AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Risk & Sustainable Management Group

Risk & Uncertainty Program Working Paper: R09#1

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Inductive reasoning about unawareness

Simon Grant

Rice University

and

John Quiggin

Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
<http://www.uq.edu.au/economics/rsmg>

Inductive reasoning about unawareness*

Abstract

We develop a model of games with awareness that allows for differential levels of awareness. We show that, for the standard modal logical interpretations of belief and awareness, a player cannot believe there exist propositions of which he is unaware. Nevertheless, we argue that a boundedly rational individual may regard the possibility that there exist propositions of which she is unaware as being supported by inductive reasoning, based on past experience and consideration of the limited awareness of others. In this paper, we provide a formal representation of inductive reasoning in the context of a dynamic game with awareness. We show that, given differential awareness over time and between players, individuals can derive inductive support for propositions expressing their own unawareness.

JEL Classification: D80, D82

Key words: unawareness, bounded rationality, induction

Simon Grant
Department of Economics, Rice University, Houston

John Quiggin
School of Economics, University of Queensland, Brisbane

*We thank, without implication, Joseph Halpern, Jeff Kline, Mamuro Kaneko, Benoît Menoni, Leandro Rêgo, and Matthew Ryan for helpful comments and criticism. The usual disclaimer applies.

1 Introduction

A number of recent papers (Dekel, Lipman and Rustichini 1998; Feinberg 2005; Halpern and Rego 2006a;b, 2007; Heifetz, Meier and Schipper 2006, 2008) have examined decision problems and strategic interactions in which boundedly rational participants are not aware of all relevant possibilities.¹ The fact that individuals are boundedly rational does not preclude sophisticated reasoning. In particular, it is reasonable to suppose that, in many situations, people are conscious of the fact that there may exist possibilities of which they are not aware, and may act on this understanding.

Although there is now a large literature on decision making under conditions of limited awareness, there is no agreement on how, if at all, propositions like ‘there exist propositions of which I am currently unaware’ can be justified and used as a basis for decisions. In particular, for any given definitions of knowledge and awareness, it is natural to ask whether an individual can know that there exist propositions of which he is unaware (Halpern and Rego 2006a).

In this paper, we show that, for the standard modal logical interpretations of knowledge and awareness, the answer to this question must be negative. Nevertheless, we argue that a boundedly rational individual may regard the possibility that there exist propositions of which she is unaware as being supported by inductive reasoning, based on past experience and consideration of the limited awareness of others. The aim of this paper is to give some substance to a notion of inductive reasoning consistent with bounded rationality and limited awareness.

The paper is organized as follows:

Section 2 presents a semantic representation of differential awareness in terms of extensive games, following the approach of Osborne and Rubinstein (1994). The crucial idea is that of a model, recursively defined in terms of a game for which at each history, the players are imputed a model representing their own awareness of the game. It is shown that this recursion is finite, and terminates in a representation, referred to as a game of common awareness, in which all players are imputed full awareness of the game in question. We show that the standard concept of sequential equilibrium may be extended to the case of games with differential awareness, and demonstrate the existence of a sequential equilibrium for the class of games modelled here. These ideas are illustrated with reference to an example first presented by Heifetz, Meier and Schipper (2006).

Section 3 presents a syntactic rendition of the same ideas, associating with each model a propositional language rich enough to specify all histories and sets of histories that arise in the game associated with the given model. Within any such model, it is possible to define belief and awareness operators using a standard modal-logical approach, in which a proposition is believed

¹ A more extensive bibliography is maintained at <http://www.econ.ucdavis.edu/faculty/schipper/unaw.htm>

by a player if it is true in all histories considered possible by that player.² We show that, the model is one in which individuals can hold false beliefs and (correctly or otherwise) impute false beliefs to others.

In Section 4 the language is extended to include existential propositions of the general form ‘there exists a proposition q , such that ...’. This development enables us to consider the process of reasoning about awareness and unawareness. We show that, in the framework developed in Section 3, and with the standard modal logical definitions of belief and awareness defined there, an individual cannot believe that there exist propositions of which he is unaware. Nevertheless, this proposition can be formulated in the richer language we consider. Moreover, in the context of games of differential awareness, it will be true, in general, that there exist propositions of which individuals are aware. The critical question, then, is how this fact may be incorporated in the reasoning of a boundedly rational, but sophisticated, individual.

In Section 5, we offer an answer to this question based on the concept of inductive support. Inductive support may be derived from past experience or from observation of others. We say that a proposition is supported by historical induction if it has been (believed) true in the past, and never been (believed) false. In particular, since everyone has the experience of becoming aware of propositions and possibilities they have not previously considered, the proposition that they will continue to do so is supported by historical induction. Similarly, a proposition which holds true for at least some individuals, and is not false for any individual is supported by induction over individuals. In a game of differential awareness, individuals believe that others are unaware of at least some propositions. Inductive reasoning suggests that the same will be true of all individuals, including the person undertaking the induction. The main results of the paper are formal statements of these arguments. We argue that decisionmakers may reasonably choose strategies subject to heuristic constraints that rule out actions if the proposition that these actions will have unforeseen consequences is supported by induction. We develop this point to derive a no-trade result for the speculative trade example developed previously.

In Section 6, we examine the relationship of this paper with previous work, and discuss possible applications to the precautionary principle often advocated as a basis for regulatory decisions regarding environmental risks. Finally, we offer some concluding comments.

² The operator defined here is more usually referred to as describing knowledge rather than belief. This is appropriate in a game of complete awareness, where there can be no false knowledge. However, situations of differential awareness allow the possibility of false belief. There may be histories not considered by the individual in question, in which the proposition in question is false.

2 Games and strategies: the semantic representation

2.1 Extensive games and restrictions

We follow Osborne and Rubinstein (1994) and Halpern and Rego (2006b) in modeling an extensive game.

Definition 1 A finite extensive game is $\Gamma = (N, A, H, P, f_c, \{\mathcal{I}^i : i \in N\}, \{v^i : i \in N\})$ where:

G1 (Player Set): $N = \{c, 1, \dots, n\}$ is a finite set of players, where c denotes the chance player and $n \geq 1$;

G2 (Actions): A is a finite non-empty set of actions;

G3 (Histories): H is a finite non-empty subset of $\langle \rangle \cup \bigcup_{m=1}^{\infty} A^m$ where $\langle \rangle$ is the null sequence and is viewed as a sub-history of every history. Moreover, if $\langle \alpha_1, \alpha_2, \dots, \alpha_k \rangle \in H$, then for each $l \leq k$ the sub-history $\langle \alpha_1, \dots, \alpha_l \rangle$ is also in H . We denote the sub-history relationship $\langle \alpha_1, \dots, \alpha_l \rangle \preceq \langle \alpha_1, \alpha_2, \dots, \alpha_k \rangle$. For a sequence $h = \langle \alpha_1, \dots, \alpha_k \rangle$, and an action α , we use $h \cdot \langle \alpha \rangle$ to denote the extended sequence $\langle \alpha_1, \dots, \alpha_k, \alpha \rangle$. We then can define the set of available actions $A_h = \{\alpha \in A : h \cdot \langle \alpha \rangle \in H\}$. We partition the set of histories H into the set of decision histories D and the set of terminal histories Z , where $h \in Z$ if and only if $A_h = \emptyset$

G4 (Player Function): $P : H \rightarrow N$ assigns to each history a player making a decision after that history;

G5 (Chance Assignment): f_c associates with every history h such that $P(h) = c$ defines a probability distribution over A_h ;

G6 (Information Partitions): For each player $i \in N - \{c\}$, $\mathcal{I}^i : H \rightarrow 2^H$ is a partition of H , such that

(a) $h \in \mathcal{I}^i(h)$

(b) If $P(h) = i$ and $\tilde{h} \in \mathcal{I}^i(h)$, then $P(\tilde{h}) = i$ and $A_h = A_{\tilde{h}}$

(c) If $\tilde{h} \notin \mathcal{I}^i(h)$, then for any $\tilde{h}' \succeq \tilde{h}$, $h' \succeq h$, $\tilde{h}' \notin \mathcal{I}^i(h')$ (Perfect recall)

G7 (Payoffs): For each player $i \in N - \{c\}$, $v^i : Z \rightarrow \mathbb{R}$ is the payoff function for player i , representing expected-utility preferences for lotteries over Z

With boundedly rational players, it is necessary to consider the possibility that, at some given history h , player $i = P(h)$, who must choose an action, may not be aware of all possible histories in the game. For example, player i may be unaware of possible future moves available to other players, to the chance player, or to herself. Player i may even be unaware of the existence of some other players. We formalize this notion as follows.

Definition 2 A game $\tilde{\Gamma} = (\tilde{N}, \tilde{A}, \tilde{H}, \tilde{P}, \tilde{f}^c, \{\tilde{\mathcal{I}}^i : i \in \tilde{N}\}, \{\tilde{v}^i : i \in \tilde{N}\})$ is a restriction of Γ , denoted $\tilde{\Gamma} \sqsubseteq \Gamma$, if

R1 $\tilde{N} \subseteq N$

R2 $\tilde{A} \subseteq A$

R3 $\tilde{H} \subseteq H$ such that $\tilde{Z} \subseteq Z$

R4 $\tilde{P} : \tilde{H} \rightarrow \tilde{N}$ is the restriction of P to $\tilde{H}$

R5 For any $h \in \tilde{H}$ such that $P(h) = c$, $\tilde{f}^c(\alpha) \geq \tilde{f}^c(\alpha)$ for all $\alpha \in \tilde{A}_h (\subseteq A_h)$.

R6 For any $h \in \tilde{H}$, $\tilde{\mathcal{I}}^i(h) = \tilde{H} \cap \mathcal{I}^i(h)$

R7 $\tilde{v}^i$ is the restriction of v^i to $\tilde{Z}$.

R1 and R2 are implied by the discussion above. R3 ensures that subhistories and the subhistory ordering are preserved in the sense that, for any $h, \tilde{h} \in H$, such that $\tilde{h} \preceq_{\Gamma} h$, $h \in \tilde{H} \implies \tilde{h} \in \tilde{H}$ and $\tilde{h} \preceq_{\tilde{\Gamma}} h$. Further no new terminal nodes are created in the restricted game allowing us to obtain the payoff function from the restriction of v^i to $\tilde{Z}$, which is imposed by R7. R4 states that the identity of the player to move is unchanged. R5 ensures that, if an action by nature is excluded from consideration at h , the probability associated with that action is distributed over the remaining actions $\alpha \in A'_h$. R6 ensures that restriction does not add information, or lose information with respect to histories of which the player is aware

An alternative approach, adopted by Grant and Quiggin (2006) is to consider bounded rationality as a generating a ‘coarsening’ of the game structure. That is, one or more histories in the complete game are mapped to a single history in the coarsened game. Both approaches have some appeal as representations of bounded rationality. However, the restriction approach, in which histories of which players are unaware are deleted from their representation, is more tractable.

