawareness perfection, Working Paper, No. 12-9, University of California, Department of Economics, Davis, CA

This Version is available at:

<https://hdl.handle.net/10419/58374>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



DEPARTMENT OF ECONOMICS

Working Paper Series

Bayesian Games with Unawareness and Unawareness Perfection

Martin Meier
Institute for Advanced Studies, Vienna
Burkhard Schipper
University of California, Davis

May 08, 2012

Paper # 12-9

Applying unawareness belief structures introduced in Heifetz, Meier, and Schipper (2012), we develop Bayesian games with unawareness, define equilibrium, and prove existence. We show how equilibria are extended naturally from lower to higher awareness levels and restricted from higher to lower awareness levels. We apply Bayesian games with unawareness to investigate the robustness of equilibria to uncertainty about opponents' awareness of actions. We show that a Nash equilibrium of a strategic game is robust to unawareness of actions if and only if it is not weakly dominated. Finally, we discuss the relationship between standard Bayesian games and Bayesian games with unawareness.

Department of Economics
One Shields Avenue
Davis, CA 95616
(530)752-0741

http://www.econ.ucdavis.edu/working_search.cfm

BAYESIAN GAMES WITH UNAWARENESS AND UNAWARENESS PERFECTION*

Martin Meier[†] Burkhard C. Schipper[‡]

This draft: May 8, 2012
First draft: June 6, 2005

Abstract

Applying unawareness belief structures introduced in Heifetz, Meier, and Schipper (2012), we develop Bayesian games with unawareness, define equilibrium, and prove existence. We show how equilibria are extended naturally from lower to higher awareness levels and restricted from higher to lower awareness levels. We apply Bayesian games with unawareness to investigate the robustness of equilibria to uncertainty about opponents' awareness of actions. We show that a Nash equilibrium of a strategic game is robust to unawareness of actions if and only if it is not weakly dominated. Finally, we discuss the relationship between standard Bayesian games and Bayesian games with unawareness.

Keywords: Unawareness, Awareness, Type-space, Incomplete information, Bayesian games, Equilibrium, Perfection, Undominated equilibrium, Weak dominance, Inattention.

JEL-Classifications: C70, C72, D80, D82.

*We thank Aviad Heifetz for his substantial input. Martin acknowledges financial support from the Spanish Ministerio de Educación y Ciencia via a Ramon y Cajal Fellowship and Research Grant SEJ2004-07861, and the Fundación Ramón Areces. Burkhard received financial support from the NSF SES-0647811, DFG SFB/TR 15, Minerva Stiftung, and IGA-UCD. We thank Leandro Chaves Rêgo, Eddie Dekel, Enrique Kawamura, Li Jing, Salvatore Modica, Klaus Nehring and participants in the Workshop on Unawareness at Stanford University, 2006, NASEM 2006, LOFT 2006, LAMES 2006, SAET 2007, TARK 2007, Stony Brook 2007, ESEM 2007, NBER GE conference 2007, GAMES 2008 and seminars at Bocconi, UCLA, St. Louis, Tucson, and Venice for helpful comments. Some of the material was circulated previously in a working paper titled “Unawareness, Beliefs and Games”.

[†]Institut für Höhere Studien, Wien. Email: martin.meier@ihs.ac.at

[‡]Department of Economics, University of California, Davis. Email: bcschipper@ucdavis.edu

1 Introduction

In standard Bayesian games, every player is able to conceive of all players, actions, states, types, and payoff functions. Players may be uncertain only about which of the conceived types of players they actually face, i.e., which state obtains. Yet, in real life people are not just uncertain about which states obtain but some contingencies may be completely out of their mind when taking decisions. Consequently, when forming beliefs and beliefs about beliefs of other players etc., those unforeseen contingencies are left out entirely and people may not even realize that those contingencies are left out. Moreover, since several players may be involved, players may also form beliefs about the unawareness of other players, beliefs about beliefs of other players about the unawareness of yet other players etc. Standard type spaces are not adequate for capturing unawareness (Modica and Rustichini, 1994, Dekel, Lipman, and Rustichini, 1998). In this paper, we apply type spaces with unawareness, so called unawareness belief structures introduced in Heifetz, Meier, and Schipper (2012), to develop Bayesian games with unawareness.

Heifetz, Meier, and Schipper (2012) showed how an unawareness belief structure consisting of a *lattice of spaces* is adequate for modeling mutual unawareness. Every space in the lattice captures one particular collection of contingencies. Higher spaces capture richer collections of contingencies, in which states correspond to situations described by a larger set of contingencies. The join of several spaces – the lowest space at least as high as every one of them – corresponds to the union of contingencies expressible in these spaces. For every player, a type mapping associates with each state ω a probabilistic belief over states in some space that might not contain ω . Conditions are imposed to relate beliefs across different spaces of the lattice. Thus, at each state, a player has a belief over all underlying uncertainties describable in the space on which this belief is concentrated. But she may be unaware regarding other uncertainties not expressible in that space. Moreover, at each state each player has beliefs about the other players' beliefs and awareness, their beliefs about other players' beliefs and awareness etc.

Unawareness belief structures capture unawareness and beliefs, beliefs about beliefs (including beliefs about unawareness), beliefs about that etc. in a parsimonious way familiar from standard type spaces. That is, hierarchies of beliefs are captured implicitly by states and type mappings. A construction of unawareness belief structures from explicit hierarchies of beliefs is complicated by the multiple awareness levels involved. A player with a given awareness level may believe that another player has a lower awareness level and believes that the first player has yet a lower awareness level etc. Yet, in Heifetz, Meier, and Schipper (2012) we also present such a hierarchical construction and show the existence of a universal unawareness type space that contains all belief hierarchies.¹

In this paper, we complement the unawareness belief structure with a set of actions and a utility function for each player. This defines Bayesian games in which players may not just be uncertain about events but also unaware of some events. We extend this

¹Heinsalu (2011) independently proves the existence of a universal unawareness type space. However, he does not present an explicit construction of hierarchies of beliefs.

definition to allow also for uncertainty about and unawareness of the existence of some actions, outcomes, and players.

The definition of a strategy in Bayesian games with unawareness is not obvious. Consider a type τ who is aware of few contingencies only, and two other types τ', τ'' with a richer awareness that agree with the quantitative beliefs of τ regarding the aspects of reality of which τ is aware; the beliefs of τ' and τ'' differ only concerning dimensions of the reality that τ does not conceive. Should the action taken by τ necessarily be some average of the actions taken by τ' and τ'' ? We believe that conceptually, the answer to this question is negative. When the player conceives of more parameters (e.g. motives for saving) as relevant to her decision, her optimal action (e.g. “invest in bonds” or “invest in stocks”) need not be related to her optimal decision (e.g. “go shopping”) when these parameters are not part of the vocabulary with which she conceives the world.²

The next step is to define Bayesian equilibrium. Analogous to standard Bayesian games, an equilibrium in a Bayesian game with unawareness is a Nash equilibrium among types. Unawareness, however, introduces a new aspect to the construction of equilibrium: A type who conceives of only few dimensions of reality does not have in mind types of other players with a wider horizon, so the optimal action of this type *does not depend* on the actions of these wider-horizon types. Those types, however, who assign a positive probability to this narrow-minded type, must take its action into account when optimizing. With finitely many states, existence follows from Nash (1950).

We apply Bayesian games with unawareness and Bayesian Nash equilibrium to analyze the robustness of equilibrium to small uncertainty about players’ awareness of actions in strategic games. We introduce a Nash equilibrium refinement, called *Unawareness Perfect Equilibrium*, and prove existence in finite strategic games. For any finite strategic game we consider a sequence of Bayesian games with unawareness over actions that converge to the finite strategic game. This represents the players’ uncertainty over the opponents’ awareness of actions. An Unawareness Perfect Equilibrium is the limit of equilibria of this sequence as uncertainty over opponents’ awareness of actions goes to zero. It turns out that our refinement characterizes undominated Nash equilibrium. At a first glance, such a characterization may look somewhat surprising because the underlying assumption of undominated Nash equilibrium is that every player’s equilibrium strategy should be robust to slight mistaken choices by opponents. Every player should be cautious or prudent with respect to the rationality of opponents, which leads him to believe that *no opponents’ actions can be excluded* from being played. This is different from our idea that there is a slight chance that due to opponents’ unawareness of some actions *any opponents’ action may be excluded* from being played. If a player is unaware of some actions, then she perceives a partial game in which these actions are missing. We assume that players are cautious or prudent in the sense of not excluding such partial games when considering the possible unawareness of opponents. An Unawareness Perfect Equilibrium strategy is robust to misperceptions or more aptly “partial perceptions” of the game by

²This is a crucial point in which our definition of a strategy differs from the one in the parallel work of Sadzik (2006).

opponents.

The most prominent equilibrium concept ruling out dominated Nash equilibrium is Trembling Hand Perfect Equilibrium by Selten (1975). Every Trembling Hand Perfect Equilibrium is undominated but the converse applies to two-player games only (see van Damme, 1991). Selten (1975, p. 35) remarks that there is an inconsistency in the interpretation of Trembling Hand Perfect Equilibrium: “There cannot be any mistakes if the players are absolutely rational. Nevertheless, a satisfactory interpretation of equilibrium points in extensive games seems to require that the possibility of mistakes is not completely excluded.” That is, Selten assumes that players are irrational with a small probability. A player is irrational if she chooses a strategy that does not maximize her payoff given her beliefs. Note that in an Unawareness Perfect Equilibrium, a player still chooses a strategy that maximizes her payoff given her beliefs but her beliefs may be constrained by her limited awareness. So, in this sense our characterization allows us to provide a justification for undominated Nash equilibrium without resorting to irrationality of players. Replacing irrational actions of opponents (i.e., trembles) by uncertainty about opponents’ unawareness is of conceptual significance because choices are endogenous but unawareness and beliefs are mostly involuntary and thus exogenously specified in a Bayesian game. The reinterpretation of undominated Nash equilibrium as equilibrium robust to “partial perceptions” of the game may be of interest to applied game theorists when contemplating which equilibrium refinement to apply in a context involving possible inattention by players.

The paper is organized as follows: In the remaining subsections of the introduction, we discuss the related literature and provide a simple example illustrating the concept of unawareness perfection. In Section 2 we present our interactive unawareness belief structure. Bayesian games with unawareness are developed in Section 3. In Section 4 we introduce the concept of unawareness perfection for strategic games, prove existence, and characterize it by undominated Nash equilibrium. In an appendix, we discuss the relationship between standard Bayesian games and Bayesian games with unawareness. All proofs are relegated to an appendix as well.

1.1 A Simplified Illustration of Unawareness Perfection

In this section we briefly discuss a simplified illustration of unawareness perfection. Although the precise construction in Section 4 is more involved, we believe that the sketched exposition put forward in this section conveys the main idea.

Consider for instance the strategic game γ^0 given by

γ^0	L	R
U	1, 1	2, 0
D	0, 2	2, 2

There are two pure equilibria, (U, L) and (D, R) . (Actually, these are all the equilibria. There is no mixed equilibrium in this game.) Which equilibrium of game γ^0 is robust to

player's uncertainty that the opponent may be unaware of some action?

Given the game γ^0 , derive a partially ordered set of restricted games by considering the set of all subsets of actions for all players partially ordered by set inclusion:

γ^1	L	γ^2	R	γ^3	L	R	γ^4	L	R
U	1, 1	U	2, 0	U	1, 1	2, 0	D	0, 2	2, 2
D	0, 2	D	2, 2						

γ^5	L	γ^6	L	γ^7	R	γ^8	R
U	1, 1	D	0, 2	U	2, 0	D	2, 2

For instance, a player in game γ^1 is unaware of column player's action R . Hence, she is unaware of games $\gamma^0, \gamma^2, \gamma^3, \gamma^4, \gamma^7$, and γ^8 . However, she can envision that her opponent may be unaware of some action in γ^1 and may view the game to be γ^5 . A player being unaware of action R only is said to have awareness level γ^1 .

Consider now the system of completely mixed beliefs over all restricted games including γ^0 in Table 1. Each row describes a completely mixed belief over games given the

Table 1: System of Completely Mixed Beliefs

Games Belief	γ^0	γ^1	γ^2	γ^3	γ^4	γ^5	γ^6	γ^7	γ^8
t^0	$1 - \varepsilon$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$	$\frac{\varepsilon}{8}$
t^1		$1 - \varepsilon$				$\frac{\varepsilon}{2}$	$\frac{\varepsilon}{2}$		
t^2			$1 - \varepsilon$					$\frac{\varepsilon}{2}$	$\frac{\varepsilon}{2}$
t^3				$1 - \varepsilon$		$\frac{\varepsilon}{2}$		$\frac{\varepsilon}{2}$	
t^4					$1 - \varepsilon$		$\frac{\varepsilon}{2}$		$\frac{\varepsilon}{2}$
t^5						1			
t^6							1		
t^7								1	
t^8									1

awareness level associated with that row. E.g., t^0 denotes the completely mixed belief over games for a player with awareness level γ^0 . Such player may believe with probability $t^0(\gamma^1)$ that the opponent's awareness level is γ^1 . Such an opponent's belief about the player's awareness level is then given in turn by t^1 . A player with awareness level γ^0 who believes with probability $t^0(\gamma^1)$ that his opponent has awareness level γ^1 also believes that such an opponent is unaware of γ^0 and believes with probability $t^1(\gamma^5)$ that the player himself has awareness level γ^5 . Etc. Essentially this corresponds to a Bayesian game with unawareness of actions.

