

2012

Risk and Uncertainty Program

**TITLE: Inductive Reasoning About
Unawareness**

AUTHORS:

Simon Grant and John Quiggin

Working Paper: R12_3

**RSMG Working
Paper Series**

**Schools of
Economics and
Political Science**

**The University of
Queensland**

St Lucia

Brisbane

Australia 4072

Web:

www.uq.edu.au

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 differential awareness. We show that, given differential awareness over time and between players, individuals can derive inductive support for propositions expressing their own unawareness. We consider the ecological rationality of heuristics to guide decisions in problems involving differential awareness.

JEL Classification: D80, D82

Key words: unawareness, bounded rationality, induction

Simon Grant
School of Economics, University of Queensland, Brisbane

John Quiggin
School of Economics, University of Queensland, Brisbane
j.quiggin@uq.edu.au

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

1 Introduction

As former US Secretary Donald Rumsfeld (2002) famously observed, the problem of unknown unknowns (things we don't know we don't know) is one of the most difficult facing any decision-maker. In standard decision-theoretic frameworks, the set of possible states of nature is known at the beginning of the problem. Learning consists of observing signals that restrict the set of possible states. Probabilities are then updated according to Bayes' rule. In reality, however, decisionmakers are frequently presented with 'surprises', that is, events they had not previously considered. So, a realistic model of choice under uncertainty must incorporate the fact that individuals are unaware of some relevant possibilities. Similarly, in a game-theoretic context, at any given stage in an extensive-form game, some players may be unaware of possible moves of which other players are aware. Furthermore, sophisticated individuals will understand that there may be possibilities of which they are unaware, even though they cannot express this understanding within the state-space model of the world available to them.

The argument of this paper is twofold. First, we show that, even when standard game-theoretic and decision-theoretic models are extended to include unawareness and differential unawareness, these standard concepts of belief, knowledge and awareness cannot encompass the idea that individuals understand their own bounded awareness. This lack of self-awareness persists even in a dynamic and interactive setting, where individuals are aware of both their own past unawareness and the bounded awareness of other individuals. We conclude that a reasonable model of individuals with a sophisticated understanding of their own bounded awareness must incorporate modes of reasoning other than deductive inference based on a fully specified state space or dynamic game tree.

Second, we argue that, in a dynamic interactive setting, it is natural to employ inductive reasoning to justify inferences about one's own awareness and unawareness. In particular, since everyone has the experience of becoming aware of propositions and possibilities they have not previously considered, standard principles of historical induction suggest that similar experiences will occur in the future, and therefore that there exist propositions of which we are currently unaware. Similarly, in a game with awareness, described by a syntactic structure, individuals believe that others are unaware of at least some possible histories and the associated propositions. Inductive reasoning suggests that the same will be true of all individuals, including the person drawing the inference.

It follows that reasonable, but boundedly rational, individuals, should not rely solely on standard Bayes-Nash reasoning to guide their decisions, whether these decisions involve games with Nature or interactions with other individuals. Rather, individuals may improve their welfare by adopting inductively justified heuristics that are, in the terminology of Gigerenzer and Goldstein

(2002), ecologically rational.

The paper is organized as follows:

Section 2 presents semantic and syntactic representations of differential awareness in terms of extensive games. The syntactic representation is constructed by associating with any given extensive form game a propositional language rich enough to specify all histories and sets of histories that arise in that game. The crucial idea in our representation of games with awareness is that players may be unaware of some possible histories of the game, and may therefore have access only to a restricted version of the game. This leads to an important modification of the standard modal-logical approach, in which a proposition is known to be true by a player if it is true in all histories considered possible by that player, given the information she has observed. In a game with awareness, a proposition may be false in the actual history, but true in all histories considered possible by a player, given their limited awareness. From the perspective of the full game (or, more generally, from the perspective of a player with greater awareness), it makes sense to characterize the player's attitude to such propositions as 'belief' rather than 'knowledge'. We show how knowledge and belief operators may be defined in a game with differential awareness, in a way that allows for false belief but not false knowledge. These ideas are illustrated with reference to an example first presented by Heifetz, Meier and Schipper (2006).