2.2 Strategies

As in Osborne and Rubinstein, a pure strategy for player i assigns to each h such that $P(h) = i$, and the associated information set $\mathcal{I}^i(h)$, an element $a^i \in A_h$. A behavioral strategy β^i for player i is a collection of independent probability measures $(\beta_i(\mathcal{I}^i(h)) : P(h) = i)$ where the support of each $\beta^i(\mathcal{I}^i(h))$ is a subset of A_h . That is, there is an independent probability measure specified for each information set controlled by player i . A behavioral strategy profile β is a set of behavioral strategies, one for each i .

A continuation behavioral strategy β_h^i for player i at h is the restriction of a behavioral strategy to the set $\{h' \in H : h' \succeq h, \text{ for some } h' \in \mathcal{I}^i(h)\}$, that is, the set of histories in H that pass through the information set $\mathcal{I}^i(h)$. A continuation behavioral strategy profile β_h is a set of continuation strategies at h , one for each i . The set of all continuation behavioral strategies for i at h is denoted $\mathcal{B}_h^i$ and the set of all continuation behavioral strategy profiles at h is $\mathcal{B}_h$.

2.3 Models and games with differential awareness

In a game where some players have bounded rationality, other players will, in general, be aware of this possibility. Hence, the standard assumption of common knowledge must be replaced with a structure in which each player imputes to the others a level of awareness which may change over a given history. We formulate this using the concept of a model.

Definition 3 *A game with differential awareness $\mathcal{G}$ is a finite partially ordered collection of models, with unique maximal element $\mathbf{M}$, where models are defined recursively as follows:*

A model $M = (\Gamma_M, \rho_M)$ consists of

1. *a game $\Gamma_M = (N_M, A_M, H_M, P_M, f_M^c, \{\mathcal{I}_M^i : i \in N_M\}, \{v_M^i : i \in N_M\})$;*
2. *a perception mapping $\rho_M : H_M \times N_M \rightarrow \mathcal{G}$ which maps each history $h \in H_M$ and player $j \in N_M$ in the given model M to another model $\tilde{M} = \rho_M(h, j)$; such that*

$$P1 \ \Gamma_{\tilde{M}} \sqsubseteq \Gamma_M$$

$$P2 \ \text{For any } h' \in \mathcal{I}_M^i(h), \ \rho_M(h, i) = M$$

$$P3 \ \text{If } h' \in \mathcal{I}_M^j(h), \text{ then } \rho_M(h', j) = \rho_M(h, j)$$

$$P4 \ \text{For any } h', h'' \in H_M, \text{ if } h' \preceq h'', \text{ then } \Gamma_{\rho_M(h', j)} \sqsubseteq \Gamma_{\rho_M(h'', j)} \text{ for all } j \in N_M$$

$$P5 \ \text{For all } \tilde{h} \in H_{\tilde{M}}, \ h' \in \mathcal{I}_M^j(h) \cap H_{\tilde{M}} \text{ if } \tilde{h} \preceq h', \text{ then } \rho_{\tilde{M}}(\tilde{h}, j) = \rho_M(h, j).$$

The properties of the perception mapping $\rho_M(\cdot, \cdot)$ are as follows:

P1 states that the game in each model is a restriction of the game in the model from which it is derived. This property implies that the recursive construction of the game with differential awareness as a sequence of models must eventually terminate in a model of common awareness (defined in more detail below).

P2 states that for the case $j = i$, player i imputes her own model to herself at h .

P3 states that the model is the same for all elements of an information set.

P4 states that awareness cannot decrease over time

P5 is consistent imputation. If, at h , the model M (for i) imputes to player j a model $\tilde{M}$, and an information set $\mathcal{I}_{\tilde{M}}^j = \mathcal{I}_M^j(h) \cap H_{\tilde{M}}$, then for any subhistory h'' of a history h' in $\mathcal{I}_{\tilde{M}}^j(\tilde{h})$, the model imputed by $\tilde{M}$ to j at h'' must be the same as the model imputed by M to j at h' . That is, informally speaking, the model i thinks j held at h'' must be the same as the model i thinks j thinks j held at h'' .

With a slight abuse of notation, we define the ordering $\sqsubseteq$ on $\mathcal{G}$ as the transitive closure of the relation generated by the requirement $M' \sqsubseteq M$ if $M' = \rho_M(h, j)$, for some $h \in \Gamma_M$.

Since the game with differential awareness $\mathcal{G}$ is a finite set of models, any maximal chain in $\mathcal{G}$ must have both a minimal element under $\sqsubseteq$, described below as a game of common awareness, and a maximal element under $\sqsubseteq$, namely $\mathbf{M} = (\mathbf{\Gamma}, \boldsymbol{\rho})$.

The recursive construction of the game with differential awareness may be described as follows. The maximal model $\mathbf{M}$ may be regarded as representing the perspective of an unboundedly rational external observer.

For each history $h \in \mathbf{H}$, and player i , the observer imputes to player i a perceived model of the game $\boldsymbol{\rho}(h, i) = M = (\Gamma_M, \rho_M)$. Since players are always assumed boundedly rational, we require $\boldsymbol{\rho}(h, i) \sqsubset \mathbf{M}$ for all h, i . That is, the subjective model $\boldsymbol{\rho}(h, i)$ available to player h at history i is always considered distinct from the maximal model $\mathbf{M}$.

The model M , in turn, consists of the game $\Gamma_M \sqsubseteq \mathbf{\Gamma}$ and a perception mapping ρ_M representing i 's beliefs about the awareness of all the players of whom she is aware (including herself).³ For each history $h \in H_M$, this mapping imputes to each player $j \in N_M$ a perceived model.

The recursive nature of the definition reflects the fact that players may impute limited awareness of the game to other players. The interpretation is that, at history h , the model $\rho_M(h, j)$ is the perception of the game with differential awareness held by the player j . The restriction $\Gamma_{M'} \sqsubseteq \Gamma_M$ reflects the fact that, in the structure considered here, players can only impute to other players an awareness that is equal to, or a restriction of, their own.

For each model M , the game Γ_M is maximal under $\sqsubseteq$ with respect to the set $\{\Gamma_{M'} : M' \sqsubseteq M\}$ that is, the set of games in models that may be imputed to some player under M . Thus, in modelling the game, each player puts themselves in the position of a maximally aware external observer.

We have:

Lemma 1 *Let $\boldsymbol{\rho}(h, i) = M$. For any $\tilde{h} \in \mathcal{I}_M^i(h)$ and any $h' \in H_M$ such that $h' \succeq \tilde{h}$, we have $\rho_M(h', i) = M$.*

Proof. From P2, $\rho_M(\tilde{h}, i) = M$. From P1 and P5, $\rho_M(h', i) = \rho_M(\tilde{h}, i) = M$ ■

That is, within the model M , players do not consider the possibility that they may become more aware in the future, even though, by P4, the model captures the fact that awareness, in general, increases over time.

Lemma 2 *Fix M , and let $\tilde{M} = \rho_M(h, j)$ for $h \in H_M$ and $j \in N_M$. For all $\tilde{h} \in H_{\tilde{M}}$, $\mathcal{I}_{\tilde{M}}^j(\tilde{h}) = \mathcal{I}_M^j(\tilde{h}) \cap H_{\tilde{M}}$.*

Proof. From P1, $\Gamma_{\tilde{M}} \sqsubseteq \Gamma_M$. Hence, by R6, $\mathcal{I}_{\tilde{M}}^j(\tilde{h}) = \mathcal{I}_M^j(\tilde{h}) \cap H_{\tilde{M}}$ as required. ■

³ For $M = \boldsymbol{\rho}(h, i)$, even if Γ_M and $\mathbf{\Gamma}$ are identical, and the perception mappings are the same, we maintain a distinction between M and $\mathbf{M}$ because the associated languages will differ with respect to existential propositions.

Lemma 2 states that no information is gained in the shift from M to $\tilde{M} = \rho_M(h, j)$. Furthermore no information is lost, except insofar as this is implied by unawareness of some histories

In general, models involve differential awareness. That is, for at least some h in M , $\rho_M(h, j) \subset M$, so that the player who is to move at h is only aware of some restriction of M . However, the special case of common awareness is of particular interest.

Definition 4 *M is a model of common awareness if, for all $h \in H_M$ and $j \in N_M$, $\rho_M(h, j) = M$*

This definition implies that, for a model of common awareness M , $(M' \sqsubseteq M) \Rightarrow (M' = M)$. Hence, any model of common awareness is minimal under $\sqsubseteq$. Conversely, if a model M is minimal under $\sqsubseteq$, then there cannot exist any $\tilde{M} \subset M$ and hence there cannot exist $h \in H_M$ and $j \in N_M$, $\rho_M(h, j) \neq M$, so that M is a model of common awareness. We record this as a lemma.

Lemma 3 *M is a model of common awareness if and only if it is minimal under $\sqsubseteq$.*

Note that if M is a model of common awareness, the associated game Γ_M is a standard extensive game. The minimal element in $\mathcal{G}$ must be a model of common awareness.

2.4 Strategies and equilibrium in games with differential awareness

In a game with differential awareness $\mathcal{G}$, a strategy profile may be defined recursively. In this case, rather than beginning with a maximal element of $\mathcal{G}$, as in the recursive definition of a model, we start with the minimal (in terms of the ordering $\sqsubseteq$) elements of $\mathcal{G}$, that is, with models of common awareness. For this case, the definition of a strategy profile coincides with that of the standard definition, since the game is common knowledge.

Hence, we are led to the following recursive definition.

Definition 5 *Let $\mathcal{G}$ be a game with differential awareness. A **strategy profile** $\beta = (\beta_M : M \in \mathcal{G})$ for $\mathcal{G}$ assigns to each model $M \in \mathcal{G}$, a behavioral strategy profile β_M for the game Γ_M with the consistency property: if $\rho_M(h, i) = \hat{M}$ where $P(h) = i$ then $\beta_M^i(\mathcal{I}_M^i(h)) = \beta_{\hat{M}}^i(\mathcal{I}_{\hat{M}}^i(h))$.*

The definition of a strategy profile for a game with differential awareness differs from the standard definition because, in general, players are not fully aware of the game. Each player must impute a model, and a strategy profile consistent with that model to other players they regard as incompletely aware. In addition, as players become more aware during the play of the game, the definition requires that their behavioral strategies in the more complete model must be the *same* for those earlier histories in which their model of the game was less complete.

For a model of common awareness M , the set of strategy profiles subject to the consistency requirement define a strategy profile for the associated standard game Γ_M . Conversely any standard behavioral strategy profile for a standard game defines a behavioral strategy profile for the associated model of common awareness.