We are interested to find out which equilibrium of the game γ^0 is robust to such beliefs about opponent's unawareness in the limit as ε goes to zero, assuming that at

each awareness level and for each ε considered, players play Bayesian Nash equilibrium. Such equilibria we call Unawareness Perfect Equilibria. To construct such an equilibrium, consider the games at the lowest levels γ^5 , γ^6 , γ^7 , and γ^8 . Since the set of outcomes is a singleton, for any ε the Bayesian equilibria are trivial in those games. Now at any higher awareness level, players must take the Bayesian equilibria of the lower games into account. Thus we can define inductively Bayesian equilibria with unawareness. When we consider games γ^8 to γ^1 in our example, this yields a unique outcome for all games except γ^2 and γ^4 . In latter two games any mixtures of the row and column players, respectively, are allowed. Since t^0 is completely mixed over all lower games, every action of the opponent is assigned some strict positive weight in equilibrium at awareness level γ^0 as long as $\varepsilon > 0$. The best reply of the type t^0 is always to play U as row player and L as column player. Taking ε to zero selects uniquely the equilibrium (U, L) as the Unawareness Perfect Equilibrium of γ^0 . This corresponds to the undominated Nash equilibrium. In the paper, we show that this equivalence holds more generally. Every Unawareness Perfect Equilibrium of a finite strategic game is an undominated equilibrium and vice versa. Section 4 develops this in a general framework.

1.2 Related Literature

There is a growing literature on unawareness both in economics and computer science.³ The independent parallel work of Sadzik (2006) is closest to ours. Building to a certain extent on our earlier work, Heifetz, Meier, and Schipper (2006), he presents a framework of unawareness with probabilistic beliefs in which the common prior on the upmost space is a primitive. In contrast, we take types as primitives and a prior may be defined on the entire unawareness belief structure as a convex combination of the type's beliefs (see Heifetz, Meier, and Schipper, 2012). Sadzik (2006) also considers Bayesian games with unawareness, but his definition of Bayesian strategy and consequently the notion of equilibrium differs from ours. As argued above, we do not confine actions of a type with a narrow horizon to be some average of actions of the corresponding types with a wider horizon, a restriction made in Sadzik (2006). As a result, in our notion of Bayesian equilibrium every type maximizes and is certain that every other type that she is aware of maximizes as well, while in the equilibrium of actions proposed in Sadzik (2006) a type may believe that another player is irrational. Sadzik (2006) does not allow for players to be unaware of other players, while we do.

A purely syntactic framework with unawareness is presented by Feinberg (2005) which he applies to games with unawareness of actions but complete information. In Section 3.2, we discuss an interesting example due to Feinberg (2005) and demonstrate that higher order awareness of unawareness in Feinberg (2005) corresponds to higher order belief of unawareness in our model. Feinberg (2009) discusses games with unawareness by modeling games and many views thereof, each (mutual) view being a finite sequence of player names i_1, \dots, i_n with the interpretation that this is how i_1 views how how i_n

³See <http://www.econ.ucdavis.edu/faculty/schipper/unaw.htm> for a bibliography.

views the game. This differs from our unawareness belief structures in which each state “encapsulates” the views of the players, their views about other players’ views etc. in a parsimonious way familiar from standard structures.

In a framework similar to Feinberg (2005, 2009), Čopić and Galeotti (2006) study two-player games with either unawareness of actions or unawareness of types (with a prior as a primitive). Yet, their notion of equilibrium differs from Bayesian equilibrium because the authors require that in equilibrium beliefs over actions and payoffs must correspond to the true joint distribution over own payoffs and the opponent’s actions.

Halpern and Rêgo (2006), Rêgo and Halpern (2012), Li (2006), Ozbay (2007), Heifetz, Meier, and Schipper (2011a, b), Meier and Schipper (2012), and Feinberg (2009) present models of extensive-form games with unawareness and discuss solution concepts for them.

Bayesian games with unawareness allow both for unawareness and probability zero beliefs. This raises the question about the differences between probability zero events and events that an agent is unaware of. In Appendix A, we show how to “flatten” a Bayesian game with unawareness by taking the union of all spaces and assigning zero probability to all states of which the individual is unaware. The “flattened” game is a standard Bayesian game with a standard type space; thus the Dekel-Lipman-Rustichini (1998) critique applies and unawareness becomes trivial. “Flattening” does not “change” the set of Bayesian Nash equilibria though but equilibria in the “flattened” game cannot be interpreted anymore with unawareness. At an epistemic level, unawareness has very different properties from probability zero belief. For instance, one property that is satisfied by unawareness is symmetry (see Heifetz, Meier, and Schipper, 2012, Proposition 5). An agent is unaware of an event if and only if she is unaware of its negation. Clearly, such a property cannot be satisfied by probability zero belief because if an agent assigns probability zero to an event, then she must assign probability one to its complement. Schipper (2012) shows that this feature captures also behavioral differences between unawareness and probability zero belief. Let’s say a decision maker chooses among different contracts for buying a firm. The second contract may differ from a first contract only in a consequence for an event E that is disadvantageous to the buyer. If the decision maker is indifferent between both contracts, then this is consistent with E being Savage null. Yet, if the decision maker is also indifferent between the first and a third contract that differs from the first only in assigning this disadvantageous consequence to the negation of the event E instead the event E itself, then this behavior is inconsistent with the negation of the event E or the event E itself being Savage null. The decision maker behaves as if both the event E and its negation are Savage null, which is impossible but consistent with unawareness of the E and of its negation. Thus, when the primitives of a decision model are fixed, unawareness has behavioral implications distinct from zero probability.

The literature on unawareness is related to the recent work in behavioral economics, finance, and macroeconomics that discusses the economic relevance of peoples inattention for various economic outcomes such as retirement savings, choice of health care plans, etc. Apart from the questions of how to design optimally economic policies to “nudge” people’s

attention (Thaler and Sunstein, 2008) or of how to optimally allocate (voluntary) inattention (e.g. Sims, 2010, Van Nieuwerburgh and Veldkamp, 2010), it should be of interest to study more generally which outcomes of strategic interaction are robust to inattention of players. While the notions of inattention discussed in behavioral economics, finance, and macroeconomics may not correspond exactly to the well-defined epistemic notion of unawareness and may additionally involve biases and features of bounded rationality, we believe that unawareness may be one component of those notions of inattention.

2 Model

2.1 State-Spaces

2.2 State-Spaces

Let $\mathcal{S} = \{S_\alpha\}_{\alpha \in \mathcal{A}}$ be a complete lattice of disjoint *state-spaces*, with the partial order \succeq on \mathcal{S} . A complete lattice is a lattice such that each subset has a least upper bound (i.e., supremum) and a greatest lower bound (i.e., infimum). If S_α and S_β are such that $S_\alpha \succeq S_\beta$ we say that “ S_α is more expressive than S_β – states of S_α describe situations with a richer vocabulary than states of S_β ”.⁴ Denote by $\Omega = \bigcup_{\alpha \in \mathcal{A}} S_\alpha$ the union of these spaces. Each $S \in \mathcal{S}$ is a measurable space, with a σ -field \mathcal{F}_S .

Spaces in the lattice can be more or less “rich” in terms of facts that may or may not obtain in them. The partial order relates to the “richness” of spaces.

2.3 Projections

For every S and S' such that $S' \succeq S$, there is a measurable surjective projection $r_S^{S'} : S' \rightarrow S$, where r_S^S is the identity. (“ $r_S^{S'}(\omega)$ is the restriction of the description ω to the more limited vocabulary of S .”) Note that the cardinality of S is smaller than or equal to the cardinality of S' . We require the projections to commute: If $S'' \succeq S' \succeq S$ then $r_S^{S''} = r_S^{S'} \circ r_{S'}^{S''}$. If $\omega \in S'$, denote $\omega_S = r_S^{S'}(\omega)$. If $D \subseteq S'$, denote $D_S = \{\omega_S : \omega \in D\}$.

Projections “translate” states from “more expressive” spaces to states in “less expressive” spaces by “erasing” facts that can not be expressed in a lower space.

2.4 Events

Denote $g(S) = \{S' : S' \succeq S\}$. For $D \subseteq S$, denote $D^\dagger = \bigcup_{S' \in g(S)} (r_S^{S'})^{-1}(D)$. (“All the extensions of descriptions in D to at least as expressive vocabularies.”)

⁴Here and in what follows, phrases within quotation marks hint at intended interpretations, but we emphasize that these interpretations are not part of the definition of the set-theoretic structure.

An *event* is a pair (E, S) , where $E = D^\uparrow$ with $D \subseteq S$, where $S \in \mathcal{S}$. D is called the *base* and S the *base-space* of (E, S) , denoted by $S(E)$. If $E \neq \emptyset$, then S is uniquely determined by E and, abusing notation, we write E for (E, S) . Otherwise, we write \emptyset^S for (\emptyset, S) . Note that not every subset of Ω is an event.

Some fact may obtain in a subset of a space. Then this fact should be also “expressible” in “more expressive” spaces. Therefore the event contains not only the particular subset but also its inverse images in “more expressive” spaces.

To illustrate the definition of event, consider Figure 1. The event $[p]$ is the union of dotted areas.

Let Σ be the set of *measurable events* of Ω , i.e., D^\uparrow such that $D \in \mathcal{F}_S$, for some state-space $S \in \mathcal{S}$. Note that unless \mathcal{S} is a singleton, Σ is not an algebra because it contains distinct \emptyset^S for all $S \in \mathcal{S}$. The event \emptyset^S should be interpreted as a “logical contradiction phrased with the expressive power available in S ”. It is quite natural to have distinct vacuous events since contradictions can be phrased with differing expressive powers.

2.5 Negation

If (D^\uparrow, S) is an event where $D \subseteq S$, the negation $\neg(D^\uparrow, S)$ of (D^\uparrow, S) is defined by $\neg(D^\uparrow, S) := ((S \setminus D)^\uparrow, S)$. Note, that by this definition, the negation of a (measurable) event is a (measurable) event. Abusing notation, we write $\neg D^\uparrow := \neg(D^\uparrow, S)$. Note that by our notational convention, we have $\neg S^\uparrow = \emptyset^S$ and $\neg \emptyset^S = S^\uparrow$, for each space $S \in \mathcal{S}$. $\neg D^\uparrow$ is typically a proper subset of the complement $\Omega \setminus D^\uparrow$. That is, $(S \setminus D)^\uparrow \subsetneq \Omega \setminus D^\uparrow$.

Intuitively, there may be states in which the description of an event D^\uparrow is both expressible and valid – these are the states in D^\uparrow ; there may be states in which its description is expressible but invalid – these are the states in $\neg D^\uparrow$; and there may be states in which neither its description nor its negation are expressible – these are the states in

$$\Omega \setminus (D^\uparrow \cup \neg D^\uparrow) = \Omega \setminus S(D^\uparrow)^\uparrow.$$

Thus our structure is not a standard state-space model in the sense of Dekel, Lipman, and Rustichini (1998).

2.6 Conjunction and Disjunction

If $\left\{ (D_\lambda^\uparrow, S_\lambda) \right\}_{\lambda \in L}$ is a collection of events (with $D_\lambda \subseteq S_\lambda$, for $\lambda \in L$), their conjunction $\bigwedge_{\lambda \in L} (D_\lambda^\uparrow, S_\lambda)$ is defined by $\bigwedge_{\lambda \in L} (D_\lambda^\uparrow, S_\lambda) := \left(\left(\bigcap_{\lambda \in L} D_\lambda^\uparrow \right), \sup_{\lambda \in L} S_\lambda \right)$. Note, that since \mathcal{S} is a *complete* lattice, $\sup_{\lambda \in L} S_\lambda$ exists. If $S = \sup_{\lambda \in L} S_\lambda$, then we have $\left(\bigcap_{\lambda \in L} D_\lambda^\uparrow \right) = \left(\bigcap_{\lambda \in L} \left((r_{S_\lambda}^S)^{-1} (D_\lambda) \right) \right)^\uparrow$. Again, abusing notation, we write $\bigwedge_{\lambda \in L} D_\lambda^\uparrow := \bigcap_{\lambda \in L} D_\lambda^\uparrow$ (we will therefore use the conjunction symbol \wedge and the intersection symbol \cap interchangeably).

Intuitively, to take the intersection of events $(D_\lambda^\uparrow, S_\lambda)_{\lambda \in L}$, we express them “most economically in the smallest language” in which they are all expressible $S = \sup_{\lambda \in L} S_\lambda$, take the intersection, and then the union of inverse images obtaining the event $(\bigcap_{\lambda \in L} ((r_{S_\lambda}^S)^{-1}(D_\lambda)))^\uparrow$ that is based in S .