In Section 3 we consider the question of how individuals can reason about their own awareness and unawareness. First, 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. Within the model of games with awareness developed in Section 2, we show that 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 with awareness, it will be true, in general, that there exist propositions of which players are unaware. Furthermore, players in games will, in general, experience the discovery of propositions of which they were previously unaware, and observe the bounded awareness of other players.

These observations lead us to consider modes of reasoning other than the Bayesian inference that characterizes standard extensive form games. We show how inductive modes of reasoning may be used by individuals to assess existential propositions about awareness. 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 that holds true for at least some individuals, and is not false for any individual is supported by induction over individuals. In a game

with awareness, players believe that others are unaware of at least some propositions. Inductive reasoning suggests that the same will be true of all players, including the person undertaking the induction.

Finally, we consider the implications of inductive reasoning about unawareness for decisions, represented in general by the choice of strategy in an extensive form game. In Section 5, 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 provide criteria under which the adoption of heuristic constraints may be ecologically rational from the perspective of the full game. The analysis is illustrated by a no-trade result for the speculative trade example developed previously. We discuss possible applications to the precautionary principle, often advocated as a basis for regulatory decisions regarding environmental risks, and to decisions regarding research and discovery.

Finally, we offer some concluding comments.

2 Games with awareness

2.1 Extensive-form games and languages: notation

In this section, we describe the notation for games and the associated languages. We will use the term ‘semantic’ to refer to the representation of the problem in terms of possible game histories, and ‘syntactic’ to refer to the representation in terms of the truth values of propositions. For standard games, there is a 1-1 correspondence between semantic and syntactic representations, including temporal structures, information and beliefs.

We begin with a(n almost) standard definition of an extensive-form game, as in Osborne and Rubinstein (1994).

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

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

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

G3 (Histories): H is a set of histories, defined as sequences of actions undertaken by the players. H is partially ordered by the subhistory relationship \preceq . The set of terminal histories is denoted Z . The set of actions available at h is denoted $A(h) \subseteq A$.

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

G5 (Chance Assignment): f^c associates with every history h such that $P(h) = c$ a probability distribution over A drawn from some set Δ , and with support $A(h)$.

G6 \mathcal{I} : $H \rightarrow 2^H$ is the information set assignment function whose range forms a partition of H and exhibits the properties that $h \in \mathcal{I}(h)$ for all $h \in H$, and for any $h' \in \mathcal{I}(h)$, $P(h) = P(h')$ and $A(h) = A(h')$.

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

So, the set of histories h is the set of all sequences of the form $\langle \alpha_1, \alpha_2, \dots, \alpha_k \rangle$, where $\alpha_j \in A(\langle \alpha_1, \alpha_2, \dots, \alpha_{j-1} \rangle)$, $j = 1, \dots, k$, including the trivial sequence $\langle \cdot \rangle$. The extensive structure of the game Γ is represented by the sub-history relationship $\langle \alpha_1, \dots, \alpha_k \rangle \preceq \langle \alpha_1, \alpha_2, \dots, \alpha_k, \dots, \alpha_l \rangle$. If $h = \langle \alpha_1, \alpha_2, \dots, \alpha_k \rangle$, we denote $h \bullet \alpha_{k+1} = \langle \alpha_1, \alpha_2, \dots, \alpha_k, \alpha_{k+1} \rangle$ and observe $h \preceq h \bullet \alpha_{k+1}$.

The process of Bayesian learning in an extensive form game works by exclusion. That is, as the game is played out, players learn that some histories in the game are no longer possible, and update their probabilities on the remaining possible histories using Bayes' rule. The information set $\mathcal{I}(h)$ describes the set of histories that have not been ruled out by the information available to player $P(h)$ at h .

A behavioral strategy β^i for player i is a collection of independent probability measures $(\beta^i(h) : P(h) = i)$ where the support of each such $\beta^i(h)$ is a subset of $A(h)$, and where for any $h' \in \mathcal{I}(h)$, $\beta^i(h') = \beta^i(h)$. That is, there is an independent probability measure over actions specified for each history at which player i is called upon to play and this probability measure is the same for all histories residing in the same information set. A behavioral strategy profile β is a set of behavioral strategies, one for each i .