Now consider a model M , in which, for each pair (h, i) with $P_M(h) = i$ either $\rho_M(h, i) = M$ or $\rho_M(h, i) = \hat{M} \sqsubset M$ is a game of common awareness. Then, whenever $\rho_M(h, i) = M$, the strategy profile defines a probability measure on $A_M(\mathcal{I}_M^i(h))$. On the other hand, if $\rho_M(h, i) = \hat{M} \sqsubset M$ is a model of common awareness, then the $\beta_{\hat{M}}$ is a strategy profile for the associated standard game $\Gamma_{\hat{M}}$ and therefore defines a probability measure on $A_{\hat{M}}(\mathcal{I}_{\hat{M}}^i(h)) \subseteq A_M(\mathcal{I}_M^i(h))$, and the consistency property entails $\beta_M^i(\mathcal{I}_M^i(h)) = \beta_{\hat{M}}^i(\mathcal{I}_{\hat{M}}^i(h))$. That is, the randomization over actions for player i at her information set at h must be consistent with the model $\hat{M}$ that the model M imputes to player i at history h .

Proceeding recursively, we can build up the overall strategy profile $\beta = (\beta_M : M \in \mathcal{G})$. Since every model can be reduced to model of common awareness with a finite number of applications of ρ , the assignment is well-defined.

2.4.1 Beliefs, assessments and sequential rationality

Definition 6 Let $\mathcal{G}$ be a game with differential awareness. A **belief system**

$\mu = ((\mu_M^i)_{i \in N_M} : M \in \mathcal{G})$ for $\mathcal{G}$ assigns to each model $M \in \mathcal{G}$, and each individual i in the associated game Γ_M , a function μ_M^i that assigns to each history h in H_M , a probability measure on the set of histories in $\mathcal{I}_M^i(h)$ with the consistency properties: (i) for any h' in H_M , if $h' \in \mathcal{I}_M^i(h)$ then $\mu_M^i(h') = \mu_M^i(h)$ and (ii) if $\rho_M(h, i) = \hat{M}$ then $\mu_M^i(h) = \mu_{\hat{M}}^i(h)$.

The interpretation of the probability measure $\mu_M^i(h)$ is that for each history $h' \in \mathcal{I}_M^i(h)$, $\mu_M^i(h)[h']$ is the probability that player i assigns to history h' , conditional on the information set $\mathcal{I}_M^i(h)$ being reached. Thus consistency condition (i) ensures that conditional on information set $\mathcal{I}_M^i(h)$ being reached, player i 's beliefs are the same no matter which history in $\mathcal{I}_M^i(h)$ actually obtained. Condition (ii) requires that the beliefs imputed to player i in model M at her information set at h must be consistent with the model that model M imputes to player i at history h .

Definition 7 Let $\mathcal{G}$ be a game with differential awareness. An **assessment** for $\mathcal{G}$ is a pair (β, μ) where β is strategy profile for $\mathcal{G}$ and μ is a belief system for $\mathcal{G}$.

We shall refer to a behavioral strategy profile β for the game of awareness $\mathcal{G}$ as being *completely mixed as can be done consistently* if it assigns in each model $M \in \mathcal{G}$, each history h in H_M , and for the individual i for which $P_M(h) = i$ and $\rho_M(h, i) = \hat{M}$, positive probability to every action in $A_{\hat{M}}(\mathcal{I}_{\hat{M}}^i(h))$. That is, every action, consistent with the model $\hat{M}$ that model M imputes to player i at history h , has assigned positive weight.

Definition 8 Let $\mathcal{G}$ be a game with differential awareness. An assessment (β, μ) is *consistent* if there exists a sequence of $((\beta^n, \mu^n))_{n=1}^\infty$ that converges pointwise to (β, μ) and has the property that each strategy profile β^n is as completely mixed as can be done consistently and that each belief system μ^n is derived from β^n using Bayes' rule.

It just remains for us to define the concept of sequential rationality for an assessment in a game with differential awareness. To do this, fix an assessment (β, μ) for the game of awareness $\mathcal{G}$ and define the outcome $O(\beta_M, \mu_M | I)$ of (β_M, μ_M) conditional on information set I , to be the distribution over terminal histories of the game Γ_M associated with the model M in $\mathcal{G}$ determined by β_M and μ_M conditional on I being reached.

Let $\hat{h} = \langle \alpha_1, \alpha_2, \dots, \alpha_K \rangle \in Z_M$ denote a terminal history of Γ_M , and let $\hat{h}^k = \langle \alpha_1, \alpha_2, \dots, \alpha_k \rangle$, $k = 1, \dots, K$ denote its subhistories. Now fix a history $h \in H_M$. According to model M , the individual $P_M(h)$ will be choosing the action in that information set, but she is imputed to be in the model $\hat{M} := \rho(h, P_M(h))$ where she would perceive herself to be in the information set $\mathcal{I}_{\hat{M}}^{P_M(h)}(h)$. Hence we have:

$$O(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\hat{M}}^{P_M(h)}(h)) [\hat{h}] = \begin{cases} 0 & \text{if } h' \preceq \hat{h} \Rightarrow h' \notin \mathcal{I}_{\hat{M}}^{P_M(h)}(h) \\ \mu_{\hat{M}}^{P_M(h)}(h) [\hat{h}^L] \prod_{k=L}^{K-1} \beta_{\hat{M}}^{P_M(h)}[\alpha_{k+1}] & \text{if } \hat{h}^L \in \mathcal{I}_{\hat{M}}^{P_M(h)}(h) \end{cases}.$$

Given the affine utility function $v_{\hat{M}}^{P_M(h)}$, the continuation expected utility imputed by the model M for the individual $P_M(h)$ who makes the choice of action at information set $\mathcal{I}_{\hat{M}}^{P_M(h)}(h)$ is therefore given by

$$\begin{aligned} & v_{\hat{M}}^{P_M(h)} \left(O(\beta_M, \mu_M | \mathcal{I}_{\rho(h, P_M(h))}^{P_M(h)}(h)) \right) \\ &= \sum_{\hat{h} \in Z_M} v_{\hat{M}}^{P_M(h)}(\hat{h}) \times O(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\hat{M}}^{P_M(h)}(h)) [\hat{h}]. \end{aligned}$$

In a standard game, an assessment is deemed sequentially rational if for every player i and every history h for which $P(h) = i$ the strategy of player i is a best response to the other players' strategies given i 's belief at the information set $\mathcal{I}^i(h)$. That is, there are no strictly profitable deviations. But as is also well-known, to show that there are no strictly profitable deviations, it is sufficient to show there are no strictly profitable one-shot deviations, that is, it is enough to show that, at each information set, the conditional expected payoff cannot be increased by deviating from the current strategy just at that information set.

For ease of exposition we shall adapt the requirement for no strictly profitable one-shot deviations to extend the concept of sequential rationality to games of awareness. To evaluate the profitability of a one-shot deviation, we need to be able to evaluate the expected utility for a decision maker at a particular information set who selects a particular action available to her in that information set and then follows the actions prescribed by the strategy profile β , thereafter. So in the model M at history $h \in H_M$, $P_M(h)$ is imputed to be in the model $\hat{M} := \rho(h, P_M(h))$, and so she has available any action α in $A_{\hat{M}}(\mathcal{I}_{\hat{M}}^{P_M(h)}(h))$. She needs to consider the continuation

of the game for each of the histories $\{h' \bullet \langle \alpha \rangle : h' \in \mathcal{I}_{\hat{M}}^{P_M(h)}\}$ she considers as possible continuation histories stemming from her choice of action α . More precisely, for each given $h' \in \mathcal{I}_{\hat{M}}^{P_M(h)}(h)$, she needs to consider the ‘outcome’:

$$O\left(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\rho_{\hat{M}}(h' \bullet \langle \alpha \rangle, P_{\hat{M}}(h' \bullet \langle \alpha \rangle))}^{P_{\hat{M}}(h' \bullet \langle \alpha \rangle)}(h' \bullet \langle \alpha \rangle)\right).$$

This is the induced distribution over terminal nodes for $(\beta_{\hat{M}}, \mu_{\hat{M}})$ (the assessment for the model $\hat{M}$) conditional on the information set $\mathcal{I}_{\rho_{\hat{M}}(h' \bullet \langle \alpha \rangle, P_{\hat{M}}(h' \bullet \langle \alpha \rangle))}^{P_{\hat{M}}(h' \bullet \langle \alpha \rangle)}(h' \bullet \langle \alpha \rangle)$ having been reached. The action choice at this information set is controlled by the individual $P_{\hat{M}}(h' \bullet \langle \alpha \rangle)$ who in turn is imputed (by the model $\hat{M}$) to be in the model $\rho_{\hat{M}}(h' \bullet \langle \alpha \rangle, P_{\hat{M}}(h' \bullet \langle \alpha \rangle))$. According to the beliefs μ , the probability that individual $P_M(h)$ is at history h' in the model $\hat{M}$ is given by $\mu_{\hat{M}}^{P_M(h)}(h)[h']$. Thus her continuation expected utility from choosing the action α may be expressed as:

$$\sum_{h' \in \mathcal{I}_{\hat{M}}^{P_M(h)}(h)} \mu_{\hat{M}}^{P_M(h)}(h') \sum_{\hat{h} \in Z_M} v^{P_M(h)}(\hat{h}) \times O\left(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\rho_{\hat{M}}(h' \bullet \langle \alpha \rangle, P_{\hat{M}}(h' \bullet \langle \alpha \rangle))}^{P_{\hat{M}}(h' \bullet \langle \alpha \rangle)}(h' \bullet \langle \alpha \rangle)\right)(\hat{h}),$$

where $\hat{M} = \rho_M(h, P_M(h))$.

Putting this all together we have:

Definition 9 Let $\mathcal{G}$ be a game with differential awareness. An **assessment** (β, μ) is **sequentially rational** if for each $M \in \mathcal{G}$ and each $h \in H_M$, and setting $\hat{M} := \rho_M(h, P_M(h))$,

$$\begin{aligned} & v_{\hat{M}}^{P_M(h)}\left(O\left(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\hat{M}}^{P_M(h)}(h)\right)\right) \\ & \geq \sum_{h' \in \mathcal{I}_{\hat{M}}^{P_M(h)}(h)} \mu_{\hat{M}}^{P_M(h)}(h') \sum_{\hat{h} \in Z_M} v^{P_M(h)}(\hat{h}) \times O\left(\beta_{\hat{M}}, \mu_{\hat{M}} | \mathcal{I}_{\rho_{\hat{M}}(h' \bullet \langle \alpha \rangle, P_{\hat{M}}(h' \bullet \langle \alpha \rangle))}^{P_{\hat{M}}(h' \bullet \langle \alpha \rangle)}(h' \bullet \langle \alpha \rangle)\right)(\hat{h}), \\ & \text{for all } \alpha \in A_{\hat{M}}\left(\mathcal{I}_{\hat{M}}^{P_M(h)}(h)\right). \end{aligned}$$

For the case of a model with common awareness, where $\rho(h, i) = M$ for all (h, i) , this is equivalent to the standard definition of a sequential equilibrium.

2.5 Equilibrium and existence

Definition 10 Let $\mathcal{G}$ be a game with differential awareness. An **assessment** (β, μ) is a **sequential equilibrium** if it is sequentially rational and consistent.