We define the relation \subseteq between events (E, S) and (F, S') , by $(E, S) \subseteq (F, S')$ if and only if $E \subseteq F$ as sets *and* $S' \preceq S$. If $E \neq \emptyset$, we have that $(E, S) \subseteq (F, S')$ if and only if $E \subseteq F$ as sets. Note however that for $E = \emptyset^S$ we have $(E, S) \subseteq (F, S')$ if and only if $S' \preceq S$. Hence we can write $E \subseteq F$ instead of $(E, S) \subseteq (F, S')$ as long as we keep in mind that in the case of $E = \emptyset^S$ we have $\emptyset^S \subseteq F$ if and only if $S \succeq S(F)$. It follows from these definitions that for events E and F , $E \subseteq F$ is equivalent to $\neg F \subseteq \neg E$ only when E and F have the same base, i.e., $S(E) = S(F)$.

Intuitively, to say “ E implies F ” we must be able to express F in the “language” used to express E . Hence, it must be that $S(F) \preceq S(E)$. The inclusion is then just $E \cap S(E) \subseteq F \cap S(E)$.

The disjunction of $\{D_\lambda^\uparrow\}_{\lambda \in L}$ is defined by the de Morgan law $\bigvee_{\lambda \in L} D_\lambda^\uparrow = \neg \left(\bigwedge_{\lambda \in L} \neg (D_\lambda^\uparrow) \right)$. Typically $\bigvee_{\lambda \in L} D_\lambda^\uparrow \subsetneq \bigcup_{\lambda \in L} D_\lambda^\uparrow$, and if all D_λ are nonempty we have that $\bigvee_{\lambda \in L} D_\lambda^\uparrow = \bigcup_{\lambda \in L} D_\lambda^\uparrow$ holds if and only if all the D_λ^\uparrow have the same base-space. Note, that by these definitions, the conjunction and disjunction of (at most countably many measurable) events is a (measurable) event.

Apart from the measurability conditions, the event-structure outlined so far is analogous to Heifetz, Meier, and Schipper (2006, 2008). An example is shown in Figure 1. It depicts a lattice with four spaces and projections. As mentioned previously, the event that p obtains is indicated by the dotted areas. The grey areas illustrate the event that not p obtains. $S_p \cup S_q$ is for instance not an event in our structure.

2.7 Probability Measures

Here and in what follows, the term ‘events’ always refers to measurable events in Σ unless otherwise stated.

Let $\Delta(S)$ be the set of probability measures on (S, \mathcal{F}_S) . We consider this set itself as a measurable space endowed with the σ -field $\mathcal{F}_{\Delta(S)}$ generated by the sets $\{\mu \in \Delta(S) : \mu(D) \geq p\}$, where $D \in \mathcal{F}_S$ and $p \in [0, 1]$.


2.8 Marginals

For a probability measure $\mu \in \Delta(S')$, the marginal $\mu|_S$ of μ on $S \preceq S'$ is defined by

$$\mu|_S(D) := \mu \left(\left(r_{S'}^{S'} \right)^{-1}(D) \right), \quad D \in \mathcal{F}_S.$$

Let S_μ be the space on which μ is a probability measure. Whenever $S_\mu \succeq S(E)$ then

Figure 1: Event Structure



we abuse notation slightly and write

$$\mu(E) = \mu(E \cap S_\mu).$$

If $S(E) \not\subseteq S_\mu$, then we say that $\mu(E)$ is undefined.

2.9 Types

I is the nonempty set of individuals. For every individual, each state gives rise to a probabilistic belief over states in some space.

Definition 1 For each individual $i \in I$ there is a type mapping $t_i : \Omega \rightarrow \bigcup_{\alpha \in \mathcal{A}} \Delta(S_\alpha)$, which is measurable in the sense that for every $S \in \mathcal{S}$ and $Q \in \mathcal{F}_{\Delta(S)}$ we have $t_i^{-1}(Q) \cap S \in \mathcal{F}_S$. We require the type mapping t_i to satisfy the following properties.⁵

- (0) *Confinement:* If $\omega \in S'$ then $t_i(\omega) \in \Delta(S)$ for some $S \preceq S'$.
- (1) If $S'' \succeq S' \succeq S$, $\omega \in S''$, and $t_i(\omega) \in \Delta(S)$ then $t_i(\omega_{S'}) = t_i(\omega)$.
- (2) If $S'' \succeq S' \succeq S$, $\omega \in S''$, and $t_i(\omega) \in \Delta(S')$ then $t_i(\omega_S) = t_i(\omega)|_S$.

⁵Recall that S_μ is the space on which μ is a probability measure. Thus, $S_{t_i(\omega)}$ is the space on which $t_i(\omega)$ is a probability measure.

(3) If $S'' \succeq S' \succeq S$, $\omega \in S''$, and $t_i(\omega_{S'}) \in \Delta(S)$ then $S_{t_i(\omega)} \succeq S$.

$t_i(\omega)$ represents individual i 's belief at state ω . Properties (0) to (3) guarantee the consistent fit of beliefs and awareness at different state-spaces. *Confinement* means that at any given state $\omega \in \Omega$ an individual's belief is concentrated on states that are all described with the same "vocabulary" - the "vocabulary" available to the individual at ω . This "vocabulary" may be less expressive than the "vocabulary" used to describe statements in the state ω .

Properties (1) to (3) compare the types of an individual in a state $\omega \in S'$ and its projection to ω_S , for some $S \preceq S'$. Property (1) and (2) mean that at the projected state ω_S the individual believes everything she believes at ω given that she is aware of it at ω_S . Property (3) means that at ω an individual cannot be unaware of an event that she is aware of at the projected state $\omega_{S'}$.

Remark 1 *Property (1) of the type mappings in Definition 1 is implied by the Properties (0), (2), and (3).*

For a proof, see Heifetz, Meier, and Schipper (2012).

Define⁶

$$Ben_i(\omega) := \left\{ \omega' \in \Omega : t_i(\omega')|_{S_{t_i(\omega)}} = t_i(\omega) \right\}.$$

This is the set of states at which individual i 's type or the marginal thereof coincides with her type at ω . Such sets are events in our structure (see Heifetz, Meier, and Schipper, 2012, for a proof):

Remark 2 *For any $\omega \in \Omega$, $Ben_i(\omega)$ is an $S_{t_i(\omega)}$ -based event, which is not necessarily measurable.⁷ We have $Ben_i(\omega) = \{\omega' \in S_{t_i(\omega)} : t_i(\omega') = t_i(\omega)\}^\uparrow = \{Ben_i(\omega) \cap S_{t_i(\omega)}\}^\uparrow$.*

Recall that by definition $t_i(\omega)(E) = t_i(\omega)(E \cap S_{t_i(\omega)})$. Moreover, recall that with event we mean measurable event in our event structure unless otherwise stated; both facts will be used throughout the paper.

Assumption 1 *If $Ben_i(\omega) \subseteq E$, for an event E , then $t_i(\omega)(E) = 1$.*

This assumption implies introspection (Property (va)) in Proposition 4 in Heifetz, Meier, and Schipper (2012). Note, that if $Ben_i(\omega)$ is measurable, then Assumption 1 is equivalent to $t_i(\omega)(Ben_i(\omega)) = 1$.

⁶The name "Ben" is chosen analogously to the "ken" in knowledge structures, see Samet (1990, p. 193).

⁷Even in a standard type-space, if the σ -algebra is not countably generated, then the set of states where a player is of a certain type might not be measurable.

Definition 2 We denote by $\underline{\mathcal{S}} := \left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, (t_i)_{i \in I} \right\rangle$ an interactive unawareness belief structure.

So far, the model outlined in this section corresponds to unawareness belief structures introduced in Heifetz, Meier, and Schipper (2012). In that paper, we also provide an explicit hierarchical construction and show the existence of a universal unawareness type space.

3 Bayesian Games with Unawareness

In this section, we generalize strategic games with incomplete information à la Harsanyi (1967/68) and Mertens and Zamir (1985, Section 5) to include also unawareness. For notational convenience, we restrict ourselves in this section to a finite set of players, finite sets of actions, and finite unawareness belief structures. A *finite* unawareness belief structure is an unawareness belief structure, where \mathcal{S} is finite, each $S \in \mathcal{S}$ is finite, and for all $S \in \mathcal{S}$, \mathcal{F}_S is the set of all subsets of S .

3.1 Unawareness of Payoff Relevant Events

For simplicity, we consider first Bayesian games with unawareness of payoff relevant events only, in which every player is aware of all of her and other's actions, and of all the players.

Definition 3 A Bayesian game with unawareness of events consists of an unawareness belief structure $\underline{\mathcal{S}} = \left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, (t_i)_{i \in I} \right\rangle$ that is augmented by a tuple $\langle (M_i)_{i \in I}, (u_i)_{i \in I} \rangle$ defined as follows: For each player $i \in I$, there is

- (i) a finite nonempty set of actions M_i , and
- (ii) a utility function $u_i : \left(\prod_{i \in I} M_i \right) \times \Omega \longrightarrow \mathbb{R}$.

The interpretation is as follows: At the beginning of a game, a state $\omega \in \Omega$ is realized. Player i does not observe the state but receives a signal $t_i(\omega)$ that provides her with some information about the state or projections thereof to lower spaces. I.e., if ω obtains, player i is of type $t_i(\omega)$. This signal is a belief about the likelihood of events on a certain space. A player's utility depends on her action, the actions chosen by other players as well as the state. Since players may be uncertain about the state ω , we assume below that the player's preference is represented by the expected value of the utility function on action-profiles of players and states, where the expectation is taken with respect to the player i 's type $t_i(\omega)$ and the types' mixed strategies. This game allows for unawareness of payoff relevant events.

Note that we allow unawareness to affect payoffs in an arbitrary way. One may consider letting $u_i(\cdot, \omega)$ be a convex combination of all $u_i(\cdot, \omega')$, $\omega' \in (r_S^{S'})^{-1}(\omega)$, for $\omega \in S$ and $S' \succeq S$. While this may be natural in some applications such as the speculation example in Heifetz, Meier, and Schipper (2012, Section 1.1), it would preclude situations in which the mere awareness of an event may reduce expected payoffs.

Let $\Delta(M_i)$ be the set of mixed strategies for player $i \in I$, that is, the set of probability distributions on the finite set M_i .

Definition 4 *A strategy of player i in a Bayesian game with unawareness of events is a function $\sigma_i : \Omega \rightarrow \Delta(M_i)$ such that for all $\omega, \omega' \in \Omega$, $t_i(\omega) = t_i(\omega')$ implies $\sigma_i(\omega) = \sigma_i(\omega')$.*

A strategy specifies for each player and state a probability distribution over her set of actions. In standard Bayesian games without unawareness, one interpretation of a strategy assumes an ex-ante point of view of the player before she knows her type. This interpretation is misleading in a game with unawareness, since if a player is aware of all her types ex-ante she should be also aware of all types interim, i.e., after learning her type (and her awareness). Hence, in the case of unawareness, the ex-ante notion of strategy is a construct for the game theorist rather than an object of choice for a player.

In Bayesian games with unawareness we subscribe to a second interpretation of Bayesian strategy from an interim point of view: Given a player i and type $t_i(\omega)$, she has an “awareness level” $S_{t_i(\omega)} \in \mathcal{S}$. That is, she can consider strategies of her opponents in $l(S_{t_i(\omega)})$, where $l(S) := \{S' \in \mathcal{S} : S' \preceq S\}$ is the complete sublattice of \mathcal{S} with S being the upmost space. This interpretation is sound precisely because of Proposition 2 and Remark 3 below: To best-respond to the strategies of the other player-types, a type of a player needs only to reason about the strategies of player-types that she is aware of. Only strategies of these player-types enter in her utility maximization problem.

Denote $\sigma_{S_{t_i(\omega)}} := \left((\sigma_j(\omega'))_{j \in I} \right)_{\omega' \in S_{t_i(\omega)}}$. A component $\sigma_j(\omega')$ of the strategy profile $\sigma_{S_{t_i(\omega)}}$ is the strategy of the *player-type* $(j, t_j(\omega'))$. $\sigma_{S_{t_i(\omega)}}$ is the profile of all player-types’ strategies in $S_{t_i(\omega)}$. (Recall that for some $\omega' \in S_{t_i(\omega)}$ and $j \in I$, $t_j(\omega')$ might have a lower awareness level than $S_{t_i(\omega)}$.)

The *expected utility* of player-type $(i, t_i(\omega))$ from the strategy profile $\sigma_{S_{t_i(\omega)}}$ is given by

$$U_{(i, t_i(\omega))}(\sigma_{S_{t_i(\omega)}}) := \int_{\omega' \in S_{t_i(\omega)}} \sum_{m \in \prod_{j \in I} M_j} \left(\prod_{j \in I} \sigma_j(\omega')(\{m_j\}) \right) \cdot u_i((m_j)_{j \in I}, \omega') dt_i(\omega)(\omega'). \quad (1)$$

$\sigma_j(\omega')(\{m_j\})$ is the probability with which the player-type $(j, t_j(\omega'))$ plays the action $m_j \in M_j$. $\prod_{j \in I} \sigma_j(\omega')(\{m_j\})$ is the joint probability with which the action profile $m = (m_j)_{j \in I}$ is played by the players. This action profile gives the utility $u_i((m_j)_{j \in I}, \omega')$ to player i in state ω' . The term $\sum_{m \in \prod_{j \in I} M_j} \left(\prod_{j \in I} \sigma_j(\omega')(\{m_j\}) \right) \cdot u_i((m_j)_{j \in I}, \omega')$ is

player i 's expected utility from the strategy profile $(\sigma_j(\omega'))_{j \in I}$ at the state ω' . However, at a state ω , the player, in general, does not know the state, but only his type $t_i(\omega)$, and so he evaluates his utility with the expectation with respect to the probability measure $t_i(\omega)$.