We define a subjective probability $\mu^i = (\mu^i(h) : P(h) = i)$ for player i to be an assignment to each history h in H for which $P(h) = i$, a probability measure on the set of histories in $\mathcal{I}(h)$ in which for any $h' \in \mathcal{I}(h)$, $\mu^i(h') = \mu^i(h)$.¹ A subjective probability system μ , is a set of subjective measures, one for each i . An assessment is a pair (β, μ) , where β is a behavioral strategy profile and μ is a subjective probability system.

Given a strategy profile β , for each h in H we denote by $\tau(\beta|h)$ the probability distribution over terminal histories induced by β , conditional on history h having been reached. Let $\tau(\beta)$ denote the unconditional probability distribution over terminal histories induced by β .

Hence given the assessment (β, μ) , the continuation value of player $i = P(h)$ at h (that is, the continuation value conditional on information set $\mathcal{I}(h)$ having been reached) is

$$V_{\Gamma}^i(h; \beta, \mu) = \sum_{h' \in \mathcal{I}(h)} \mu^i(h)[h'] \left(\sum_{h'' \in Z} \tau(\beta|h')[h''] v^i(h'') \right),$$

¹ Osborne and Rubinstein use the term 'belief system'. Our terminology has been chosen to avoid confusion with the belief operator to be defined below.

and we denote the expected value of the (entire) game for player i by $V_{\Gamma}^i(\beta) = \sum_{h \in Z} \tau(\beta)[h] v^i(h)$.

An assessment (β, μ) is *sequentially rational* if each player's strategy is a best response at every information set at which she is called upon to play. That is, we require, for each player $i \in N$ and every h , such that $P(h) = i$,

$$V_{\Gamma}^i(h; \beta, \mu) \geq V_{\Gamma}^i\left(h; \left(\hat{\beta}^i, \beta^{-i}\right), \mu\right), \text{ for every strategy } \hat{\beta}^i \text{ of player } i.$$

We refer to a behavioral strategy profile that assigns positive probability to every information set as completely mixed. An assessment is *consistent* if it is the limit of a sequence $((\beta_m, \mu_m))_{m=1}^{\infty}$ in which each strategy profile β_m is completely mixed and each belief system μ_m is derived from β_m using Bayes' rule. A *sequential equilibrium* is a consistent and sequentially rational assessment. As in Osborne and Rubinstein, it is straightforward to demonstrate the existence of a sequential equilibrium.

We now consider a syntactic rendition of the same structure. With each game Γ , we associate a language \mathcal{L}_{Γ} . The language must be rich enough to encompass the sequential structure of H , the information in \mathcal{I} and the valuations v^i . The relationship between the properties of the language \mathcal{L}_{Γ} and of the game Γ may be formalized using the semantic approach to modal logic first presented by Kripke (1963) and developed in relation to the logic of knowledge by Fagin and Halpern (1988).

Space does not permit a detailed exposition here, but a brief outline will be useful. The central idea is to use properties of the game Γ to define (in Kripke's terminology) accessibility relations between histories $h \in H$. More precisely, the property of common membership of an information set defines an equivalence relation while the temporal structure defines a partial ordering. These relations define a class of Kripke frames for which an appropriate axiomatization of the language \mathcal{L}_{Γ} can be shown to be complete (permitting derivation of all theorems applicable to games Γ) and sound (ensuring that every proposition derivable in \mathcal{L}_{Γ} is valid in Γ).

Definition 2 For a game Γ , the game language \mathcal{L}_{Γ} is a set of sentences closed under the logical operators \wedge and \neg , and derived from:

L1 (Player terms): Terms $\mathcal{N} = \{p_i : i \in N\}$ representing players, with c denoting Nature (the 'chance' player), where p_i is read as ' i is a player in the game';

L2 (Actions): A set of elementary propositions $\mathcal{A} = \{p_{\alpha} : \alpha \in A\}$ where p_{α} is read as 'action α has just been taken'

L3 (Histories): A set of elementary propositions $\mathcal{H} = \{p_h : h \in H\}$ read as 'the current history is h '.