We now demonstrate the existence of a sequential equilibrium for games with differential awareness.

Proposition 4 A sequential equilibrium exists for any game with differential awareness.

Proof. Define the order n of a game $\mathcal{G}$ as the maximum length of a chain in $\mathcal{G}$. We can now derive the result by induction, starting with the observation that any maximal chain in a game with differential awareness $\mathcal{G}$ must have, as its minimal element under $\sqsubseteq$, a model of common awareness.

For the case $n = 1$, $\mathcal{G}$ is a model with common awareness. A model with common awareness corresponds to a standard finite game with perfect recall and therefore the standard existence result applies.

For the general case, suppose the result is true for $n = 1 \dots N$, and let $\mathcal{G}$ be of order $n = 1 \dots N + 1$ with maximal model $\mathbf{M}$ and associated game Γ . Since the maximal model $\mathbf{M}$ is not accessible to any of the players, the recursive construction of models ensures that for each h , $\rho(\mathbf{P}(h), h) = M \sqsubset \mathbf{M}$ where M defines a game of order n . For each M in the latter class, choose an imputed strategy profile that is a sequential equilibrium for M .

Now consider the game Γ' in which, for each h such that $\rho(P(h), h) \sqsubset M$, the move by player $j = P(h)$ is replaced by a move by Nature such that $f_c(h) = \beta_{\rho(h,j),h}^j$ the equilibrium strategy for j in the game associated with $\tilde{M} = \rho(h, j)$. This is a game of common awareness, and therefore the standard existence result applies. Further, the equilibrium derived in this way is a sequential equilibrium for $\mathbf{M}$ as defined above. ■

2.6 Example

To illustrate the recursive construction of models, we adapt the speculative trade example of Heifetz, Meier and Schipper (2006). In this example, a buyer (player 1) and an owner (player 2) may contract the sale of the owner's firm at a price of 1. The value of the firm depends on two contingencies; the possibility of a lawsuit which would reduce the value by L and a business opportunity which would increase the value by G . If neither occurs, the value remains unchanged at 1. We represent the maximal game Γ as follows. Nature has two initial moves determining whether the lawsuit and business opportunity arise. Before learning about Nature's moves, the buyer chooses whether to make an offer of 1. If an offer is made, the owner chooses whether to accept it, also before learning about Nature's moves. At the terminal nodes, players receive their net payoffs and Nature's moves are revealed.

We first describe the game in the maximal model. The initial history is $\langle \rangle$. Nature's first move is a choice from the set $\{\alpha_n, \alpha_0\}$, (innovation or null action). Nature's second move is a choice from the set $\{\alpha_l, \alpha_0\}$ (lawsuit or null action). So there are now four histories $\langle \alpha_n, \alpha_l \rangle$, $\langle \alpha_n, \alpha_0 \rangle$, $\langle \alpha_0, \alpha_l \rangle$, $\langle \alpha_0, \alpha_0 \rangle$, forming an information set which we shall denote $\mathbf{I}_1$. At $\mathbf{I}_1$, player 1 chooses from the set $\{\alpha_1, \alpha_0\}$ (offer 100 or null action). If 1 chooses α_0 , the game terminates. If 1 chooses α_1 the information set becomes $\mathbf{I}_2 = \{\langle \alpha_n, \alpha_l, \alpha_1 \rangle, \langle \alpha_n, \alpha_0, \alpha_1 \rangle, \langle \alpha_0, \alpha_l, \alpha_1 \rangle, \langle \alpha_0, \alpha_0, \alpha_1 \rangle\}$ and 2 chooses from the set $\{\alpha_A, \alpha_R\}$ (accept or reject the offer). The maximal game is illustrated in

figure 1.

<INSERT FIGURE 1 around here>

As in Heifetz, Meier and Schipper (2006), we suppose that the buyer is unaware of the possibility of a lawsuit while the seller is unaware of the possibility of an innovation. Thus, at each non-terminal history h , the maximal model imputes to the buyer a model $\rho(h, 1) = M^1 = (\Gamma^1, \rho^1)$. In Γ^1 all histories containing the move α_ℓ are deleted. Similarly, at each non-terminal history h , the maximal model imputes to the owner a model $\rho(h, 2) = M^2 = (\Gamma^2, \rho^2)$ in which all histories containing the move α_n are deleted. The games are illustrated in figures 2 and 3.

<INSERT FIGURES 2 and 3 here>

Both parties impute to the other a restriction of their own model. The buyer is unaware of the lawsuit, and assumes the owner to be unaware of the innovation (at all non-terminal histories), while the converse is true for the owner. The two restricted models denoted $M^{12} = \rho^1(h, 2)$ and $M^{21} = \rho^2(h, 1)$ are identical, and will be denoted M^3 . Complete (that is, terminal) histories in Γ^3 begin with two null moves α_0 by Nature, followed by the decisions of the two players. For all histories h , and for $j = 1, 2$, $\rho^3(h, j) = M^3$. That is, M^3 as illustrated in figure 4 is a model of common awareness.

<INSERT FIGURE 4 around here>

We have not yet described the imputed model at terminal nodes in games other than M^3 . Since Nature's moves are revealed, the players must become aware of histories incorporating those moves, and of their own past unawareness, as well as that of the other player. Hence, for any terminal history h that begins with $\langle \alpha_n, \alpha_\ell \rangle$ and any player j , the maximal model imputes full awareness of the game. That is $\rho(h, j) = M^0$, where $\Gamma^0 = \Gamma$.

By contrast, for terminal histories h beginning with $\langle \alpha_n, \alpha_0 \rangle$, the owner must become aware of the innovation, but there is no reason for the buyer to become aware of the unrealized possibility of the lawsuit. Thus, for such histories, we have $\rho(h, 1) = M^1$, $\rho(h, 2) = \mathbf{M}$. The converse applies for terminal histories h beginning with $\langle \alpha_0, \alpha_\ell \rangle$, so that $\rho(h, 1) = \mathbf{M}$, $\rho(h, 2) = M^2$. Histories beginning with $\langle \alpha_0, \alpha_0 \rangle$ produce no change in awareness when Nature's move is revealed, that is, $\rho(h, 1) = M^1$, $\rho(h, 2) = M^2$.

Similarly, in the buyer's model M^1 , for terminal histories h in which Nature's move reveals the innovation, the owner becomes aware of this and $\rho^1(h, 2) = M^1$. But if the innovation does not take place, the owner remains unaware and $\rho^1(h, 2) = M^{12} = M^3$. The converse applies for the owner.

Heifetz, Meier and Schipper (2006) propose a dominance principle that is sufficient to ensure that trade takes place in this model. In our model this corresponds to a sequential equilibrium $(\hat{\beta}, \hat{\mu})$, in which

$$\begin{aligned}\beta_{\mathbf{M}}^1(\mathbf{I}_1)[\alpha_1] &= \beta_{M^1}^1(I_1^1)[\alpha_1] = \beta_{M^2}^1(I_1^2)[\alpha_1] = \beta_{M^3}^1(\langle\alpha_0, \alpha_0\rangle)[\alpha_1] = 1 \\ \beta_{\mathbf{M}}^2(\mathbf{I}_2)[\alpha_A] &= \beta_{M^1}^2(I_2^1)[\alpha_A] = \beta_{M^2}^2(I_2^2)[\alpha_A] = \beta_{M^3}^2(\langle\alpha_0, \alpha_0\rangle)[\alpha_A] = 1\end{aligned}$$

and

$$\begin{aligned}\mu_{\mathbf{M}}^1(\mathbf{I}_1)[\langle\alpha_n, \alpha_0\rangle] + \mu_{\mathbf{M}}^1(\mathbf{I}_1)[\langle\alpha_n, \alpha_\ell\rangle] &= \mu_{M^1}^1(I_1^1)[\langle\alpha_n, \alpha_0\rangle] > 0 \\ \mu_{\mathbf{M}}^1(\mathbf{I}_1)[\langle\alpha_0, \alpha_0\rangle] + \mu_{\mathbf{M}}^1(\mathbf{I}_1)[\langle\alpha_0, \alpha_\ell\rangle] &= \mu_{M^1}^1(I_1^1)[\langle\alpha_0, \alpha_0\rangle] > 0 \\ \mu_{M^1}^2(I_2^1)[\langle\alpha_n, \alpha_0, \alpha_1\rangle] &= 0 \\ \mu_{M^1}^2(I_1^1)[\langle\alpha_0, \alpha_0, \alpha_1\rangle] &= \mu_{M^3}^2(\langle\alpha_0, \alpha_0, \alpha_1\rangle)[\langle\alpha_0, \alpha_0, \alpha_1\rangle] = 1 \\ \mu_{\mathbf{M}}^2(\mathbf{I}_2)[\langle\alpha_n, \alpha_\ell, \alpha_1\rangle] + \mu_{\mathbf{M}}^2(\mathbf{I}_2)[\langle\alpha_0, \alpha_\ell, \alpha_1\rangle] &= \mu_{M^2}^2(I_2^2)[\langle\alpha_0, \alpha_\ell, \alpha_1\rangle] > 0 \\ \mu_{\mathbf{M}}^2(\mathbf{I}_2)[\langle\alpha_n, \alpha_0, \alpha_1\rangle] + \mu_{\mathbf{M}}^2(\mathbf{I}_2)[\langle\alpha_0, \alpha_0, \alpha_1\rangle] &= \mu_{M^2}^2(I_2^2)[\langle\alpha_0, \alpha_0, \alpha_1\rangle] > 0 \\ \mu_{M^2}^1(I_1^2)[\langle\alpha_0, \alpha_\ell\rangle] &= 0 \\ \mu_{M^2}^1(I_1^2)[\langle\alpha_0, \alpha_0\rangle] &= \mu_{M^3}^1(\langle\alpha_0, \alpha_0\rangle)[\langle\alpha_0, \alpha_0\rangle] = 1.\end{aligned}$$

To check sequential rationality, notice that at $I_1^1 = \mathbf{I}_1 \cap H_{M^1}$ (in M^1) the action α_0 (no offer) leads to a payoff of zero for the buyer in all histories. The action α_1 yields a net payoff of G in the history $\langle\alpha_n, \alpha_0, \alpha_1, \alpha_A\rangle$ and a net payoff of zero in all other histories (those where the innovation is not realized or the owner rejects the offer).

For the owner at $I_2^2 = \mathbf{I}_2 \cap H_{M^2}$ (in M^2) the action α_A yields a sure net payoff of 0, while α_R yields a net payoff of $-L$ for the history $\langle\alpha_0, \alpha_\ell, \alpha_1, \alpha_R\rangle$ and 0 for $\langle\alpha_0, \alpha_0, \alpha_1, \alpha_R\rangle$. Finally, in M^3 all actions yield the payoff pair $(0, 0)$.

3 Models, propositions and languages: the syntactic representation

Fix a game with differential awareness $\mathcal{G}$ with associated maximal element $\mathbf{M} = (\Gamma, \rho)$. With each model $M = (\Gamma_M, \rho_M) \in \mathcal{G}$, we associate a syntactic structure, defined as propositions in a formal language.