Definition 5 (Equilibrium) *Given a Bayesian game with unawareness of events $\left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, (t_i)_{i \in I}, (M_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$, define the associated strategic game by:*

(i) $\{(i, t_i(\omega)) : \omega \in \Omega \text{ and } i \in I\}$ *is the set of players,*

and for each player $(i, t_i(\omega))$,

(ii) *the set of mixed strategies is $\Delta(M_i)$, and*

(iii) *the utility function is given by Equation (1).*

A profile $(\sigma_i)_{i \in I}$ is an equilibrium of the Bayesian game with unawareness of events if and only if the following is an equilibrium of the associated strategic game: $(i, t_i(\omega))$ plays $\sigma_i(\omega)$, for all $i \in I$ and $\omega \in \Omega$.

An equilibrium of a Bayesian game with unawareness is a Nash equilibrium of a strategic game in which types of players are the “players”. The actions available to the type of player i at state ω are the actions of player i . The utility function of the type of player i at ω is the expected utility function, given player i 's awareness and belief over states at ω . In an equilibrium of a Bayesian game with unawareness of events, the type of every player chooses an optimal mixture of actions, given her awareness, belief and the choices of the types of the other players she is aware of. This is analogous to equilibrium in Bayesian games without unawareness of actions.

3.2 Allowing for Unawareness of Actions

Bayesian games with unawareness of events in Definition 3 do not allow us to model properly unawareness of actions. In standard Bayesian game theory, ignorance of actions is modeled by the assumption that players will never use such actions, because extremely low payoffs (i.e., highly negative) are assigned to those actions (see the discussion in Harsanyi, 1967, p. 168). We do not follow this convention here. Even in standard Bayesian games this convention is questionable, because it applies only to rational players. If there is lack of common belief of rationality then a player's type being ignorant of an action is indeed different from her obtaining a very low payoff from playing this action (see Hu and Stuart, 2001, for a discussion). In this subsection we introduce unawareness of actions, and discuss the notion of strategy in Bayesian games with unawareness.

Denote $[t_i(\omega)] := \{\omega' \in \Omega : t_i(\omega') = t_i(\omega)\}$. This is the set of states at which player i has the same type as in state ω .

Definition 6 A Bayesian game with unawareness of events and actions consists of a unawareness belief structure $\left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, (t_i)_{i \in I} \right\rangle$ that is augmented by a tuple $\langle (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \rangle$ defined as follows:

(i) a finite nonempty set of actions M_i , for $i \in I$, and a correspondence $\mathcal{M}_i : \Omega \longrightarrow 2^{M_i} \setminus \{\emptyset\}$, for $i \in I$, such that for any nonempty subset of actions $M'_i \subseteq M_i$, $[M'_i] := \{\omega \in \Omega : M'_i \subseteq \mathcal{M}_i(\omega)\}$ is an event (in the unawareness belief structure), and $\omega', \omega'' \in [t_i(\omega)] \cap S_{t_i(\omega)}$ implies $\mathcal{M}_i(\omega') = \mathcal{M}_i(\omega'')$, for all $\omega \in \Omega$,

(ii) for every $i \in I$, a utility function $u_i : \bigcup_{\omega \in \Omega} \left(\prod_{j \in I} \mathcal{M}_j(\omega) \right) \times \{\omega\} \longrightarrow \mathbb{R}$.

This definition allows for unawareness of events as well as actions. Which actions a player i has available at what state is described explicitly by the correspondence \mathcal{M}_i . Any set of available actions is associated with an event in our unawareness belief structure. We require that, for each type of each player, the sets of available actions are identical across states in the space on which this type is defined and at which the player's type coincides with this type. Note, that if $\omega \notin S_{t_i(\omega)}$, then it is possible that $\mathcal{M}_i(\omega')$ is a proper subset of $\mathcal{M}_i(\omega)$, for $\omega' \in [t_i(\omega)] \cap S_{t_i(\omega)}$. This allows in addition to unawareness of other players' actions also for unawareness of a player's own actions. Note that we exclude that at a state, a player considers it possible that she has an action available, which, in fact, is not available to her in this state. This is to avoid the conceptional problem of defining what should happen if a player is to take an action that is not available to her. Finally, observe that since outcomes consist of profiles of actions, players in a Bayesian game with unawareness of actions may also be unaware of outcomes in the game.

Definition 7 A strategy of player i in a Bayesian game with unawareness of actions is a function $\sigma_i : \Omega \longrightarrow \Delta(M_i)$ such that for all $\omega \in \Omega$,

(i) $\sigma_i(\omega) \in \Delta\left(\mathcal{M}_i(\omega_{S_{t_i(\omega)}})\right)$, and

(ii) $t_i(\omega') = t_i(\omega)$ implies $\sigma_i(\omega') = \sigma_i(\omega)$.

Example 1 (Feinberg, 2005) The following interesting game due to Feinberg (2005) is an example of unawareness of actions. It allows us also to compare our unawareness belief structures with the work by Feinberg (2005). Consider the strategic 3x3 game

		Colin		
		b_1	b_2	b_3
Rowena	a_1	0, 2	3, 3	0, 2
	a_2	2, 2	2, 1	2, 1
	a_3	1, 0	4, 0	0, 1

This game has a unique dominance solvable Nash equilibrium, (a_2, b_1) . Consider now a game with unawareness: The set of players remains unchanged, Rowena, R , and Colin, C . There are two state-spaces, S and S' with $S \succ S'$. In particular, $S = \{\omega_1, \omega_2\}$ and $S' = \{\omega_3\}$. The information structure is given by the type mappings

$$\begin{aligned} t_R(\omega_1)(\{\omega_2\}) &= t_R(\omega_2)(\{\omega_2\}) = t_R(\omega_3)(\{\omega_3\}) = 1, \\ t_C(\omega_1)(\{\omega_1\}) &= t_C(\omega_2)(\{\omega_3\}) = t_C(\omega_3)(\{\omega_3\}) = 1. \end{aligned}$$


Actions are specified by

$$\begin{aligned} \mathcal{M}_R(\omega_1) &= \mathcal{M}_R(\omega_2) = \{a_1, a_2, a_3\}, \mathcal{M}_R(\omega_3) = \{a_1, a_2\}, \\ \mathcal{M}_C(\omega_1) &= \mathcal{M}_C(\omega_2) = \mathcal{M}_C(\omega_3) = \{b_1, b_2, b_3\}. \end{aligned}$$

The information structure is depicted in Figure 2. The solid arrows and ellipses represent Rowena's information structure, while Colin's information structure is depicted with intermitted arrows and ellipses. At states ω_1 and ω_2 , payoffs are given by the above payoff matrix. At state ω_3 , payoffs are given by the sub-matrix spanned by rows a_1 and a_2 and columns b_1, b_2 , and b_3 in the above matrix, i.e.,

		Colin		
		b_1	b_2	b_3
Rowena	a_1	0, 2	3, 3	0, 2
	a_2	2, 2	2, 1	2, 1

Figure 2: Information Structure in Example 1



We claim that

$$(\sigma_R(\omega), \sigma_C(\omega)) = \begin{cases} (a_3, b_3) & \text{if } \omega = \omega_1 \\ (a_3, b_2) & \text{if } \omega = \omega_2 \\ (a_1, b_2) & \text{if } \omega = \omega_3 \end{cases}$$

is an equilibrium. To see this, note that the game at ω_3 has two pure equilibria, (a_2, b_1) and (a_1, b_2) in the S' -partial game, where the latter is payoff dominant. At ω_3 , both

players are unaware of action a_3 in ω_3 . The unique dominance solvable Nash equilibrium (a_2, b_1) of the original game (without unawareness of actions) remains an equilibrium because none of the players is unaware of an equilibrium action and equilibrium actions remain best responses if some other actions are deleted. Moreover, after deleting action a_3 (the action both players are unaware of at state ω_3 in S'), the game has another Nash equilibrium (a_1, b_2) . At ω_1 , both players are aware of all actions but Rowena believes that Colin is unaware of action a_3 . Hence Rowena believes that Colin thinks that (a_1, b_2) is a Nash equilibrium. Rowena's best response to Colin playing b_2 is a_3 . Moreover, since at ω_1 Colin is aware of all actions and he believes that Rowena believes that Colin is unaware of action a_3 , his best response to Rowena playing a_3 is b_3 . Note that in this equilibrium at ω_1 , both receive a low payoff (compared to the Nash equilibria discussed previously).

Feinberg (2005) obtains (a_3, b_3) as an equilibrium if both players are aware of all actions, Rowena is 'unaware' that Colin is aware of all of her actions, and Colin is 'aware' that Rowena is 'aware' of Colin being unaware of a_3 .⁸ That is, in Feinberg (2005) a player can be aware of an event but unaware that somebody else is aware of it. This is in contrast to our unawareness belief structure, where according to Proposition 8, 1., in Heifetz, Meier, and Schipper (2008) a player is aware of an event if and only if she is aware that somebody else could be aware of it. That is, if a player can reason about some issue then she can also reason that somebody else can reason about that issue. We obtain (a_3, b_3) as an equilibrium if both players are aware of all actions, Rowena does not believe that Colin is aware of a_3 , and Colin believes that Rowena believes that Colin is unaware of a_3 . The example suggests, that higher order 'awareness' in Feinberg (2005) operates like belief in our unawareness belief structure. Note however, that Feinberg (2005) does not define a notion of belief in his framework. \square

3.3 Allowing for Unawareness of Players

So far, we did not allow for unawareness of players. In standard Bayesian game theory, ignorance of players is modeled by dummy players. This is distinct from being unable to conceive of a player at all. In this subsection we allow for unawareness of players. This requires that we generalize our interactive unawareness belief structure such that a player may exist only at some states but not at others.

Definition 8 *A Bayesian game with unawareness is a tuple*

$$\Gamma(\underline{\mathcal{S}}) := \left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$$

defined as follows:

⁸When writing '...', we indicate that those notions differ from our notions used in this paper.

- (0) $\mathcal{S} = \{S_\alpha\}_{\alpha \in \mathcal{A}}$ is as before a complete lattice of spaces with surjective and commuting projections $(r_{S_\beta}^{S_\alpha})$, for $S_\beta \preceq S_\alpha$ (see Section 2).
- (i) $\mathcal{E} : I \longrightarrow \Sigma$ is the “existence” correspondence that assigns to each player $i \in I$ an event in which she exists. Moreover, $\mathcal{S}_i := \{S \in \mathcal{S} : \mathcal{E}(i) \cap S \neq \emptyset\}$ is the complete sublattice of spaces with states in which player i exists.
- (ii) For every player $i \in I$, $t_i : \mathcal{E}(i) \longrightarrow \bigcup_{S \in \mathcal{S}_i} \Delta(S)$ is a type mapping that satisfies Properties (0) to (3) (see Section 2) such that for every $\omega \in \mathcal{E}(i)$, $t_i(\omega)(\mathcal{E}(i)) = 1$.
- (iii) For every $i \in I$, M_i is a finite nonempty set of actions. $\mathcal{M}_i : \mathcal{E}(i) \longrightarrow 2^{M_i} \setminus \{\emptyset\}$ is a correspondence, for $i \in I$, with the following properties:
- a.) For every M'_i with $\emptyset \neq M'_i \subseteq M_i$: If there is a state ω with $\omega \in \mathcal{E}(i)$, then the set $\{\omega' \in \Omega : M'_i \subseteq \mathcal{M}_i(\omega')\}$ is an event.
 - b.) For $i \in I$: If $\omega \in \mathcal{E}(i)$ and $\omega', \omega'' \in S_{t_i(\omega)} \cap [t_i(\omega)]$, then $\mathcal{M}_i(\omega') = \mathcal{M}_i(\omega'')$.
- (vi) Further, we impose introspection as follows: For $\omega \in \mathcal{E}(i)$, $t_i(\omega)(\text{Ben}_i(\omega)) = 1$, where $\text{Ben}_i(\omega) := \{\omega' \in \mathcal{E}(i) : t_i(\omega')|_{S_{t_i(\omega)}} = t_i(\omega)\}$.
- (v) For $i \in I$, $u_i : \bigcup_{\omega \in \mathcal{E}(i)} \left(\prod_{j \in I(\omega)} \mathcal{M}_j(\omega) \right) \times \{\omega\} \longrightarrow \mathbb{R}$ is the utility function of i , where $I(\omega) := \{i \in I : \omega \in \mathcal{E}(i)\}$.

This game allows for unawareness of events, actions, outcomes, and players. For every player $i \in I$, the “existence” correspondence \mathcal{E} assigns to i the event in which she exists. Consequently we restrict player i ’s type mapping to states at which she exists. Moreover, player i ’s type is concentrated only on states in which she exists. A player can not assign strict positive probability to states at which she does not exist. The correspondence \mathcal{M}_i assigns a non-empty set of actions for player i only to the set of states in which player i exists. The dimension of the domain of a utility function may vary from state to state, since players may exist in some states but not in others, and each players utility at a state depends on the actions of all the players that exist in that state.