L4 (Player Assignment) A set of elementary propositions $\mathcal{P} = \{p_{h,i} : h \in H, i \in N\}$ where $p_{h,i}$ is read as 'player i moves at h '

L5 (Probabilities): A set of probability propositions $\{p_{\pi} : \pi \in \Delta\}$ where Δ is the probability simplex

over A and p_π is read as ‘the probability distribution over A is given by $\pi \in \Delta$ ’

L6 (Knowledge operators): $\{k_i p : i \in N\}$ read as ‘ i knows p ’

L7 (Payoffs): For each player $i \in N$, and $v \in V$, an elementary proposition $p_{v,i}$ read as ‘player i receives payoff v ’. We will assume, without details, that the language is rich enough to allow players to do standard arithmetic, for example, to express propositions like ‘my expected payoff is less than v ’

L8 (Temporal structure) The temporal structure of the game is given by an operator w (where $w p$ is read as ‘ p was true in the past’)

Given an appropriate semantic interpretation, the language \mathcal{L}_Γ represents a syntactic rendition of the game Γ . That is, for any given play of the game, the truth value of any proposition $p \in \mathcal{L}_\Gamma$ can be inferred. Conversely, the structure of the game can be inferred from the set of propositions in \mathcal{L}_Γ that are tautologically true, and any set of truth values consistent with those tautologies determines a play of the game.

The semantic interpretation works as follows. For any $p \in \mathcal{L}_\Gamma$, the statement ‘ p is true at h ’ is written $h \models_\Gamma p$. The relation \models_Γ can be derived, using the standard rules of logic and the interpretation rules given below:

Definition 3 A semantic-syntactic game representation $(\Gamma, \mathcal{L}_\Gamma, \models_\Gamma)$ consists of an extensive form game Γ the associated language \mathcal{L}_Γ and a relation \models_Γ such that, for all p, h , either $h \models_\Gamma p$ or $h \models_\Gamma \neg p$, and:

S1 (Players): $h \models_\Gamma p_i, \forall h \in H, i \in N$

S2 (Actions): $h \models_\Gamma p_\alpha \Leftrightarrow \exists h', s.t. h = h' \bullet \alpha$

S3 (History): $\forall h, h' \in H, (i) h \models_\Gamma p_h, (ii) h' \neq h \Leftrightarrow h' \models_\Gamma \neg p_h, (iii) h \models_\Gamma w p_h \Leftrightarrow h' \preceq h$

S4 (Player Assignment) $h \models_\Gamma p_{h,i} \Leftrightarrow P(h) = i$

S5 (Probability): Given a strategy profile β , $h \models_\Gamma p_\pi$ iff β assigns π to A at h

S6 (Knowledge): $h \models_\Gamma k_i p$, if and only if $P(h) = i$ and $h' \models_\Gamma p$ for all $h' \in \mathcal{I}(h)$;

S7 (Payoffs): $h \models_\Gamma p_{v,i} \Leftrightarrow h \in Z \wedge v^i(h) = v$

S8 (Temporal structure) $h \models_\Gamma w p \Leftrightarrow \exists h' \prec h, h' \models_\Gamma p$

Of these rules, that of the knowledge operator k_i is the only one that requires special attention. It states that, at the history h in the game, player i knows the proposition p is true if and only if p is true in all histories considered possible by player i at h . Thus, the semantic interpretation of the language relates knowledge directly to the information sets of players in the game.

It is straightforward to show that the knowledge operator satisfies the standard set of axioms referred to in the literature on modal logic as S5. The most important of these are Necessitation $p \Rightarrow k_i p$, commonly denoted K, Distribution $k_i(p \Rightarrow q) \Rightarrow (k_i p \Rightarrow k_i q)$, commonly denoted D,

Truth ($k_i p \Rightarrow p$), commonly denoted T, Positive Introspection ($k_i p \Rightarrow k_i k_i p$), commonly denoted 4, and negative introspection, ($\neg k_i p \Rightarrow k_i \neg k_i p$), commonly denoted 5. Fagin et al (1985) show that, in a ‘possible worlds’ representation of knowledge, S5 is a complete and sound axiomatization for the knowledge operator defined as above. However, this result does not encompass the temporal structure of knowledge in an extensive form game, where the notion corresponding to the set of worlds is the set of partial histories at which a given player is to move. We conjecture that a sound and complete axiomatization could be obtained by combining the S5 axioms for knowledge, a suitable set of axioms for the temporal operator w and a ‘perfect recall’ axiom, of the form $wk_i p \Rightarrow k_i w k_i p$. However, this question is beyond the scope of the present paper.