For any $M \in \mathcal{G}$, we begin with a set of elementary propositions P_M . P_M is the closure, under standard logical operators $\vee, \wedge$ and $\neg$, of the set $\{p_h : h \in H_M\}$, where the elementary proposition p_h may be stated as ‘the current history is h ’. Thus, the syntactic structure P_M is directly related to the game structure Γ_M .

For $p \in P_M$, truth is relativized to histories h : for $h \in H_M$, the statement ‘ p is true at h (in the model M)’ is written $h \models_M p$. The relation $\models_M$ can be derived, using the standard rules

of logic, from the elementary requirements $h \models_M p_h$ and $\tilde{h} \models_M \neg p_h$ for $\tilde{h} \neq h$, for $h, \tilde{h} \in H_M$. Conversely, truth is determined by a valuation function $V_M : P_M \times H_M \rightarrow \{True, False\}$, where $V_M(p, h) = True$ if and only if $h \models_M p$. It is easy to check that truth values satisfy the usual logical properties.

3.1 Time, belief and awareness

We now extend P_M to a richer language Q_M which incorporates modal operators describing time, belief and awareness. More precisely, Q_M is the closure of P_M under the operators w , b , a and u to be defined below.

First, the language Q_M includes a temporal operator w^h for which the valuation is derived from $\models_M$ and the structure of the game Γ_M . The temporal operator w^h (read as ‘was true at h ’) is defined as follows.

Definition 11 *Fix a model $M \in \mathcal{G}$. For any $p \in P_M$, $h, h' \in H_M$, $h' \models_M w^h p$ if $h \preceq h'$ and $h \models_M p$.*

Interest in temporal operators derives from the fact that, under the definitions of $\models_M$ and $\mathbf{V}_M$, a given proposition p may be true at one history, but false at some subsequent history.^{4, 5}

Next, we consider beliefs. A proposition is believed true by an individual in a given model, if it is true in all histories in the information set imputed to the individual in the model. Hence we have:

Definition 12 *Fix a model $M \in \mathcal{G}$, $j \in N_M$ and $h \in H_M$. Let $\tilde{M} = \rho_M(h, j)$ and recall $\mathcal{I}_{\tilde{M}}^j(h) = \mathcal{I}_M^j(h) \cap H_{\tilde{M}}$ is the information set imputed to j at h .*

Then, $h \models_M b_j p$, if $\tilde{h} \models_{\tilde{M}} p$ for all $\tilde{h} \in \mathcal{I}_{\tilde{M}}^j(h)$. ($b_j p$ is read as ‘in the model $\tilde{M} = \rho_M(h, j)$ imputed to j in the model M at h , j believes p)

In the special case $\tilde{M} = M$, the operator $b_j p$ corresponds to the standard modal definition of the knowledge operator as in Halpern (2003). If $\tilde{M} \sqsubset M$, however, $h \models_M b_j p$ is consistent with $h \models_M \neg p$. That is, proposition p may be false at the history h but true for all histories in the information set $\mathcal{I}_{\tilde{M}}^j(h)$, given the model $\tilde{M}$ imputed to j at h .⁶ Hence the model is one in which individuals can hold false beliefs and (correctly or otherwise) impute false beliefs to others.

⁴ As will be shown below, the fact that the same proposition may be evaluated differently at h and at h' for $h \preceq h'$ is crucial in considering inductive reasoning, which derives support for a proposition from the fact that it has held true in the past. Induction would be trivial if propositions were evaluated on complete (terminal) histories and invariant for subhistories.

⁵ On the use of temporal logic in game theory, see Bonanno (2001).

⁶ Under R6 above, this can only be true if $h \notin H_{\tilde{M}}$, that is, if the actual history is one not considered by j (in the model imputed to j by ρ_M).

We can now derive syntactic notions of awareness and unawareness. The standard definition of awareness is that an individual is aware of a proposition if they know its truth value, or know that they do not know its truth value. Since the belief operator b corresponds to the knowledge operator, we propose:

Definition 13 Fix a model $M \in \mathcal{G}$ and choose $h \in H_M$, $p \in Q_M$, and $j \in N_M$. Then $h \models_M a_j p$ if $h \models_M b_j p \vee b_j \neg p \vee b_j (\neg b_j p \wedge \neg b_j \neg p)$

That is, j is aware of q in the model M , if, given the perceived model $\tilde{M}$ imputed to j in the model M , one of the following holds: (i) j believes p , (ii) j believes $\neg p$ or (iii) j believes that she does not hold a belief about p . (Notice that $b_j p \vee b_j \neg p \vee b_j (\neg b_j p \wedge \neg b_j \neg p)$ may be stated more compactly as $b_j p \vee b_j (\neg b_j p)$). We denote unawareness as the negation of awareness, that is, $u_j p$ is a synonym for $\neg a_j p$.

The construction of the operators w, b, a and u incorporates an extension of $\models_M$ from P_M to Q_M . Correspondingly, we can extend the valuation operator V_M to the domain $Q_M \times H_M$ by setting $V_M(p, h) = \text{True}$ if and only if $h \models_M p$, as before.

As can be inferred from the results of Dekel, Lipman and Rustichini (1998), for a partitional information structure (which we have assumed), $a_i p$ is trivially true when $\rho_M(h, j) = \tilde{M} = M$.

Proposition 5 Fix a model $M \in \mathcal{G}$ and history $h \in H_M$. For any $j \in N_M$ such that $\rho_M(h, j) = M$, and any $p \in Q_M$, $h \models_M a_j p$. Conversely if $p \in Q_M - Q_{\tilde{M}}$ and $\rho_M(h, j) = \tilde{M}$, $h \models_M u_j p$.

Proof. For any $p \in Q_M$, let $\mathfrak{S}_M(p) = \{h' \in H_M : V_M(p, h') = \text{True}\}$. If $\mathcal{I}_M^j(h) \subseteq \mathfrak{S}_M(p)$, then $h \models_M b_j p$ and otherwise $h \models_M \neg b_j p$. Since the information structure is partitional, $\neg b_j p \Rightarrow b_j (\neg b_j p)$ so $h \models_M a_j p$ as required.

For the converse, observe that by construction of $\models_{\tilde{M}}$, we have $h \models_{\tilde{M}} p \vee h \models_{\tilde{M}} \neg p$ if, and only if, $p \in Q_{\tilde{M}}$. Hence, from Definition 12, $b_j p \Rightarrow p \in Q_{\tilde{M}}$ and therefore, from Definition 13, $a_j p \Rightarrow p \in Q_{\tilde{M}} \Rightarrow p \notin Q_M - Q_{\tilde{M}}$. ■

In particular, consider $\mathbf{Q}$, the language associated with the maximal model $\mathbf{M}$. For any $M \sqsubseteq \mathbf{M}$, $P_M \subseteq \mathbf{P}$, so that $Q_M \subseteq \mathbf{Q}$ and (since $\mathbf{M} \sqsubseteq \mathbf{M}$)

$$\mathbf{Q} = \bigcup_{M \sqsubseteq \mathbf{M}} Q_M$$

It follows that for any M , $p \in Q_M$, $h \in \mathbf{H}$, and i such that $\rho(i, h) = M$, such that $\Gamma_M = \Gamma$ we have $h \models_M a_i p$. That is, if individual i has access to the game Γ , she is aware of any proposition expressible in any restriction Γ_M .

Also, by P4, $h' \prec h \in H$, $\rho(h', i) \sqsubseteq \rho(h, i)$. Hence, if $h' \models_{\mathbf{M}} a_i p$, then also $h \models_{\mathbf{M}} a_i p$, so that the interpretation of P.4 as stating that awareness increases over time is justified. If there is a strict inclusion $\rho(h', i) \subset \rho(h, i)$ then there exists $p \in P_M$ such that $h' \models_{\mathbf{M}} u_i p \wedge h \models_{\mathbf{M}} a_i p$. If,

in addition, $h \models_{\mathbf{M}} b_i p$, then individual i , at h , believes that her previous model did not include, even as a possibility, the proposition p which she now believes to be true.

4 Existential propositions and unawareness

In this section we consider a question raised by Halpern and Rego (2006a): Can an individual know (or, in our terminology, believe) that there exists a proposition of which he is unaware? We show that, with the specifications of believe and awareness presented above, the answer to this question is negative. We therefore consider alternative ways of representing the individual's understanding of his limited awareness, and introduce a consciousness operator. We assume that individuals are sufficiently sophisticated that they may reason about propositions of the general form 'there exist propositions, having some implications, of which I am currently unaware'.

4.1 The existential quantifier

Although the language Q_M is sufficient to describe the model M , it is inadequate to describe propositions a reasonable individual might entertain about models M^* such that $M \sqsubset M^*$ and, in particular, about the maximal model $\mathbf{M}$. To describe such propositions we need an extended language $\hat{Q}_M$ which allows for reasoning about sets of propositions, some of which may not be included in Q_M .

In addition to the logical and model operators used in the construction of Q_M , the language $\hat{Q}_M$ incorporates the existential quantifier $\exists$, used in conjunction with a formula for substitution. We define $\hat{Q}_M$ as logical closure of the union of Q_M with the set of existential propositions of the form

$$\exists q \in Q_{\hat{M}} (\theta(q)),$$

where $\hat{M} \in \{M : \tilde{M} \sqsubseteq M\} \cup \mathbf{M}$ and $\theta(q)$ is a Boolean combination of the free proposition q and propositions in Q_M .⁷ That is, the existential quantifier $\exists$ may be applied either to models that are imputed to some player j at a history h in the model M , or to the maximal model $\mathbf{M}$ taken to represent unbounded rationality.

As an illustration, for given $p \in Q_M$ we can write,

$$\exists q \in Q_{\hat{M}} ((q \Rightarrow p) \wedge \neg(p \Rightarrow q))$$

which we may interpret as saying that there is some (non-equivalent) proposition q that implies p . For example, in a criminal investigation, the fact that a person is classed as a suspect typically means that, if some additional evidence were obtained, that person's guilt could be inferred. However, investigators will not, in general, know the exact nature of the evidence they are looking

⁷ We thank Joe Halpern for suggesting this formulation of the existential quantifier.

for. The evidence could be either propositional (X was at the scene of the crime) or epistemological (X knew that the gun was loaded).

4.2 Beliefs about existential propositions

We first consider how the relation $\models_M$ and the function V_M may be extended to $\hat{Q}_M$. We begin by considering V_M , and propose the mapping $\hat{V}_M : \hat{Q}_M \times H_M \rightarrow \{True, False, Undecided\}$. We begin by observing that, for any $Q_{\tilde{M}}$, the existential proposition $\exists q \in Q_{\tilde{M}}, \theta(q)$ may be rewritten as

$$q^* \equiv \bigvee_{q \in Q_{\tilde{M}}} \theta(q). \quad (1)$$

Since $Q_{\tilde{M}}$ contains only finitely many semantically distinct elements, this proposition may be regarded as a finite disjunction and therefore as an element of $Q_{\tilde{M}}$.