Note that if $\mathcal{E}(i) = \Omega$ for all $i \in I$, then we obtain an unawareness belief structure and a Bayesian game with unawareness of events and actions as defined before.

Note further that if $\omega \in \mathcal{E}(i)$, then $[t_i(\omega)] := \{\omega' \in \Omega : t_i(\omega') = t_i(\omega)\} \subseteq \mathcal{E}(i)$.

Definition 9 A strategy of player i in a Bayesian game with unawareness is a function $\sigma_i : \mathcal{E}(i) \longrightarrow \Delta(M_i)$ such that for all $\omega \in \mathcal{E}(i)$,

- (i) $\sigma_i(\omega) \in \Delta(\mathcal{M}_i(\omega_{S_{t_i}(\omega)}))$.
- (ii) $t_i(\omega') = t_i(\omega)$ implies $\sigma_i(\omega') = \sigma_i(\omega)$.

3.4 Equilibrium

In this section we define Bayesian equilibrium, prove existence and some properties for the most general games defined so far, Bayesian games with unawareness of players and actions.

Denote $\sigma_{S_{t_i}(\omega)} := \left((\sigma_j(\omega'))_{j \in I(\omega')} \right)_{\omega' \in S_{t_i}(\omega)}$. The *expected utility* of player-type $(i, t_i(\omega))$ from the strategy profile $\sigma_{S_{t_i}(\omega)}$ is given by

$$U_{(i, t_i(\omega))}(\sigma_{S_{t_i}(\omega)}) := \int_{\omega' \in S_{t_i}(\omega)} \sum_{m \in \prod_{j \in I(\omega')} \mathcal{M}_j \left(\omega'_{S_{t_j}(\omega')} \right)} \left(\prod_{j \in I(\omega')} \sigma_j(\omega')(\{m_j\}) \right) \cdot u_i((m_j)_{j \in I(\omega')}, \omega') dt_i(\omega)(\omega'). \quad (2)$$

Definition 10 (Equilibrium) *Given a Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}})$, define the associated strategic game by*

(i) $\{(i, t_i(\omega)) : \omega \in \Omega \text{ and } i \in I(\omega)\}$ *is the set of players,*

and for each player $(i, t_i(\omega))$,

(ii) *the set of mixed strategies is $\Delta(\mathcal{M}_i(\omega_{S_{t_i}(\omega)}))$, and*

(iii) *the utility function is given by Equation (2).*

A profile $(\sigma_i)_{i \in I}$ is an equilibrium of the Bayesian game with unawareness of events if and only if the following is an equilibrium of the associated strategic game: $(i, t_i(\omega))$ plays $\sigma_i(\omega)$, for all $i \in I(\omega)$ and $\omega \in \Omega$.

Proposition 1 (Existence) *Let $\Gamma(\underline{\mathcal{S}}) = \left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$ be a Bayesian game with unawareness. If I , Ω , and $(M_i)_{i \in I}$ are finite, then there exists an equilibrium.*

PROOF. By Nash's (1950) theorem. □

Recall $l(S) := \{S' \in \mathcal{S} : S' \preceq S\}$. $l(S)$ is a sublattice with S as the least upper bound.

Definition 11 *Given a Bayesian game with unawareness*

$$\Gamma(\underline{\mathcal{S}}) = \left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle,$$

we can define an S' -partial Bayesian game with unawareness

$$\Gamma(\underline{l(S')}) = \left\langle \underline{l(S')}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha \preceq S'}, \mathcal{E}', (t_i)_{i \in I(\Omega')}, (M_i)_{i \in I(\Omega')}, (\mathcal{M}'_i)_{i \in I(\Omega')}, (u_i)_{i \in I(\Omega')} \right\rangle,$$

in which $\mathcal{E}'(i) = \mathcal{E}(i) \cap \Omega'$, where $\Omega' = \bigcup_{S'' \in l(S')} S''$, and for any $i \in I(\Omega') := \bigcup_{\omega \in \Omega'} I(\omega)$, \mathcal{M}'_i is \mathcal{M}_i restricted to $\mathcal{E}'(i)$.

Note that contrary to an ordinary Bayesian game, the game is not “common knowledge” among the players. Let $\Gamma(\underline{\mathcal{S}}) = \left\langle \underline{\mathcal{S}}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$ be a Bayesian game with unawareness. At $\omega \in \Omega$, the game conceived by player j is $\Gamma(\underline{l(S_{t_j(\omega)})}) = \left\langle \underline{l(S_{t_j(\omega)})}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha \preceq S_{t_j(\omega)}}, \mathcal{E}', (t_i)_{i \in I(\Omega')}, (M_i)_{i \in I(\Omega')}, (\mathcal{M}'_i)_{i \in I(\Omega')}, (u_i)_{i \in I(\Omega')} \right\rangle$, where the lattice of spaces is replaced with the sublattice $\underline{l(S_{t_j(\omega)})}$ with $S_{t_j(\omega)}$ as the up-most space, and \mathcal{E}' (and hence the domains of t_i and u_i) is restricted to $\Omega' = \bigcup_{S \in \underline{l(S_{t_j(\omega)})}} S$. Type $t_j(\omega)$ of player j can conceive of all events expressible in the spaces of the sublattice $\underline{l(S_{t_j(\omega)})}$.

The following proposition shows that we can naturally extend equilibria from “lower awareness levels to higher awareness levels” by taking the equilibrium strategies at the “lower awareness levels” fixed and looking for a fixed point at “higher awareness levels”.

Proposition 2 (“Upwards Induction”) *Given a Bayesian game with unawareness $\left\langle \underline{\mathcal{S}}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$, consider for $S', S'' \in \mathcal{S}$ with $S' \preceq S''$ the S' -partial (resp. S'' -partial) Bayesian game with unawareness. If I , Ω , and $(M_i)_{i \in I}$ are finite, then for every equilibrium of the S' -partial Bayesian game, there is an equilibrium of the S'' -partial Bayesian game in which equilibrium strategies of player-types in $\{(i, t_i(\omega)) : \omega \in \Omega' = \bigcup_{S \in l(S')} S \text{ and } i \in I(\Omega')\}$ are identical with the equilibrium strategies in the S' -partial Bayesian game.*

This proposition suggests a procedure for constructing equilibria in Bayesian games with unawareness. We start with an equilibrium in the \hat{S} -partial Bayesian game with unawareness, where \hat{S} denotes the greatest lower bound (the meet) of the lattice, and extend it step-by-step to higher spaces by finding a fixed-point taking the strategies of player-types in the lower spaces as given.

For some strategic situations, Proposition 2 suggests that players which are unaware may have commitment power (although they do not understand that they are committed) compared to players with a “higher awareness level”. This is so because types with “lower awareness levels” do not react to types of which they are unaware. Types with “higher awareness” must take strategies of types with “lower awareness” as given. It is easy to construct examples of Bayesian games with unawareness in which the “value of awareness” may be negative. For instance, in a simultaneous-move linear Cournot

duopoly, a player who is unaware of his opponent can obtain the Stackelberg leader profit if the opponent knows that the first player is unaware of him.

We can also restrict an equilibrium from higher awareness levels to lower awareness. This is so, because if player-types play an equilibrium in a game that allows for “higher awareness levels”, then those player-types still play optimally at “lower awareness levels given that they exist there”. This is stated more formally in Remark 3.

Remark 3 Let $\left\langle \mathcal{S}, \left(r_{S_\beta}^{S_\alpha} \right)_{S_\beta \preceq S_\alpha}, \mathcal{E}, (t_i)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \right\rangle$ be a Bayesian game with unawareness. Consider for $S', S'' \in \mathcal{S}$ with $S' \preceq S''$ the S' -partial (resp. S'' -partial) Bayesian game with unawareness. Then for every equilibrium of the S'' -partial Bayesian game there is a unique equilibrium of the S' -partial Bayesian game in which the equilibrium strategies of player-types in $\{(i, t_i(\omega)) : \omega \in \Omega' = \bigcup_{S \in I(S')} S \text{ and } i \in I(\Omega')\}$ are identical to the equilibrium strategies of the S'' -partial Bayesian game.

4 Unawareness Perfection

In this section, we apply our framework to analyze Nash equilibria of strategic games that are robust to small uncertainty about awareness of actions. The main idea is to associate with a strategic game a sequence of Bayesian games with unawareness of actions and then consider the limit of equilibria of these games as uncertainty over awareness of actions vanishes.

We start by defining a sequence of specific Bayesian games with unawareness of actions. Let $\tilde{\gamma} = \langle I, (\tilde{M}_i)_{i \in I}, (\tilde{v}_i)_{i \in I} \rangle$ be a finite strategic game with a finite set of players $I = \{1, \dots, n\}$, for each player $i \in I$ a finite nonempty set of actions \tilde{M}_i and a payoff function $\tilde{v}_i : \tilde{M} \rightarrow \mathbb{R}$, where $\tilde{M} := \prod_{i \in I} \tilde{M}_i$.

Given such a strategic game, we append for each player $i \in I$ a “default” action d_i and extend the payoff functions such that any player’s default action is strictly dominated by every of her other actions. Moreover, whenever some player $i \in I$ takes her default action, all other players are indifferent among their actions (except their default action). I.e., for all $i \in I$, let $M_i = \tilde{M}_i \cup \{d_i\}$. For all $i \in I$, define $v_i : M \rightarrow \mathbb{R}$ with $v_i(m) = \tilde{v}_i(m)$ for all $m \in \tilde{M}$ and $v_i(d_i, m_{-i}) < v_i(m_i, m_{-i})$ for all $m_i \in \tilde{M}_i$ and all $m_{-i} \in M_{-i} = \prod_{j \in I \setminus \{i\}} M_j$, and $v_i(m_i, m_{-i}) = v_i(m'_i, m_{-i})$ for all $m_i, m'_i \in \tilde{M}_i$ and $m_{-i} = (m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n) \in M_{-i}$ with $m_j = d_j$ for some $j \in I \setminus \{i\}$. We write $\gamma = \langle I, (M_i), (v_i) \rangle$ for the strategic game with the default actions. The default action is a technical device to obtain a unique “lowest” game. Intuitively, a player should always have some action available even if this action just amounts to “do nothing”.

We call a subset L' of a lattice L a meet-sublattice if it is a lattice with respect to the order induced by L and the meet of any two elements of L' is the meet of the two elements in the lattice L . (Note that we do not require the join of any two elements of L' to be the join in L .)

Given the strategic game γ , we define a set of *restricted games* as follows: For nonempty M'_i with $M'_i \subseteq M_i$ and $d_i \in M'_i$, $M' := \prod_{i \in I} M'_i$, define a restricted strategic game $\langle I, (M'_i)_{i \in I}, (v_i|_{M'})_{i \in I} \rangle$ where $v_i|_{M'}$ is player i 's payoff function in γ restricted to outcomes in M' . Note that for each player the action set of a restricted game contains her default action. Let \perp denote the trivial game in which each player just has her default action. A set of restricted games derived from γ is rich, if for all $(m_i)_{i \in I} \in \tilde{M}$ the restricted games $\langle I, (\{m_i, d_i\})_{i \in I}, (v_i|_{\prod_{j \in I} \{m_j, d_j\}})_{i \in I} \rangle$, γ , and \perp belong to this set of restricted games. A partial order on the set of all restricted games is defined by set inclusion of $\prod_{i \in I} M_i$. Note that the set of all restricted games is a finite lattice. Thus, the meet is defined by the intersection of action sets. Let $G(\gamma)$ denote a rich meet-sublattice of restricted games derived from γ , and let \preceq denote the partial order on $G(\gamma)$.⁹ If $\alpha \in G(\gamma)$, we denote by $G(\alpha)$ the sublattice of strategic games $\beta \in G(\gamma)$ for which $\beta \preceq \alpha$.

For notational convenience, define $I^0 := I \cup \{0\}$. For each $\alpha \in G(\gamma)$, we define a space of states

$$S_\alpha = \{\omega : \omega = (\alpha_i)_{i \in I^0} \text{ with } \alpha_0 = \alpha \text{ and } \alpha_i \in G(\alpha) \text{ for every } i \in I\}.$$

Intuitively, a state in S_α shall describe which game each player perceives together with the index $\alpha_0 = \alpha$ for the state space S_α . In a state in S_α no player can perceive a game $\beta \in G(\gamma)$ that is “more expressive than or incomparable to” α . Let $\mathcal{S}_\gamma = \{S_\alpha\}_{\alpha \in G(\gamma)}$. By definition (by the first component of the states), these spaces are disjoint. The set \mathcal{S}_γ is a lattice of disjoint state-spaces, where the partial order is defined by extending \preceq to \mathcal{S}_γ by $S_\alpha \succeq S_\beta$, $S_\alpha, S_\beta \in \mathcal{S}_\gamma$ if and only if $\alpha \succeq \beta$, $\alpha, \beta \in G(\gamma)$. As before, the union of all spaces is denoted by Ω .

For all $\alpha, \beta \in G(\gamma)$ with $\alpha \succeq \beta$, the projections $r_\beta^\alpha : S_\alpha \rightarrow S_\beta$ are defined by the following rule: If $\omega = (\alpha_i)_{i \in I^0} \in S_\alpha$, then $r_\beta^\alpha(\omega) = (\beta_i)_{i \in I^0}$ with $\beta_i = \inf\{\alpha_i, \beta\}$ for all $i \in I^0$. The inf does always exist since $G(\gamma)$ is a lattice. The proof of the following remark is contained in the appendix.