We can now consider 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. That is, $a_i p \Leftrightarrow k_i p \vee k_i \neg p \vee k_i (\neg k_i p \wedge \neg k_i \neg p)$ with unawareness being the negation of awareness, that is, $u_i p$ is a synonym for $\neg a_i p$. (Notice that $k_i p \vee k_i \neg p \vee k_i (\neg k_i p \wedge \neg k_i \neg p)$ may be stated more compactly as $k_i p \vee k_i (\neg k_i p)$).

As observed by Modica and Rustichini (1994), for a partitioned information structure (which we have assumed), $a_i p$ is true for all $p \in \mathcal{L}_\Gamma$. Also we have $a_j a_i p$ and so forth. Because of this property we refer to a standard extensive form game as a game of full common awareness. Since $u_i p$ is trivially false in a game of full common awareness, we will not use the definition above, but will define awareness and unawareness in the context of a game with differential awareness, which we now construct.

2.2 Restrictions

Under the standard assumptions of unbounded rationality and common knowledge, all players in a game are aware of the structure of the game, of each other’s awareness, of others’ awareness of their own awareness, and so on. With boundedly rational players, however, 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 the less than full awareness of the player $P(h)$ at h by ascribing to her a game that is a *restriction* of the full awareness game. Essentially, the restriction is obtained by deleting some of the terminal histories and all partial histories of those terminal histories of the original game and then constructing the restricted game to be consistent with the original game in terms of its information structure.

Definition 4 Fix a game $\Gamma = (N \cup \{c\}, A, H, P, f^c, \mathcal{I}, \{v^i : i \in N\})$, where $Z \subset H$ denotes the set of terminal histories in Γ . A non-empty subset of terminal histories $\tilde{Z} \subset Z$ is deemed admissible, if $\Gamma_{\tilde{Z}} = (N_{\tilde{Z}} \cup \{c\}, A_{\tilde{Z}}, H_{\tilde{Z}}, P_{\tilde{Z}}, f_{\tilde{Z}}^c, \mathcal{I}_{\tilde{Z}}, \{v_{\tilde{Z}}^i : i \in N_{\tilde{Z}}\})$ constitutes a restriction of the game Γ in

which:

$$\begin{aligned}
H_{\tilde{Z}} &= \{h \in H : \exists z \in \tilde{Z}, h \preceq z\}; \\
N_{\tilde{Z}} &= \{n \in N : \exists h \in H_{\tilde{Z}}, P(h) = n\}; \\
\text{for each } h &\in H_{\tilde{Z}}: \\
\mathcal{I}_{\tilde{Z}}(h) &= H_{\tilde{Z}} \cap \mathcal{I}(h); \\
A_{\tilde{Z}}(h) &= \{a \in A(h) : h \bullet a \in H_{\tilde{Z}}\}; \\
P_{\tilde{Z}}(h) &= P(h); \\
f_{\tilde{Z}}^c(h)(a) &= \frac{f^c(h)(a)}{\sum_{\tilde{a} \in A_{\tilde{Z}}(h)} f^c(h)(\tilde{a})}, \\
&\text{and,} \\
\text{for each } i &\in N_{\tilde{Z}} \text{ and } z \in \tilde{Z}, v_{\tilde{Z}}^i(z) = v^i(z).
\end{aligned}$$

We denote the relation that $\Gamma_{\tilde{Z}}$ is a restriction of the game Γ by $\Gamma_{\tilde{Z}} \sqsubseteq \Gamma$.

Furthermore, in a slight abuse of notation, we shall also extend the domains of $A_{\tilde{Z}}$, $P_{\tilde{Z}}$, $f_{\tilde{Z}}^c$, $\mathcal{I}_{\tilde{Z}}$ to H , by setting for each h in $H - H_{\tilde{Z}}$:

$$\begin{aligned}
\mathcal{I}_{\tilde{Z}}(h) &:= H_{\tilde{Z}} \cap \mathcal{I}(h); \quad A_{\tilde{Z}}(h) := \{a \in A(h) : \exists h' \in \mathcal{I}_{\tilde{Z}}(h) \text{ s.t. } h' \bullet a \in H_{\tilde{Z}}\}; \quad P_{\tilde{Z}}(h) := P(h); \text{ and,} \\
f_{\tilde{Z}}^c(h)(a) &:= \begin{cases} f^c(h)(a) / \sum_{\tilde{a} \in A_{\tilde{Z}}(h)} f^c(h)(\tilde{a}) & \text{if } A_{\tilde{Z}}(h) \neq \emptyset \\ 0 & \text{if } A_{\tilde{Z}}(h) = \emptyset \end{cases}.
\end{aligned}$$