For the maximal language $\mathbf{Q}$, we have $\mathbf{Q} = \hat{\mathbf{Q}}$ since every proposition in $\hat{\mathbf{Q}}$ may be regarded as a finite disjunction of elements of $\mathbf{Q}$.

For any M such that $\tilde{M} \sqsubseteq M$ (and therefore $Q_{\tilde{M}} \subseteq Q_M$) the valuation V_M may be derived directly. $\hat{V}_M(q^*, h) = True$ if $V_M(\theta(q), h) = True$ for some $q \in Q_{\tilde{M}}$, and $\hat{V}_M(q^*, h) = False$ otherwise. More generally, if $V_M(\theta(q), h) = True$ for some $q \in Q_{\tilde{M}} \cap Q_M$, then $V_M(q^*, h) = True$. However, if $Q_{\tilde{M}} \not\subseteq Q_M$, then, for $q \in Q_{\tilde{M}} - Q_M$, and any $h \in H_M$ neither $V_M(\theta(q), h) = True$ nor $V_M(\theta(q), h) = False$. Hence, in this case, we set $\hat{V}_M(q^*, h) = Undecided$.⁸

Now we can extend the relation $\models_M$ to $\hat{Q}_M$ to state that, for $h \in H_M, q \in \hat{Q}_M$, $h \models_M q$ if and only if $V_M(q^*, h) = True$. Otherwise we may write $h \not\models_M q$. Note that for $q \in Q_M$, $h \not\models_M q \Leftrightarrow h \models_M \neg q$, but this is not true for $q \in \hat{Q}_M$ in general.

Given the extension of the relation $\models_M$ to $\hat{Q}_M$, the definitions of the modal operators w, b, a, u also extend to $\hat{Q}_M$. The belief operator is of particular interest. From the argument above, we have:

Lemma 6 (i) If for some $p \in Q_M$, $h \models_M b_i \theta(p)$, then

$$h \models_M b_i (\exists q \in Q_M (\theta(q))).$$

(ii) If, for all $p \in Q_M$, $h \models_M \neg b_i \theta(p)$, then

$$h \models_M b_i \neg (\exists q \in Q_M (\theta(q))).$$

On the other hand, suppose that the existential proposition $p \in \hat{Q}_M$ is not logically equivalent to any elementary proposition $p \in Q_M$. This means that no possible knowledge regarding primitive propositions, that is, regarding current and past histories, is sufficient to infer the truth or falsity of q . In this case, we say $\neg b_i q$.

⁸ Note that we regard as undecided propositions that may be tautologically true such as $q^* \vee \neg q^*$ if the components q^* are themselves undecided.

4.3 Beliefs and unawareness

We now present a key result, showing that an individual can never believe (in the modal-logical sense formalized above) that there exists a proposition of which he is unaware.

Proposition 7 *For any $h \in H_M$, and any $\tilde{M}$,*

$$h \not\models_M b_i (\exists q \in Q_{\tilde{M}} (u_i q)).$$

Proof. If $\tilde{M} \subseteq M$ then by Proposition 5, $h \models_M \neg u_i p$ for all $p \in Q_{\tilde{M}}$, so, by 6(ii), $h \models_M \neg b_i (\exists q \in Q_{\tilde{M}} (u_i q))$, which implies $h \not\models_M b_i (\exists q \in Q_{\tilde{M}} (u_i q))$.

If $M \subset \tilde{M}$, then $V((\exists q \in Q_{\tilde{M}} (u_i q)), h) = \text{Undecided}$, so $h \not\models_M (\exists q \in Q_{\tilde{M}} (u_i q))$ and, once again, $h \not\models_M b_i (\exists q \in Q_{\tilde{M}} (u_i q))$. ■

Proposition 7 is central to our analysis and to the concerns of Halpern and Rêgo (2006a). Halpern and Rêgo (2006a) raise the question of whether there exists an extension of the logic of awareness such that it would be possible to say, for example, that an agent knows that there exists a fact of which he is unaware, but of which some other agent is aware. Proposition 7 shows that the belief operator b_i which generates the set of propositions Q_M is not powerful enough to permit statements of the form ‘I believe that there exists some proposition q of which I am currently unaware’. Since, as we have shown, the belief operator in a model with differential awareness is the analog of the knowledge operator in a model of common awareness, this result provides a negative answer to the question raised by Halpern and Rêgo.

We must therefore consider whether a boundedly rational, but nevertheless sophisticated, individual might be able to reason about their own limited awareness, using methods outside the scope of the modal-logical framework considered thus far. One such method is that of induction.

5 Inductive reasoning about unawareness

We now address the central question for any account of limited awareness: in what sense can an individual derive support, from experience or observation, for the proposition that there exist propositions of which she is unaware?

We begin by considering reasoning based on historical induction. Informally, the principle of historical (or temporal) induction states that if a proposition has been found to be true in many past instances, this fact provides support for belief that it will hold true in the future. For example, the fact that the proposition ‘the sun will rise tomorrow’ was true yesterday, the day before and the day before that and so on, provides inductive support for the belief that the same proposition is true today.⁹

⁹ Notice that in changing the history to which the proposition refers we may also change the subjective model and the temporal reference of terms like ‘tomorrow morning’.

Formally, we state

Definition 14 (support by historical induction) Fix a model M , $q \in \widehat{Q}_M$, $h \in H_M$ and i such that $\rho(h, i) = M$

(i) for at least some $h' \preceq h$, $h \models_M b_i w^{h'} q$; and

(ii) for no $h'' \preceq h$, $h \models_M b_i (w^{h''} \neg q)$,

then $h \models_M t_i q$ [read as ‘at h individual i regards q as supported by historical induction’].

That is, suppose at history h that individual i believes q to have been true for some past h' and does not believe q to have been false for any past h'' , then i regards q as supported by historical induction.

The definition adopted here allows for inductive reasoning on the basis of limited evidence, provided there is no disconfirming evidence. Consistent with this relatively modest requirement for evidential support, we do not require that a proposition that is supported by induction be believed true, but merely that it be regarded as possible.¹⁰

Our account of historical inductive reasoning about unawareness has a structure similar to that of more familiar examples of historical induction. Over time, individuals become aware of propositions that, previously, they have not considered. Thus, at any history h , individuals know that there exist propositions of which they were unaware at some previous h' . That is, they know that the existential proposition ‘there exists a proposition of which I am unaware’ was true at h' . Indeed, they know that this existential proposition has been true at all past histories, except perhaps recent histories in a period in which no new propositions have been discovered. Hence, the proposition that the future, like the past, will be characterized by the discovery of new propositions, is supported by induction.¹¹

Suppose that for $h' \prec h$, individual i becomes aware of previously unconsidered propositions in the course of the partial history from h' to h , so that $Q_{\rho_M(h', i)} \subset Q_{\rho_M(h, i)} = Q_M$ where the inclusion $\subset$ is strict. Then for propositions $q \in Q_{\rho_M(h, i)} - Q_{\rho_M(h', i)}$ we have $h \models_M a_i q$ and $h \models_M w^{h'} u_i(q)$. Further, these evaluations hold for any $\tilde{h} \in \mathcal{I}_M^j(h)$, so

$$h \models_M b_i \left(\exists q \in Q_M \left(a_i q \wedge w^{h'} u_i(q) \right) \right). \quad (2)$$

¹⁰ Halpern (2003) discusses reasoning based on judgements about possibility, which may be, but need not be, supported by induction. As shown by Halpern, reasoning based on concepts of possibility may be extended to ordinal rankings of relative likelihood, yielding some, but not all, of the power of a fully-developed probabilistic model. For example, a proposition known true at many past histories might be regarded as more likely than one known true at only a single history as required by Definition 14. We will not explore these issues here.

¹¹ A closely related argument is prominent in philosophical debates over ‘realism’, namely, the view that the success of science reflects its correspondence to objective truth. Critics such as Laudan (1981) argue on the basis of historical experience that, since successful theories have been proven false in the past, the success of a theory cannot be regarded as evidence for its truth. Similarly, in our analysis, the fact that models used with some success in decisionmaking have nonetheless been discovered to be incomplete in the past supports the view that the model currently held by any given decisionmaker is also unlikely to be complete. However, we allow for increasing awareness over time, and therefore for a model of the world that gradually converges towards the true model.

That is, individual i believes at $\tilde{h}$ that there exist propositions that she is aware of now, but was unaware of in the past. Then is,

Definition 15 Fix M , i and $h \in H_M$, and suppose $\rho(h, i) = M$. If (2) holds for some $h' \prec h$, individual i displays increasing awareness at h .

We have established that individual i at h cannot express and therefore cannot believe (in the modal logical sense of b_i) propositions of the form $\exists q \in Q_{\rho_M(h)} : u_i(q)$, or their negations. Nevertheless, given past experience of discovery, it seems reasonable to suppose that the individual may judge such propositions to be an appropriate basis for actions and decisions. Given the dynamic temporal structure of the model developed here, it is natural to consider whether historical/temporal induction can be used as a basis for such judgements. Our next result provides a positive answer to this question.

Proposition 8 Fix M , i and $h \in H_M$, and suppose $\rho(h, i) = M$. If individual i displays increasing awareness at h , then $h \models_M \mathbf{t}_i(\exists q \in \mathbf{Q} : u_i q)$.

Proof. Denote by $\hat{q}$ the existential proposition $\exists q \in \mathbf{Q} : u_i(q)$, for which the induction will be undertaken at h . For any $h' \preceq h$ consider the proposition $w^{h'}(\exists q \in \mathbf{Q}(u_i(q)))$. Let $M = \rho(h, i)$, $M' = \rho(h', i)$. By P1, $M' \sqsubseteq M$.

If $M' = M$, then previous results show $h \models_M (\neg b_i \hat{q} \wedge \neg b_i \neg \hat{q})$

If $M' \sqsubset M$, then $\exists \tilde{q} \in Q_M - Q_{M'} \subseteq \mathbf{Q} - Q_{M'}$ and hence $h \models_M b_i w^{h'} u_i(\tilde{q})$, so $h \models_M b_i w^{h'} \hat{q}$.

Hence provided that for some past h' , $M' \sqsubset M$, we have $h \models_M \mathbf{t}_i \hat{q}$ as required. ■

That is, given non-trivial increasing awareness, the individual believes at h that, for at least some $h' \prec h$, the proposition $\exists q \in \mathbf{Q}(u_i q)$ was true. On the other hand, this proposition can never be known false. Hence, by historical induction, the individual judges that her awareness is incomplete.

5.1 Induction over individuals

We have shown that, given the interpretations of b , a and u adopted here, which is consistent with the framework of Heifetz, Meier and Schipper (2006), an individual i cannot believe that she may in the future become aware of propositions of which she is currently unaware. Similarly, individual i cannot believe that some other individual j is aware of propositions of which i is unaware. By contrast, individual i can believe that j is unaware of propositions of which i is aware.