Remark 4 For all $\alpha \in G(\gamma)$, $r_\alpha^\alpha = id_{S_\alpha}$. Projections commute, i.e., for any $\alpha, \beta, \delta \in G(\gamma)$ with $\alpha \succeq \beta \succeq \delta$, $r_\delta^\alpha = r_\delta^\beta \circ r_\beta^\alpha$.

Next, we define for each player $i \in I$ a sequence of type mappings $t_i^k : \Omega \rightarrow \bigcup_{\alpha \in G(\gamma)} \Delta(S_\alpha)$ satisfying the following properties: For each $k = 0, 1, \dots$,

- (0') If $\omega = (\alpha_j)_{j \in I^0} \in S_\alpha$ then $t_i^k(\omega) \in \Delta(S_{\alpha_i})$. (Note that $\alpha_i \preceq \alpha$.)
- (1') If $\omega = (\delta_j)_{j \in I^0} \in S_\alpha$, $t_i^k(\omega) \in \Delta(S_{\delta_i})$ then $t_i^k(\omega_\beta) = t_i^k(\omega)$ for all $\beta \in G(\gamma)$ s.t. $\alpha \succeq \beta \succeq \delta_i$.

⁹Although $G(\gamma)$ is a subset of the lattice of all restricted games given γ , and although the partial order on $G(\gamma)$ is the partial order of the lattice of all restricted games, $G(\gamma)$ is not necessarily a sublattice of the lattice of all restricted games. This is because for example $\alpha, \beta \in G(\gamma)$, $\alpha \vee \beta$, might not be in $G(\gamma)$, but only a game larger than $\alpha \vee \beta$. That is, the join of α and β in $G(\gamma)$ may not coincide with the join of α and β in the lattice of all restricted games given γ .

(2') If $\omega = (\beta_j)_{j \in I^0} \in S_\alpha$, $t_i^k(\omega) \in \Delta(S_{\beta_i})$, $\alpha \succeq \beta_i \succeq \delta$, then $t_i^k(\omega_\delta) = t_i^k(\omega)|_{S_\delta}$.

These properties imply properties (0) to (3) of type mappings in Definition 1 (see the appendix for a proof):

Lemma 1 *For each $k = 0, 1, \dots$, the type mapping satisfies the properties (0) to (3) of Definition 1. More specific, (0') implies (0), (1') implies (1), (2') implies (2), and (0') implies (3).*

Let $\bar{\alpha} := (\alpha_i)_{i \in I^0}$ with $\alpha_i = \alpha \in G(\gamma)$ for all $i \in I^0$. That is, a bar over α signifies a “monomorphic” state in which all components are identical to α . Note that by Property (0'), for all $\alpha \in G(\gamma)$, $t_i^k(\bar{\alpha}) \in \Delta(S_\alpha)$ for all $i \in I$ and $k = 0, 1, \dots$.

In the current context we impose additional properties. For each $i \in I$ and each $k = 0, 1, \dots$,

(4) If $\omega = (\alpha_i)_{i \in I^0}$ and $\omega' = (\beta_i)_{i \in I^0}$ with $\alpha_i = \beta_i$, then $t_i^k(\omega) = t_i^k(\omega')$.

That is, for each awareness level, each agent has only one type. This together with (2') implies that Introspection (Assumption 1) is satisfied.

(5) Uncertainty about opponents' awareness: If $t_i^k(\omega) \in \Delta(S_\alpha)$, then $\omega' = (\beta_j)_{j \in I^0} \in S_\alpha$ with $\beta_i = \alpha$ and $\perp \prec \beta_j \prec \alpha$, for all $j \in I \setminus \{i\}$ implies $t_i^k(\omega)(\{\omega'\}) > 0$.

(6) Certainty in the limit: For every $i \in I$, $\lim_{k \rightarrow \infty} t_i^k(\bar{\gamma})(\{\omega\})$ exists for all $\omega \in S_\gamma$, and $t_i^k(\bar{\gamma})(\{\bar{\gamma}\}) \rightarrow 1$ as $k \rightarrow \infty$.

Remark 5 *Properties (0') and (4) imply property (1').*

Property (4) states that if at two states player i has the same awareness, then she has the same beliefs at those two states. This assumption is not necessary but simplifies the analysis. The next two properties are crucial for the analysis. Uncertainty about opponents' awareness, Property (5), requires that player i considers it possible that any opponent is unaware of any nonempty subset of (non-default) actions. In particular it implies that for any (non-default) action and any opponent, a player does not exclude the possibility that the opponent is unaware of this action. Finally, Property (6), certainty in the limit, just means that in the limit at the objective true state $\bar{\gamma}$ there is common certainty that all players are aware of all actions. That is, in the limit the players at state $\bar{\gamma}$ are commonly certain that they are playing the original strategic game γ . Note that properties of the type mapping imply properties analogous to Property (6) for all “monomorphic” states $\bar{\alpha}$ for $\alpha \in G(\gamma)$.

For each $i \in I$, define a correspondence $\mathcal{M}_i : \Omega \rightarrow 2^{M_i}$ by $\mathcal{M}_i(\omega) = M_i^\alpha$ if $\omega \in S_\alpha$, where M_i^α is player i 's set of actions in the game $\alpha \in G(\gamma)$.

Finally, for each player $i \in I$, define a utility function $u_i : \bigcup_{\omega \in \Omega} \left(\prod_{j \in I} \mathcal{M}_j(\omega) \right) \times \{\omega\} \rightarrow \mathbb{R}$ in the following manner: For all $\alpha \in G(\gamma)$, if $\omega \in S_\alpha$ then $u_i(\cdot, \omega) = v_i|_{M^\alpha}$ where M^α is the action space in game α .

With these definitions we have (see the appendix for a proof):

Remark 6 *For each finite strategic game γ and each $k = 0, 1, \dots$ we have that*

$$\Gamma(\underline{\mathcal{S}}_\gamma^k) := \langle \mathcal{S}_\gamma, (r_\beta^\alpha)_{\alpha \succeq \beta; \alpha, \beta \in G(\gamma)}, (t_i^k)_{i \in I}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \rangle$$

is a finite Bayesian game with unawareness of events and actions.

Proposition 1 and Remark 6 imply:

Corollary 1 *For each finite strategic game γ and each $k = 0, 1, \dots$ there exists a Bayesian Nash equilibrium of $\Gamma(\underline{\mathcal{S}}_\gamma^k)$.*

Denote by $\sigma^k = ((\sigma_i^k(\omega))_{i \in I})_{\omega \in \Omega}$ a strategy combination in the game $\Gamma(\underline{\mathcal{S}}_\gamma^k)$.

The next lemma says that the limit at state $\bar{\gamma}$ of a convergent sequence of Bayesian Nash equilibria of games $\Gamma(\underline{\mathcal{S}}_\gamma^k)$ as $k \rightarrow \infty$ is an equilibrium of γ . (See the appendix for a proof.)

Lemma 2 *If $\sigma^k \in E(\Gamma(\underline{\mathcal{S}}_\gamma^k))$ and $\sigma^k(\bar{\gamma}) \rightarrow \nu$ as $k \rightarrow \infty$, then ν is an equilibrium of γ .*

Definition 12 (Unawareness Perfect Equilibrium) *An Unawareness Perfect Equilibrium ν of the strategic game γ is a Nash equilibrium of γ for which there exists a sequence of Bayesian games with unawareness $(\Gamma(\underline{\mathcal{S}}_\gamma^k))_{k=0}^\infty$ (as defined above, with certainty in the limit) with Bayesian Nash equilibria $\sigma^k \in E(\Gamma(\underline{\mathcal{S}}_\gamma^k))$, $k = 0, 1, \dots$, for which $\sigma^k(\bar{\gamma}) \rightarrow \nu$ as $k \rightarrow \infty$.*

That is, an Unawareness Perfect Equilibrium of a strategic game is a Nash equilibrium that is robust to uncertainty over unawareness of actions as this uncertainty vanishes.

Using Lemma 2 and the existence result for Bayesian Nash equilibrium (Proposition 1), we prove in the appendix the following:

Theorem 1 *For every finite strategic game, an Unawareness Perfect Equilibrium exists.*

It turns out that unawareness perfection is closely related to undominated actions. An action $m_i \in M_i$ is weakly dominated in the strategic game γ if there exists a mixed action $\mu_i \in \Delta(M_i)$ such that $v_i(m_i, m_{-i}) \leq v_i(\mu_i, m_{-i})$ for all $m_{-i} \in M_{-i}$ and $v_i(m_i, m_{-i}) < v_i(\mu_i, m_{-i})$ for some $m_{-i} \in M_{-i}$. This is the standard definition.¹⁰ An action that is not

¹⁰Note that the default action won't change the set of undominated actions since it is strictly dominated and payoff-neutral to opponents.

weakly dominated is undominated. A Nash equilibrium is undominated if every player's (mixed) equilibrium action assigns strict positive probability to undominated actions only. We prove in the appendix the following characterization:

Theorem 2 *A Nash equilibrium of a finite strategic game is an Unawareness Perfect Equilibrium if and only if it is undominated (i.e., not weakly dominated).*

That every Unawareness Perfect Equilibrium is undominated essentially follows from continuity of von Neumann-Morgenstern utilities and the fact that it is a limit of Bayesian Nash equilibria in which opponents' types may be forced to play any non-default action due to unawareness of alternative actions. Any action played with strict positive probability in an unawareness perfect equilibrium of a strategic game must also get strict positive probability in a Bayesian Nash equilibrium of $\Gamma(S_\gamma^k)$ at some point during the sequence for sufficiently large k . This implies that the action is a best reply against a full support belief over all action profiles that do not contain the default actions. Hence, by Pearce (1984, Lemma 4) the action is undominated in the original game. The proof of the converse is more involved as we construct a sequence of Bayesian games in which we mimic the full support belief that rationalizes a players undominated equilibrium strategy by beliefs about opponents' restricted awareness and their play restricted to actions that they are aware of. Any undominated Nash equilibrium is a best reply against the opponents' equilibrium actions and again by Pearce (1984, Lemma 4), a full support belief over opponents' actions. We show that for any undominated Nash equilibrium of a finite strategic game we can construct a sequence of Bayesian games with unawareness with a corresponding sequence of Bayesian Nash equilibria. Any fully aware player's Bayesian equilibrium action is a best reply to possibly unaware opponent's Bayesian equilibrium strategies and, in the limit, to the mixed strategies of the others in an equilibrium of the true strategic game.

A Connection to Standard Bayesian Games

In this section we compare standard Bayesian games to Bayesian games with unawareness. In particular, we show how to derive a standard type space with zero probability from our unawareness structure by “flattening” our lattice of spaces. “Flattening” the game is a purely technical procedure. While we can show a correspondence between equilibria in a Bayesian game with unawareness and equilibria in a standard Bayesian game, the equilibrium in the standard Bayesian game cannot be interpreted anymore under unawareness because the “language” required to identify events of which a player could be unaware is essentially “erased”. Since a flattened structure is a standard type-space, the “Dekel-Lipman-Modica-Rustichini critique” applies (Modica and Rustichini, 1994, Dekel, Lipman, and Rustichini, 1998). Hence unawareness is trivial in the flattened

game.¹¹

Definition 13 $G \subseteq \Omega$ is a measurable set if and only if for all $S \in \mathcal{S}$, $G \cap S \in \mathcal{F}_S$.

Note that a measurable subset of Ω is not necessarily an event in our special event structure.

Remark 7 The collection of measurable sets forms a sigma-algebra on Ω .

Remark 8 Let \mathcal{S} be at most countable and G be a measurable set, $p \in [0, 1]$ and $i \in I$. Then $\{\omega \in \Omega : t_i(\omega)(G) \geq p\}$ is a measurable set.

Let $\underline{\mathcal{S}}$ be an unawareness belief structure. We define the *flattened type-space* associated with the unawareness belief structure $\underline{\mathcal{S}}$ by

$$F(\underline{\mathcal{S}}) := \langle \Omega, \mathcal{F}, (t_i^F)_{i \in I} \rangle,$$

where Ω is the union of all state-spaces in the unawareness belief structure $\underline{\mathcal{S}}$, \mathcal{F} is the collection of all measurable sets in $\underline{\mathcal{S}}$, and $t_i^F : \Omega \rightarrow \Delta(\Omega, \mathcal{F})$ is defined by

$$t_i^F(\omega)(E) := \begin{cases} t_i(\omega)(E \cap S_{t_i(\omega)}) & \text{if } E \cap S_{t_i(\omega)} \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

A *standard type-space* on Y for the player set I is a tuple

$$\underline{Y} := \langle Y, \mathcal{F}_Y, (t_i)_{i \in I} \rangle,$$

where Y is a nonempty set, \mathcal{F}_Y is a sigma-field on Y , and for $i \in I$, t_i is a $\mathcal{F}_Y - \mathcal{F}_{\Delta(Y)}$ -measurable function from Y to $\Delta(Y, \mathcal{F}_Y)$, the space of countably additive probability measures on (Y, \mathcal{F}_Y) , such that for all $\omega \in Y$ and $E \in \mathcal{F}_Y$: $[t_i(\omega)] \subseteq E$ implies $t_i(\omega)(E) = 1$, where $[t_i(\omega)] := \{\omega' \in Y : t_i(\omega') = t_i(\omega)\}$.