By construction the subhistories and the subhistory ordering are preserved in the sense that, for any $h, \tilde{h} \in H$, such that $\tilde{h} \preceq_{\Gamma} h$, if h is in the restricted set of histories $H_{\tilde{Z}}$ then \tilde{h} must also be in $H_{\tilde{Z}}$ and furthermore we retain $\tilde{h} \preceq_{\Gamma_{\tilde{Z}}} h$ in the restricted game. In addition, 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} . The definition of the chance assignment function $f_{\tilde{Z}}^c$ ensures that if an action by nature is excluded from consideration at h , the relative probabilities of the remaining actions $\alpha \in A_{\tilde{Z}}(h)$ are unchanged. Notice it also follows from the definition that a restricted game does not add information, or lose information with respect to the histories in the restricted game.

Remark 1 This relation \sqsubseteq is a partial ordering on the set of games, corresponding to the subset ordering on sets of terminal histories. In particular, if $\tilde{\tilde{Z}} \subseteq \tilde{Z} \subseteq Z$, then $\Gamma_{\tilde{\tilde{Z}}} \sqsubseteq \Gamma_{\tilde{Z}} \sqsubseteq \Gamma$.

A parallel definition applies for the associated language \mathcal{L}_{Γ} , and to the interpretation \models_{Γ} . If $\Gamma_{\tilde{Z}} \sqsubseteq \Gamma$, then $\mathcal{L}_{\Gamma_{\tilde{Z}}} \sqsubseteq \mathcal{L}_{\Gamma}$. For elementary propositions $p \in \mathcal{L}_{\Gamma_{\tilde{Z}}}$ (those incorporating only terms about players, actions and histories in the restricted game $\Gamma_{\tilde{Z}}$), we have, whenever $h \in \tilde{H} \subseteq H$, $h \models_{\Gamma_{\tilde{Z}}} p$ if and only if $h \models_{\Gamma} p$. Things are different when we come to consider knowledge and belief. It is possible to have a situation where $h' \models_{\Gamma_{\tilde{Z}}} p$ for all $h' \in \mathcal{I}_{\tilde{Z}}(h)$, but nevertheless,

$h \models_{\Gamma} \neg p$. That is, p is true for all the histories considered possible by player $i = P(h)$ at $\mathcal{I}_{\tilde{Z}}(h)$ in the restricted game $\Gamma_{\tilde{Z}}$, but that the restricted information set does not include the actual history h , at which p is false.

Thus, in a restricted game, players may hold false beliefs about propositions p . More generally, what appears to the player as reliable knowledge about a proposition p may be an unreliable belief if there are histories h'' that cannot be ruled out on the basis of the information available at h , of which the player is unaware and at which p is false. We write ' $b_i p$ ' read as ' i believes p ' and define what it means for a player to believe a proposition is true in the game they perceive to be playing from the viewpoint of the unrestricted game.

Definition 5 Fix a game Γ and an admissible subset of terminal histories $\tilde{Z} \subseteq Z$ with associated restricted game $\Gamma_{\tilde{Z}} \sqsubseteq \Gamma$. For any proposition $p \in \mathcal{L}_{\Gamma_{\tilde{Z}}}$, any history $h \in H$ and $i = P(h)$, $h \models_{\Gamma, \tilde{Z}} b_i p$ iff $h' \models_{\Gamma_{\tilde{Z}}} p$ for all $h' \in \mathcal{I}_{\tilde{Z}}(h)$.