This point may also be expressed in terms of false beliefs. The individual can believe that other individuals have false beliefs. But, within their own model, they cannot believe that they have false beliefs. In particular, i can impute to j a model in which j falsely believes that i is unaware of propositions of which j is aware. There is an obvious asymmetry here, about which boundedly

rational (but sophisticated) individuals may reason. Thus, i may reason about propositions such as ‘ j is aware of some p of which I am unaware’ and ‘my model of j ’s awareness is incomplete’.

The general principle of induction considered in the philosophical literature states that observations of members of some set S , all of which satisfy some property θ , provide inductive support for the proposition ‘All members of set S satisfy property θ ’. For example, if a number of ravens are observed to be black, and none are observed to be any other color we derive inductive support for the proposition ‘All ravens are black’.¹² In the application here, the sets to which induction is to be applied will consist of the set of individuals N_M , and the properties will involve awareness of propositions.

To formalize this idea in our notation, fix a model M . Let a propositional variable $q(j)$, containing the operators a_j and u_j , range over elements of $\hat{Q}_M$ such that $q(\tilde{j})$ is the proposition in which each instance of a_j in $q(j)$ is replaced by $a_{\tilde{j}}$ and each instance of u_j is replaced by $u_{\tilde{j}}$, for $\tilde{j} \in N_M$.

Definition 16 Fix a model M , an individual i and history $h \in H_M$, such that $\rho(h, i) = M$, and a propositional variable $q(j)$ in which

(i) for at least some $j \in N_M$, $h \models_M b_i q(j)$

(ii) for all $j \in N_M$, $h \not\models_M b_i \neg q(j)$

then, $h \models_M i_i (\wedge_{j \in N_M} q(j))$ [read as ‘at h individual i regards $q(j)$ as supported by induction over the set of individuals N_M ’].

Notice that, if $i_i (\wedge_{j \in N_M} q(j))$, then for any particular $j \in N_M$, $i_i q(j)$. That is, if a proposition is inductively supported for the set of individuals N_M , it is supported for any member of that set.

Proposition 9 Let $\rho(h, i) = M$. If M is not a model of common awareness then

(i) $h \models_M i_i (\exists p \in \mathbf{Q}, u_i p)$

(ii) $h \models_M i_i (\exists p \in \mathbf{Q}, j \in N_M, u_i p \wedge a_j p)$.

Proof. (i) Define $q(j) \equiv \exists p \in \mathbf{Q}, u_j p$. Since M is not a model of common awareness, $\exists \tilde{j} \in N_M$, such that $\rho_M(h, \tilde{j}) = \tilde{M} \sqsubset M$. Hence, for this $\tilde{j}$, $\exists p \in Q_M$, $h \models_M u_{\tilde{j}} p$. So (bearing in mind that $\rho_M(\tilde{j}, \cdot)$ is constant on the information set $I_j(h)$), we have $h \models_M b_i q(j)$. On the other hand, for any j (including i), we have by 7 $h \not\models_M b_j (\exists p \in \mathbf{Q}, u_j p)$. Hence, $h \models_M i_i (\exists p \in \mathbf{Q}, u_i p)$ as required.

(ii) Define $q(j) \equiv (\exists p \in \mathbf{Q}, j \in N_M, u_i p \wedge a_j p)$. As in (i), since M is not a model of common awareness, $\exists \tilde{j} \in N_M$, such that $\rho_M(h, \tilde{j}) = \tilde{M} \sqsubset M$, where we recall that M is the model held by i . Hence, for this $\tilde{j}$, $\exists p \in Q_M$, $h \models_M u_{\tilde{j}} p \wedge a_i p$. On the other hand, for any j, k (including i), we have, by Proposition 7, $h \not\models_M b_i (\exists p \in \mathbf{Q}, u_j p)$ and hence $h \not\models_M b_i (\exists p \in \mathbf{Q}, u_j p \wedge a_k p)$. ■

¹² As the famous example of black swans shows, inductive reasoning is never conclusive.

Part (i) of Proposition 9 is similar in its structure to Proposition 14. Given the belief on the part of i that some individuals are unaware of some propositions, and the absence of belief that any individual is fully aware of all propositions in $\mathbf{Q}$, inductive reasoning supports the proposition that all individuals, including i herself, are unaware of some propositions.

If the future and past selves are considered as other agents, there is a natural linkage between this idea and that of historical induction. The models held by the individual in the past were restrictions of their current model, and this proposition is known to the individual at h . Similarly, other individuals as modelled at h must have models that are restrictions of M .

Part (ii) takes the analysis a step further. In the case of a two-player game it may be seen as an application of symmetry between individuals. Suppose i believes that the game is one of differential awareness, in which there exist propositions and associated possible moves in the game, of which she is aware but j is not. If she is sophisticated enough to consider the implications of her own limited awareness, it is natural to entertain the converse possibility, that there exist propositions and associated possible moves in the game, of which j is aware but i herself is not.

In summary, although the question raised by Halpern and Rêgo (2006a) may be answered in the negative as regards the standard knowledge and awareness operators, inductive reasoning and symmetry arguments can provide a basis for a judgement that others may be aware of possibilities we have not considered.

5.2 Heuristic constraints and equilibrium

We develop a basis for making decisions that depend upon judgements about propositions that cannot be expressed explicitly in the language of the model of that player but are nonetheless supported either on historical inductive grounds or by induction over individuals. We refer to these to as heuristic constraints.

A heuristic constraint for a particular player associated with a model is a rule precluding the adoption by that player of an action whenever a certain proposition that involves the player taking that action is supported either on historical inductive grounds or by induction over individuals. For example, in the speculative trade example of Heifetz, Meier and Schipper (2006), we can construct a heuristic constraint based on the proposition, ‘my opponent is aware of something that I am not, and whatever it is, may result in me incurring a loss from trading with him.’ The corresponding heuristic constraint is not to engage in trade if the aforementioned proposition is supported either on historical inductive grounds or by induction over individuals.

Formally we define a heuristic constraint as follows.

Definition 17 Fix a model M , a history h in H_M . For the player $i = P_M(h)$, let $q_\alpha \in Q_M$ denote the proposition ‘ i takes action α at $\mathcal{I}_M^i(h)$ ’ and let $\theta(q_\alpha) \in \hat{Q}_M$ denote a Boolean combination of q_α and propositions in $\hat{Q}_M$. The heuristic constraint which precludes the adoption by player i of

the action $\alpha \in (A_M)_h$ at information set $\mathcal{I}_M^i(h)$ when the proposition $\theta(q_\alpha)$ is supported either on historical inductive grounds or by induction over individuals is denoted by:

$$NOT_M^i(\alpha : h \models_M \mathbf{t}_i\theta(q_\alpha) \vee \mathbf{i}_i\theta(q_\alpha)).$$

Imposing heuristic constraints on a game with differential awareness means that for any model in which heuristic constraints may apply, the game associated with that model is taken to be the restriction of the original game associated with that model obtained by deleting all histories that contain any actions that are precluded by the heuristic constraints. As the modified game is a restriction of the original game, this means that its set of terminal histories are a subset of those from the original game. This in turn implies that at any information set in the original game for which a heuristic constraint precludes the choice of at least one action, there is at least one other available action that is not precluded by any of the heuristic constraints that may apply. Putting this all together we define games with differential awareness subject to heuristic constraints as follows.

Definition 18 Fix $\mathcal{G}$, a game with differential awareness, and fix $\mathcal{H}$ a set of heuristic constraints to which models in $\mathcal{G}$ may be subject. We denote by $\hat{\mathcal{G}} = (\mathcal{G}, \mathcal{H})$, the game with differential awareness subject to heuristic constraints, in which for each model $M \in \mathcal{G}$, is associated the game $\hat{\Gamma}_M \sqsubseteq \Gamma_M$, in which history $h'' \in H_M$ is not in $\hat{H}_M$, if there exists a history $\tilde{h}$ in H_M , such that for $i = P_M(\tilde{h})$

- (i) the heuristic constraint $NOT_M^i(\alpha : \tilde{h} \models_M \mathbf{t}_i\theta(q_\alpha) \vee \mathbf{i}_i\theta(q_\alpha))$ is in $\mathcal{H}$;
- (ii) $\tilde{h} \models_M \mathbf{t}_i\theta(q_\alpha) \vee \mathbf{i}_i\theta(q_\alpha) \rho_M(h, i)$;
- (iii) $h' \preceq h''$, where $h' = \tilde{h}' \cdot \langle \alpha \rangle$, for some $\tilde{h}' \in \mathcal{I}_M^i(\tilde{h})$; and,
- (iv) there exists $\hat{\alpha}$ in $(A_M)_{\tilde{h}}$ such that no heuristic constraint satisfying (i)-(iii) holds for $\hat{\alpha}$.

Condition (i) states the presence of the heuristic constraint. Condition (ii) states that the required conditions for inductive justification of $\theta(q_\alpha)$ hold at $\tilde{h}$ and condition (iii) requires that the history h'' passes through $\tilde{h}' \cdot \langle \alpha \rangle$, for some $\tilde{h}' \in \mathcal{I}_M^i(\tilde{h})$. Finally, condition (iv), the existence requirement, ensures that the set of actions available at $\tilde{h}$ is non-empty.

We revise the definition of a strategy profile, so that in the model M in which heuristic constraints may apply, the strategy profile is defined for the game $\hat{\Gamma}_M$. Thus, our recursive definition of a strategy profile is modified as follows.

Definition 19 Let $\hat{\mathcal{G}}$ be a game with differential awareness and with models subject to heuristic constraints. A **strategy profile** $\hat{\beta} = (\hat{\beta}_M : M \in \hat{\mathcal{G}})$ for $\hat{\mathcal{G}}$ assigns to each model $M \in \hat{\mathcal{G}}$ a behavioral strategy profile $\hat{\beta}_M$ for the game $\hat{\Gamma}_M \sqsubseteq \Gamma_M$, with the consistency property: if for some $h \in \hat{H}_M$, where $P(h) = i$, if $\rho_M(h, i) = \hat{M}$, then $\beta_M^i(\mathcal{I}_M^i(h)) = \hat{\beta}_{\hat{M}}^i(\mathcal{I}_{\hat{M}}^i(h))$.