Heifetz, Meier, and Schipper (2012), Proposition 8, showed that if $\underline{\mathcal{S}}$ is a unawareness belief structure, then $F(\underline{\mathcal{S}})$ is a standard type-space. Moreover, it has the following property: For every $p > 0$, measurable set $E \in \mathcal{F}$ and $i \in I$: $\{\omega \in \Omega : t_i(\omega)(E) \geq p\} = \{\omega \in \Omega : t_i^F(\omega)(E) \geq p\}$ (and hence $\{\omega \in \Omega : t_i(\omega)(E) > p\} = \{\omega \in \Omega : t_i^F(\omega)(E) > p\}$.)

A flattened unawareness structure is just a standard type-space. To derive such a type-space, one extends a player's type mapping by assigning probability zero to measurable sets that she is unaware of.

Heifetz, Meier, and Schipper (2012) also showed that the converse is not true. I.e., given a standard type space with zero probability, it is not always possible to find some unawareness structure with non-trivial unawareness whose flattening is this standard type space.

¹¹Being able to “flatten” a Bayesian game with unawareness into a standard Bayesian game and showing a correspondence of the equilibria does not imply that unawareness has no behavioral implications. Rather, “flattening” eliminates the primitives required to reveal those behavioral implications. For behavioral implications of unawareness in a decision theoretic framework, see Schipper (2012).

Definition 14 (Flattened Game) *Given a Bayesian game with unawareness of events and (possibly) actions $\Gamma(\underline{\mathcal{S}})$, we can associate a standard Bayesian game $F(\Gamma(\underline{\mathcal{S}}))$ played on a standard type-space (with possibly allowing for varying action sets of the players across different types) in the following manner:*

If $\Gamma(\underline{\mathcal{S}}) = \langle \underline{\mathcal{S}}, (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \rangle$, where $\underline{\mathcal{S}} = \langle \mathcal{S}, (r_{S_\beta}^{S_\alpha})_{S_\beta \preceq S_\alpha}, (t_i)_{i \in I} \rangle$ is a unawareness belief structure, then set $F(\Gamma(\underline{\mathcal{S}})) := \langle F(\underline{\mathcal{S}}), (M_i)_{i \in I}, (\mathcal{M}_i)_{i \in I}, (u_i)_{i \in I} \rangle$, where $F(\underline{\mathcal{S}})$ is the flattened structure associated with $\underline{\mathcal{S}}$, and $(M_i)_{i \in I}$, $(\mathcal{M}_i)_{i \in I}$, and $(u_i)_{i \in I}$ remain unchanged.

The flattened game is a standard Bayesian game.

Proposition 3 *Since the strategy sets and the utility functions remain unchanged, we have that any strategy profile is a Bayesian equilibrium in $\Gamma(\underline{\mathcal{S}})$ if and only if it is a Bayesian equilibrium in $F(\Gamma(\underline{\mathcal{S}}))$.*

The interpretation of a flattened game may be flawed in several ways. For instance, we can have types of players who are certain of their set of actions, but consider it possible that they have a larger set of actions even though they don't have a larger set of actions. This leads to serious conceptual problems, if a player were to choose such an action.¹²

B Proofs

B.1 Proof of Proposition 2

Let $\sigma_{|\Omega'}^*$ be an equilibrium in the S' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}')$. For $S'' \succeq S'$ we define a strategic form game with

- $I(\Omega'' \setminus \Omega') := \{(i, t_i(\omega)) : \omega \in \Omega'', i \in I(\omega)\} \setminus \{(i, t_i(\omega)) : \omega \in \Omega', i \in I(\omega)\}$ being the set of players,
- the set of strategies of player $(i, t_i(\omega)) \in I(\Omega'' \setminus \Omega')$ is $\Delta(\mathcal{M}_i(\omega_{S_{t_i(\omega)}}))$,¹³
- the payoff function of player $(i, t_i(\omega))$ is given by equation (2) but fixing the strategy of each (dummy) player in $\{(i, t_i(\omega')) : \omega' \in \Omega', i \in I(\omega')\}$ to her respective equilibrium strategy $\sigma_i^*(\omega)$ of the S' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}')$.

Since I , Ω , and $(M_i)_{i \in I}$ are finite, this strategic game has an equilibrium by Nash's (1950) theorem. Fix one equilibrium of this game.

¹²A player could then “test” his own beliefs by trying to choose such actions.

¹³Note that for $(i, t_i(\omega)) \in I(\Omega'' \setminus \Omega')$ we have either $S_{t_i(\omega)} \succ S'$ or $S_{t_i(\omega)}$ and S' are incomparable.

Consider now the strategy profile $\sigma_{|\Omega''}^*$ in which players in $\{(i, t_i(\omega)) : \omega \in \Omega', i \in I(\omega)\}$ play their component of the profile $\sigma_{|\Omega'}^*$ and players in $I(\Omega'' \setminus \Omega')$ play the equilibrium strategies of the equilibrium in the above defined strategic game.

We need to show that $\sigma_{|\Omega''}^*$ is an equilibrium of the S'' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}'')$. Suppose not, then for some player $(i, t_i(\omega)) \in I(\Omega'') = \{(i, t_i(\omega')) : \omega' \in \Omega'', i \in I(\omega')\}$ there exists $\sigma_i(\omega) \in \Delta(\mathcal{M}_i(\omega_{S_{t_i}(\omega)}))$ with $\sigma_i(\omega) \neq \sigma_i^*(\omega)$ such that for $\sigma_{S_{t_i}(\omega)} := (\sigma_i(\omega), (\sigma_j^*(\omega'))_{\omega' \in S_{t_i}(\omega), j \in I(\omega) \setminus \{i\}})$ we have

$$U_{(i, t_i(\omega))}(\sigma_{S_{t_i}(\omega)}) > U_{(i, t_i(\omega))}(\sigma_{S_{t_i}(\omega)}^*),$$

i.e., there exists a profitable deviation from $\sigma_{|\Omega''}^*$ for some player-type $(i, t_i(\omega))$ with $\omega \in \Omega''$ and $i \in I(\omega)$ given that all other player-types in $I(\Omega'')$ play their equilibrium strategy.

If $(i, t_i(\omega)) \in I(\Omega'' \setminus \Omega')$ then her strategy is not an equilibrium strategy in the above defined strategic game, a contradiction. If $(i, t_i(\omega)) \in \{(i, t_i(\omega')) : \omega' \in \Omega', i \in I(\omega')\}$, then since her payoffs are identical in both games, her strategy is not an equilibrium strategy in the S' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}')$, a contradiction. Hence $\sigma_{|\Omega''}^*$ must be an equilibrium of the S'' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}'')$. \square

B.2 Proof of Remark 3

Let $\sigma_{|\Omega''}^*$ be an equilibrium of the S'' -partial Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}'')$. Moreover, let $\sigma_{|\Omega'}^*$ be a profile of strategies that is identical with $\sigma_{|\Omega''}^*$ for all $(i, t_i(\omega)) \in I(\Omega')$.

Since for each player type $(i, t_i(\omega))$ with $i \in I(\omega)$ and $\omega \in \Omega'$ the payoff depends only (also in the S'' -partial game) on the choices of the player types $(j, t_j(\omega'))$ with $j \in I(\omega')$ and $\omega' \in \Omega'$. Fixing the behavior of the other player types, he faces the same maximization problem in the S' -partial game as in the S'' -partial game. \square

B.3 Proof of Remark 4

Let $\omega = (\alpha_0, (\alpha_i)_{i \in I})$ with $\alpha_0 = \alpha$, and let $\alpha \succeq \beta \succeq \delta$. We have to show for every $i \in I^0$ that $\inf\{\alpha_i, \delta\} = \inf\{\delta, \inf\{\alpha_i, \beta\}\}$. Since $\delta \preceq \beta$, we have $\inf\{\alpha_i, \delta\} \preceq \inf\{\alpha_i, \beta\}$ and hence $\inf\{\alpha_i, \delta\} \preceq \inf\{\delta, \inf\{\alpha_i, \beta\}\}$. We have $\inf\{\delta, \inf\{\alpha_i, \beta\}\} \preceq \delta$, by definition, but since $\inf\{\alpha_i, \beta\} \preceq \alpha_i$, we also have $\inf\{\delta, \inf\{\alpha_i, \beta\}\} \preceq \alpha_i$. This implies $\inf\{\alpha, \inf\{\alpha_i, \beta\}\} \preceq \inf\{\alpha_i, \delta\}$. \square

B.4 Proof of Lemma 1

We prove only the last claim, since the rest is obvious.

Let $\alpha \succeq \beta \succeq \delta$, $\omega = (\alpha_j)_{j \in I^0} \in S_\alpha$ and $t_i^k(\omega_\beta) \in \Delta(S_\delta)$. By definition $t_i^k(\omega) \in \Delta(S_{\alpha_i})$. We have to show that $\alpha_i \succeq \delta$. Since $t_i^k(\omega_\beta) \in \Delta(S_\delta)$, we have by (0') that $\omega_\beta = (\beta_j)_{j \in I^0}$ is such that $\beta_i = \delta$. By the definition of r_β^α , we have $\beta_i = \inf\{\alpha_i, \beta\}$. Hence $\delta \preceq \alpha_i$. \square

B.5 Proof of Remark 6

We define $\alpha \cap \beta$ to be the restricted game such that $M_i^{\alpha \cap \beta} := M_i^\alpha \cap M_i^\beta$, for all $i \in I$. Note that if $\alpha, \beta \in G(\gamma)$, then $\alpha \cap \beta = \alpha \wedge \beta$ since $G(\gamma)$ is a meet-sublattice of the lattice of all restricted games.

Let $\emptyset \neq M'_i \subseteq M_i$, where M_i is the action set of player i in γ . We have to show that $\{\omega \in \Omega : \mathcal{M}_i(\omega) \supseteq M'_i\}$ is an event.

By definition, $G(\gamma)$ is a finite meet-sublattice of the lattice of all restricted games given γ ordered by set inclusion of $\prod_{i \in I} M_i$. Fix a player $i \in I$. Recall that $\mathcal{M}_i(\omega) = M_i^\alpha$, for all $\omega \in S_\alpha$, where M_i^α is the action set of player i in the restricted game α .

Let $A = \{\alpha \in G(\gamma) : M_i^\alpha \supseteq M'_i\}$. Since $G(\gamma)$ is a finite meet-sublattice, $\bigcap_{\alpha \in A} \alpha =: \alpha(M'_i) \in G(\gamma)$. We have $M'_i \subseteq M_i^{\alpha(M'_i)}$ and $M_i^{\alpha(M'_i)} \subseteq \mathcal{M}_i(\omega)$, for all ω such that $\mathcal{M}_i(\omega) \supseteq M'_i$. Since for $\omega \in S_\alpha$, $\mathcal{M}_i(\omega) = M_i^\alpha$, we have that $[M'_i] = \{\omega \in \Omega : \mathcal{M}_i(\omega) \supseteq M'_i\} = (S_{\alpha(M'_i)})^\uparrow$, which is an event. \square

B.6 Proof of Lemma 2

Suppose for some $i \in I$, ν_i would not be a best reply to ν_{-i} . Then there exists $m_i \in M_i$ such that

$$\varepsilon \leq v_i(m_i, \nu_{-i}) - v_i(\nu_i, \nu_{-i}),$$

for some $\varepsilon > 0$. By continuity of the utility functions in mixed strategies and in beliefs on types there exists a k_1 such that

$$\left| v_i(\nu_i, \nu_{-i}) - U_{(i, t_i^k(\bar{\gamma}))}(\sigma_i^k, \sigma_{-i}^k) \right| < \frac{\varepsilon}{3},$$

for all $k \geq k_1$. And likewise there exists k_2 such that

$$\left| U_{(i, t_i^k(\bar{\gamma}))}(m_i, \sigma_{-i}^k) - v_i(m_i, \nu_{-i}) \right| < \frac{\varepsilon}{3},$$

for all $k \geq k_2$.

Let $k \geq \max\{k_1, k_2\}$, then

$$\begin{aligned} & U_{(i, t_i^k(\bar{\gamma}))}(m_i, \sigma_{-i}^k) - U_{(i, t_i^k(\bar{\gamma}))}(\sigma_i^k, \sigma_{-i}^k) \\ &= \underbrace{U_{(i, t_i^k(\bar{\gamma}))}(m_i, \sigma_{-i}^k) - v_i(m_i, \nu_{-i})}_{\geq -\frac{\varepsilon}{3}} + \underbrace{v_i(m_i, \nu_{-i}) - v_i(\nu_i, \nu_{-i})}_{\geq \varepsilon} \end{aligned}$$

$$\begin{aligned}
& + \underbrace{v_i(\nu_i, \nu_{-i}) - U_{(i, t_i^k(\bar{\gamma}))}(\sigma_i^k, \sigma_{-i}^k)}_{\geq -\frac{\varepsilon}{3}} \\
& \geq \frac{\varepsilon}{3},
\end{aligned}$$

that is, σ_i^k is not a best reply to σ_{-i}^k , for sufficiently large k . This implies that σ^k is not an equilibrium of the game $\Gamma(\underline{\mathcal{S}}_\gamma^k)$, for sufficiently large k , a contradiction. \square

B.7 Proof of Theorem 1

For every $k = 0, 1, \dots$, the Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}_\gamma^k)$ has an equilibrium by Corollary 1. Since the set of mixed strategy combinations at $\bar{\gamma}$ is a closed and bounded subset of an Euclidean space, the sequence of equilibria $(\sigma^k(\bar{\gamma}))_{k=0}^\infty$ has a subsequence that converges to some ν . By Lemma 2, ν is an unawareness perfect equilibrium of γ . \square

B.8 Proof of Theorem 2

“ \Rightarrow ”: Let ν be an unawareness perfect equilibrium of γ . Then there exists a sequence of Bayesian games with unawareness $(\Gamma(\underline{\mathcal{S}}_\gamma^k))_{k=0}^\infty$ with corresponding Bayesian Nash equilibria $\sigma^k \in E(\Gamma(\underline{\mathcal{S}}_\gamma^k))$, $k = 0, 1, \dots$, for which $\sigma^k(\bar{\gamma}) \rightarrow \nu$ as $k \rightarrow \infty$.

From Pearce (1984, Lemma 4) follows that a Nash equilibrium profile ν is undominated if and only if for every $i \in N$, ν_i is a best response to a completely mixed strategy profile of opponents. Note that $\sum_{\omega \in S_\gamma} t_i^k(\bar{\gamma})(\{\omega\})\sigma_{-i}^k(\omega)$ is equivalent to a completely mixed strategy profile of opponents for any $k = 0, 1, \dots$, in the game γ . Since $\sigma_i^k(\bar{\gamma}) \rightarrow \nu_i$, for any $m_i \in M_i$ with $\nu_i(m_i) > 0$, there exists a sufficiently large $k(m_i)$ such that for all $k \geq k(m_i)$, $\sigma_i^k(\bar{\gamma})(m_i) > 0$. Since M_i is finite, there is a k_{\max} such that $\sigma_i^k(\bar{\gamma})(m_i) > 0$ for all $m_i \in M_i$ with $\nu_i(m_i) > 0$. Thus, any such $m_i \in M_i$ is a best reply to $\sum_{\omega \in S_\gamma} t_i^k(\bar{\gamma})(\{\omega\})\sigma_{-i}^k(\omega)$, for all $k \geq k_{\max}$. Hence, ν_i is a best reply to the completely mixed belief equivalent to $\sum_{\omega \in S_\gamma} t_i^{k_{\max}}(\bar{\gamma})(\{\omega\})\sigma_{-i}^{k_{\max}}(\omega)$.

“ \Leftarrow ”: Recall that from Pearce (1984, Lemma 4) it follows that a Nash equilibrium profile ν is undominated if and only if for every $i \in N$, ν_i is a best response to a completely mixed strategy profile of opponents $\tilde{\nu}_{-i}$. We will use these completely mixed strategy profiles $\tilde{\nu}_{-i}$, $i \in I$, to construct a sequence of Bayesian games with unawareness $(\Gamma(\underline{\mathcal{S}}_\gamma^k))_{k=0}^\infty$ with Bayesian Nash equilibria $\sigma^k \in E(\Gamma(\underline{\mathcal{S}}_\gamma^k))$, $k = 0, 1, \dots$, for which $\sigma^k(\bar{\gamma}) \rightarrow \nu$ as $k \rightarrow \infty$.

Given the strategic game γ , consider any restricted strategic game α with $d_i \in M_i^\alpha$ for all $i \in I$ and $1 \leq |M_i^\alpha| \leq 2$. Note that by construction in any such restricted game, each player has at most one non-default action and if there is a non-default action for player i , then this non-default action is the strict dominant Nash equilibrium action of α . Let $G(\gamma)$ be the set comprising of all such games, \perp , and γ itself. Note that $G(\gamma)$ is rich (see page 23) and a lattice.

We now construct a sequence of Bayesian games with unawareness of actions by defining for each player $i \in I$ the type mapping as follows: For $m_{-i} \in \tilde{M}_{-i}$, define $\ell(m_{-i})$ to be the number of profiles $(\alpha_j)_{j \in I \setminus \{i\}}$ with $\alpha_j \in G(\gamma) \setminus \{\perp, \gamma\}$, $j \in I \setminus \{i\}$, for which $\{m_{-i}\} = \prod_{j \in I \setminus \{i\}} (M_j^{\alpha_j} \setminus \{d_j\})$. (Note that in such a profile we have $|M_j^{\alpha_j}| = 2$, for all $j \in I \setminus \{i\}$). Let ℓ_0 be the number of profiles $(\beta_j)_{j \in I \setminus \{i\}}$ such that $M_j^{\beta_j} = \{d_j\}$ for at least one $j \in I \setminus \{i\}$.

Fix $\varepsilon \in (0, 1)$. If $\omega = (\alpha_j)_{j \in I^0}$ with $\alpha_j \in G(\gamma) \setminus \{\perp, \gamma\}$ such that $|M_j^{\alpha_j}| = 2$, for $j \in I \setminus \{i\}$ and $\alpha_j = \gamma$ for $j \in \{0, i\}$, let $t_i^k(\bar{\gamma})(\{\omega\}) = \frac{\varepsilon^k}{\ell(m_{-i})} \tilde{\nu}_{-i}(m_{-i})$ with $\{m_{-i}\} = \prod_{j \in I \setminus \{i\}} (M_j^{\alpha_j} \setminus \{d_j\})$.

If $\omega = (\beta_j)_{j \in I^0}$ such that $\beta_0 = \beta_i = \gamma$ and $M_j^{\beta_j} = \{d_j\}$ for at least one $j \in I \setminus \{i\}$, define $t_i^k(\bar{\gamma})(\{\omega\}) = \frac{\varepsilon^{2k}}{\ell_0}$. Note that in such a state player j is only aware of d_j , which forces him to play d_j . From the way we extended the players' utility functions to the default auctions, this makes all other players indifferent between all their non-default actions. Hence, for equilibrium considerations, for the type of each player in state $\bar{\gamma}$ these states play no role, even if he assigns positive probability to them.

All the remaining probability mass, $1 - \varepsilon^k - \varepsilon^{2k}$, is assigned to $t_i^k(\bar{\gamma})(\{\bar{\gamma}\})$. We impose properties (0'), (1'), (2'), and (4). Property (3) is implied by Lemma 1. Properties (5) and (6) are satisfied by construction. Beliefs of player i in states different from $\bar{\gamma}$ are then completely determined by beliefs in $\bar{\gamma}$ via properties (2') and (4).

Next, we construct a sequence of Bayesian Nash equilibria whose limit in state $\bar{\gamma}$ is the undominated Nash equilibrium ν . For any player i and any k , let $\sigma_i^k(\omega) = \nu_i$ for all $\omega \in S_\gamma$ with $\omega = (\beta_j)_{j \in I^0}$ such that $\beta_i = \gamma$. Moreover, set $\sigma_i^k(\perp)(d_i) = 1$, and $\sigma_i^k(\omega)(m_i) = 1$ with $m_i \neq d_i$ in any other state. (Recall that in such a state player i just has one non-default action.) Since in the latter states, $\sigma_i^k(\omega) \in \Delta(M_i^\alpha)$, for $\alpha \prec \gamma$, we extend $\sigma_i^k(\omega)$ to $\tilde{\sigma}_i^k(\omega) \in \Delta(M_i)$ by setting $\tilde{\sigma}_i^k(\omega)(m_i) := \sigma_i^k(\omega)(m_i)$ for all $\omega \in \Omega_\gamma$ and $m_i \in M_i^\alpha$.

Note that $\sum_{\omega \in S_\gamma} t_i^k(\bar{\gamma})(\{\omega\}) \tilde{\sigma}_{-i}^k(\omega) \rightarrow \nu_{-i}$ for $k \rightarrow \infty$. Moreover, for any k , $\sum_{\omega \in S_\gamma} t_i^k(\bar{\gamma})(\{\omega\}) \tilde{\sigma}_{-i}^k(\omega)$ is a convex combination of $\tilde{\nu}_{-i}$ and ν_{-i} (and mixtures where at least one player chooses his default action, but this can be disregarded).

To see that σ^k is a Bayesian Nash equilibrium for any $k = 0, 1, \dots$, note that $\sigma_i^k(\omega) = \nu_i$ with $\omega = (\beta_j)_{j \in I^0}$ such that $\beta_i = \gamma$ is a best reply to ν_{-i} since ν is Nash equilibrium. Since ν is an undominated Nash equilibrium, we also noted above that ν_i is a best reply against $\tilde{\nu}_{-i}$. Hence ν_i is also a best reply to any convex combination of ν_{-i} and $\tilde{\nu}_{-i}$. It follows that for any $k = 0, 1, \dots$, σ_i^k is a Bayesian Nash equilibrium mixture of player i in the Bayesian game with unawareness $\Gamma(\underline{\mathcal{S}}_\gamma^k)$. Since $\sigma_i^k(\bar{\gamma}) = \nu_i$ for all k and $i \in I$, ν is an Unawareness perfect equilibrium of γ . \square

References

- [1] Čopič, J. and A. Galeotti (2006). Awareness as an equilibrium notion: Normal-form games, mimeo.
- [2] Dekel, E., Lipman, B., and A. Rustichini (1998). Standard state-space models preclude unawareness, *Econometrica* **66**, 159–173.
- [3] Feinberg, Y. (2009). Games with unawareness, mimeo.
- [4] Feinberg, Y. (2005). Games with incomplete awareness, mimeo.
- [5] Halpern, J. and L. Rêgo (2008). Interactive unawareness revisited, *Games and Economic Behavior* **62**, 232–262.
- [6] Halpern, J. and L. Rêgo (2006). Extensive games with possibly unaware players, in: *Proc. Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, 744–751.
- [7] Harsanyi, J. (1967/68). Games with incomplete information played ‘Bayesian’ players, Part I, II, and II, *Management Science* **14**, 159-182, 320-334, 486–502.
- [8] Heifetz, A., Meier, M., and B.C. Schipper (2012). Unawareness, beliefs, and speculative trade, mimeo.
- [9] Heifetz, A., Meier, M., and B.C. Schipper (2011a). Dynamic unawareness and rationalizable behavior, mimeo.
- [10] Heifetz, A., Meier, M., and B.C. Schipper (2011b). Prudent rationalizability in generalized extensive-form games, mimeo.
- [11] Heifetz, A., Meier, M., and B.C. Schipper (2008). A canonical model for interactive unawareness, *Games and Economic Behavior* **62**, 304–324.
- [12] Heifetz, A., Meier, M., and B.C. Schipper (2006). Interactive unawareness, *Journal of Economic Theory* **130**, 78–94.
- [13] Heinsalu, S. (2011). Universal type spaces with unawareness, Yale University.
- [14] Hu H. and H. W. Stuart, Jr. (2001). An epistemic analysis of the Harsanyi transformation, *International Journal of Game Theory* **30**, 517–525.
- [15] Li, J. (2006b). Dynamic games with perfect awareness information, mimeo.
- [16] Meier, M., and B.C. Schipper (2012). Conditional dominance in games with unawareness, mimeo.
- [17] Meier, M. and B.C. Schipper (2009). Speculative trade under unawareness - The infinite case, mimeo.

- [18] Mertens, J.-F. and S. Zamir (1985). Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory* **14**, 1–29.
- [19] Modica, S. and A. Rustichini (1999). Unawareness and partitional information structures, *Games and Economic Behavior* **27**, 265–298.
- [20] Modica, S. and A. Rustichini (1994). Awareness and partitional information structures, *Theory and Decision* **37**, 107–124.
- [21] Nash, J. (1950). Equilibrium points in n-person games, *Proceedings of the National Academy of Sciences* **36**, 48–49.
- [22] Ozbay, E. (2007). Unawareness and strategic announcements in games with uncertainty, in: Samet, D. (ed.), *Proceedings of the 11th conference on Theoretical Aspects of Rationality and Knowledge*, Presses Universitaires de Louvain, pp. 231–238.
- [23] Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52**, 1029–1050.
- [24] Rêgo, L. and J. Halpern (2012). Generalized Solution Concepts in Games with Possibly Unaware Players, *International Journal of Game Theory* **41**, 131–155.
- [25] Sadzik, T. (2006). Knowledge, awareness and probabilistic beliefs, mimeo.
- [26] Schipper, B.C. (2012). Awareness-dependent subjective expected utility, *International Journal of Game Theory*, forthcoming.
- [27] Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory* **4**, 25–55.
- [28] Sims, C. (2010). Rational inattention and monetary policy, in: Friedman, B.H. and F. H. Hahn (Eds.), *Handbook of Monetary Policy*, Elsevier, 155–181.
- [29] Thaler, R.H. and C. R. Sunstein (2008). *Nudge: Improving decisions about health, wealth, and happiness*, Yale University Press.
- [30] Van Damme, E. (1991). *Stability and perfection of Nash equilibria*, Berlin et al.: Springer-Verlag.
- [31] Van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and under-diversification, *Review of Economic Studies* **77**, 779–805.