The definitions of belief systems, assessments and sequential rationality from section 2.4.1 can be modified accordingly with each instance of the game Γ_M for model M , being replaced by the restricted game $\hat{\Gamma}_M$ that excludes all histories that are precluded by the adoption of the heuristic constraints. Thus we can extend the definition of a sequential equilibrium to apply to games with differential awareness subject to heuristic constraints. And as an immediate corollary to Proposition 4, we have:

Corollary 10 *A sequential equilibrium exists for any game with differential awareness subject to heuristic constraints.*

5.3 Example Part 2:

We now return to the speculative trade example of Heifetz, Meier and Schipper (2006). By Proposition 9, for each party $i = 1, 2$, the proposition $\exists q \in \mathbf{Q}(u_j q \wedge a_i q)$ is true for $j \neq i$. Moreover, in the given example, the proposition can be extended to

$$\theta(q_{\alpha^i}) \equiv \exists q \in \mathbf{Q}\left(u_j q \wedge a_i q \wedge \left(q \Rightarrow q_{loss}^j\right)\right), \quad (3)$$

where α^j is the action necessary for j to take in order to transact with i and q_{loss}^j is the proposition the transaction produces a loss for j .¹³

Now let party i consider the proposition $\exists q \in \mathbf{Q}(u_i q \wedge a_j q)$, that is, that there exists a proposition of which she is unaware, but the other party is aware. As shown above, inductive reasoning, embodying the idea of symmetry, provides support by induction over the set of individuals for the existential proposition $\exists q \in \mathbf{Q}(u_i q \wedge a_j q \wedge (q \Rightarrow q_{loss}^i))$, so we have for both $h^1 = \langle \alpha_0, \alpha_0 \rangle$ and $h^2 = \langle \alpha_0, \alpha_0, \alpha_1 \rangle$,

$$h^i \models_{M^i} i; \theta(q_{\alpha^i}).$$

Hence the heuristic constraint

$$\mathcal{H}^i = NOT(\alpha^i : h \models_M i; \theta(q_{\alpha^i}))$$

precludes the adoption of α^i by party i . That is, the decision by i to participate in the transaction is precluded by this heuristic principle. Hence, if either the buyer adopts the principle $\mathcal{H}^1$ or the owner adopts the principle $\mathcal{H}^2$, the transaction will not take place.

Now compare the dominance principle proposed by Heifetz, Meier and Schipper (2006). Heifetz, Meier and Schipper propose that if (i) in all histories h' an agent considers possible at h , action α leads to at least as good an outcome as α' , and (ii) in some possible history, action α leads to a

¹³ Recall in the specification of the game that for party 1 the buyer, α^1 corresponds to the action ‘make an offer of 1’ (action α_1 in figures 1–4) and for party 2, the owner, α^2 is the action ‘accept the offer of 1’ (action α_A in figures 1–4).

better outcome, then the agent should choose α . As shown above, this principle leads the players to engage in trade, even though the trade is zero-sum.

In the given example, this principle does not appear compelling. Even though this condition is satisfied for the modal-logical interpretation of ‘considers possible’ (namely $h' \in \mathcal{I}^i(h)$), observation of the limited awareness of other individuals, combined with principles of symmetry between individuals, provides inductive support for the proposition in equation (3). Note that, evaluated in the maximal model, (3) is in fact true for both individuals.

6 Related literature

Our construction of the game with differential awareness $\mathcal{G}$ is close to that of Halpern and Rego (2006b, 2007) and Heifetz, Meier and Schipper (2006, 2009). Although we follow Halpern and Rego (2006b) in using an Osborne–Rubinstein setup for the games Γ_M , our recursive construction of $\mathcal{G}$ is closer to the approach adopted by Heifetz, Meier and Schipper (2008).

The associated family of languages $\{Q_M : M \in \mathcal{G}\}$ is related to the corresponding set of games $\{\Gamma_M : M \in \mathcal{G}\}$ in a way that is broadly similar to the syntax of awareness developed by Halpern and Rego (2008). Moreover, our awareness operator satisfies the properties postulated by Feinberg (2005).

However, our construction of $\hat{Q}_M$ represents an approach to reasoning about unawareness quite different from that proposed by Halpern and Rego (2007) and endorsed by Heifetz, Meier and Schipper (2008). This difference in turn reflects the fact, that these authors are primarily concerned with the extension of Bayes–Nash solution concepts for games of common awareness to the case of games of differential awareness. In this setting, it is natural to propose, as do Halpern and Rego (2007), that awareness of one’s own unawareness should be modelled by incorporating ‘virtual moves’ in the game tree to represent possibilities of which one is unaware. Thus, at each history in the model, all players act as if they (though perhaps not all other players) are fully aware of the game, as represented by the augmented structure incorporating virtual actions. In their equilibrium choices, players make no distinction between the moves in the maximal game of which they are aware and the additional virtual moves that have been added to take account of their bounded awareness.

By contrast, we are concerned with the idea that consciousness of one’s limited awareness may reasonably be seen as a constraint on the adoption of Bayes–Nash reasoning and solution concepts. That is, if an individual does not believe that they have considered all relevant possibilities, they may be unwilling to commit to a path of action that would be optimal if their awareness were complete. As we show in Grant and Quiggin (2008), unawareness does not preclude the possibility, under appropriate ‘small world’ conditions, that it may be appropriate to apply Bayes–Nash reasoning to particular decision problems. However, as has been shown here, there is, in general,

an important role for inductive modes of reasoning that cannot simply be reduced to Bayesian decision theory.

7 The precautionary principle

The precautionary principle, presented as a guide to environmental policy decisions in the presence of uncertainty, has been the subject of vigorous debate (Wingspread 1998). However, discussion of the principle as a decision-theoretic rule has mostly relied on the (normally implicit) assumption that decision makers are unboundedly rational and aware of all possible contingencies. In this context, the precautionary principle has been criticized as involving inconsistency (Marchant and Mossman 2005, Sunstein 2005) or excessive risk aversion (Miller and Conko 2005) and defended as a way of capturing option value (Gollier, Jullien and Treich 2000). It is evident, however, that in a fully specified decision-theoretic model, with all contingencies taken into account, and an appropriately specified objective function, there should be no need for additional heuristic rules such as those of the precautionary principle.

When the limited awareness of participants in decision processes is taken into account, however, the precautionary principle seems more appealing. Given the bounded rationality of human agents, it is impossible to enumerate all relevant possibilities. This point is sometimes expressed with reference to ‘unknown unknowns’, that is, relevant possibilities of which we are unaware.

The case for the precautionary principle arises when a decisionmaker, such as a regulator, is faced with a choice between alternatives, one of which leads to consequences for which the relevant elements of the state space are well understood and the other which leads to consequences that depend to a significant extent on ‘unknown unknowns’. If most surprises are unpleasant, a risk analysis based only on known risks will underestimate the costs of choices of the second kind. That is, standard risk analysis leads to a bias in favour of taking chances on poorly-understood risks. The precautionary principle may be seen as a rule designed to offset such biases.

In a multi-agent context, regulatory decisions typically involved assessment of proposed actions seen as raising possible risks. In this context, the precautionary principle may be understood, as a procedural constraint, putting the burden of proof on to proponents of decisions involving poorly-understood risks. If the proponent can provide sufficient information to satisfy the regulator that all relevant contingencies have been considered, standard principles of decision analysis may be applied to justify a proposal. If not, the regulator may choose to apply the precautionary principle and reject the project even in the absence of a negative benefit–cost evaluation.

8 Concluding comments

We have shown that, under plausible conditions induction from past experience, or from observation of others, supports the belief that there exist possibilities of which we are unaware. Given that belief, we may choose to replace or supplement Bayes–Nash reasoning with heuristic rules of behavior. Once this possibility is admitted, it is evident that the decision as to whether to apply Bayes–Nash reasoning or heuristic rules in a given instance must itself be governed by heuristic considerations (since otherwise the entire process would collapse to a more complicated version of Bayes–Nash). In this context, our approach may be interpreted as an attempt to address the question of when it is (in)appropriate for boundedly rational individuals to adopt Bayes–Nash reasoning (see also Grant and Quiggin 2008)

Where players consider that they are fully aware of all relevant moves, the model’s predictions correspond with those of the standard Bayes–Nash equilibrium, after taking account of the possible unawareness of other players of moves available to them. On the other hand where players judge, for example on the basis of induction from past experience, that there may exist actions they have not considered, our model allows for a wide range of possible responses.

References

- Bonanno, G. (2001), ‘Branching time logic, perfect information games and backward induction’, *Games and Economic Behavior*, 36(1), 57–73.
- Dekel, E., Lipman, B. and Rustichini, A. (1998), ‘Standard state-space models preclude unawareness’, *Econometrica*, 66, 159–173.
- Feinberg, Y. (2005), ‘Games with incomplete awareness’, Stanford University
- Grant, S. and Quiggin, J. (2008), ‘Bounded rationality and small worlds’, Risk and Uncertainty Program Working Paper WP2R08, Risk and Sustainable Management Group, University of Queensland.
- Grant, S. and Quiggin, J. (2006), ‘Learning and Discovery’, 12th International Conference on the Foundations and Applications of Utility, Risk and Decision Theory (FUR XII 2006), Rome, 26 June.
- Gollier, C., Jullien, B. and Treich, N. (2000), ‘Scientific progress and irreversibility: An economic interpretation of the “precautionary principle”’, *Journal of Public Economics*, 75, 229–53.
- Halpern, J. (2003), *Reasoning About Uncertainty*, The MIT Press, Cambridge, Massachusetts.
- Halpern, J. and Rego, L. (2006a), ‘Reasoning about knowledge of unawareness’, Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2–5.

- Halpern, J. and Rego, L. (2006b), ‘Extensive games with possibly unaware players’, Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, Hakodate, Japan, May 8-12.
- Halpern, J. and Rego, L. (2007), ‘Generalized solution concepts in games with possibly unaware players’, Proceedings of the Eleventh Conference on Theoretical Aspects of Rationality and Knowledge, Brussels, June 25-27.
- Halpern, J. and Rêgo, L. (2008), ‘Interactive unawareness revisited’, *Games and Economic Behavior*, 62(1), 232–262.
- Heifetz, A., Meier, M. and Schipper, B. (2006), ‘Interactive unawareness’, *Journal of Economic Theory*, 130(1), 78–94.
- Heifetz, A., Meier, M. and Schipper, B. (2009), ‘Dynamic unawareness and rationalizable behavior’, Working Paper, UC Davis..
- Laudan, L. (1981), ‘A confutation of convergent realism’, *Philosophy of Science*, 48(1), 19–49.
- Marchant, G.E. and Mossman, K.L. (2005) *Arbitrary and Capricious: The Precautionary Principle in the European Union Courts*, International Policy Press a division of International Policy Network, London.
- Miller, H. and Conko, G. (2005), ‘The UN, biotechnology, and the poorest of the poor’, *Hoover Digest*, 2005(2), <http://www.hoover.org/publications/digest/2993011.html>.
- Osborne, M.J. and Rubinstein, A. (1994), *A Course in Game Theory*, MIT Press, Cambridge, Mass.
- Sunstein, C.R. (2005) *Laws of Fear: Beyond the Precautionary Principle (the Seeley Lectures)*, Cambridge University Press,
- Wingspread Conference (1998), ‘Wingspread Statement on the Precautionary Principle’, Press release, Racine, Washington, February.

Figure 1: The maximal game Γ which starts at the chance node in the center.

Figure 2: Γ^1 , the Game the Buyer Perceives He is Playing, which starts at the chance node in the middle.