Remark 2 Since the restricted game $\Gamma_{\tilde{Z}}$ is itself a game, it remains true that, for any $h' \in H_{\tilde{Z}}$, $j = P_{\tilde{Z}}(h')$, and any $p \in \mathcal{L}_{\Gamma_{\tilde{Z}}}$, $h' \models_{\Gamma_{\tilde{Z}}} k_j p$ if and only if $h'' \models_{\Gamma_{\tilde{Z}}} p$ for all $h'' \in \mathcal{I}_{\tilde{Z}}(h')$. That is, within any game, the knowledge operator is defined as usual. However, when considering a restriction of a game from the viewpoint of that game, we use the belief operator b_i . To sum this up

$$h \models_{\Gamma, \tilde{Z}} b_i p \Leftrightarrow h' \models_{\Gamma_{\tilde{Z}}} k_i p \text{ for any (and hence all) } h' \in \mathcal{I}_{\tilde{Z}}(h).$$

Note that $b_i p$ is a proposition in Γ referring to the restriction $\Gamma_{\tilde{Z}}$, while the proposition $k_i p$ is well-defined within each of Γ and $\Gamma_{\tilde{Z}}$, but has different meanings. In particular it is possible that $h \models_{\Gamma} \neg k_i p$ while $h' \models_{\Gamma_{\tilde{Z}}} k_i p$ for all $h' \in \mathcal{I}_{\tilde{Z}}(h)$ (and in the latter case if $h \notin \mathcal{I}_{\tilde{Z}}(h)$ it may be, but need not be true that $h \models_{\Gamma} \neg p$). That is, from the perspective of Γ 'knowledge' in $\Gamma_{\tilde{Z}}$ may be unreliable or false, which is why we call it belief. Only if $\mathcal{I}_{\tilde{Z}}(h) = \mathcal{I}(h)$ will knowledge and belief coincide.

It follows from the above remark that the belief operator does not satisfy the Truth Axiom T. Even more importantly, the belief operator does not satisfy negative introspection (Axiom 5 in S5). Consider any history $h' \in \mathcal{I}(h)$, such that $h' \notin H_{\tilde{Z}}$ and therefore $h' \notin \mathcal{I}_{\tilde{Z}}(h)$. Then, $h \models_{\Gamma, \tilde{Z}} \neg b_i p_h$ but also (since $\neg p_h$ is not a proposition in $\mathcal{L}_{\Gamma_{\tilde{Z}}}$) $h \not\models_{\Gamma, \tilde{Z}} b_i \neg p_h$ and $h \not\models_{\Gamma, \tilde{Z}} b_i b_i \neg p_h$.

Lemma 1 The belief operator satisfies the properties of the system $KD4$, namely

Necessitation (K) $p \Rightarrow b_i p$,

Distribution (D) $b_i(p \Rightarrow q) \Rightarrow (b_i p \Rightarrow b_i q)$,

Positive Introspection (4) $b_i p \Rightarrow b_i b_i p$

Since the belief operator displays non-trivial unawareness, we treat awareness and unawareness as derived operators. We write $a_i p$, read as ' i is aware of p ', and write $u_i p$, read as ' i is unaware

of p , and define what it means for a player to be aware (respectively, unaware) of a proposition in the game they perceive to be playing from the viewpoint of the unrestricted game.

Definition 6 Fix a game Γ , an admissible subset of terminal histories $\tilde{Z} \subseteq Z$, with associated restricted game $\Gamma_{\tilde{Z}} \sqsubseteq \Gamma$. For any proposition $p \in \mathcal{L}_{\Gamma_{\tilde{Z}}}$, any history $h \in H$ and $i = P(h)$,

- (i) $h \models_{\Gamma, \tilde{Z}} a_i p$ iff $h \models_{\Gamma, \tilde{Z}} b_i p \vee b_i (\neg k_i p)$
(ii) $h \models_{\Gamma, \tilde{Z}} u_i p$ iff $h \models_{\Gamma, \tilde{Z}} \neg a_i p$

In words, we say at a history h in the game Γ , player $i = P(h)$ who perceives to be playing the game $\Gamma_{\tilde{Z}}$, is aware of a proposition if from the perspective of the unrestricted game Γ , he either believes the proposition is true or believes he does not know the proposition is true. Notice in this formulation, the awareness operator preserves semantic equivalence. That is, if two propositions $p, q \in \mathcal{L}_{\Gamma_{\tilde{Z}}}$ are semantically equivalent then i is aware of p if and only if she is aware of q .²

2.3 Perception mapping

We are concerned with games with awareness, in which players, at any given history where they are called on to play, may be unaware of some possible histories. To represent this, we need to relate the game actually being played to the game as perceived by the players. We shall define the evolution of awareness of the players through the game by the perception mapping:

Definition 7 Fix a game Γ . A perception mapping is a function $\tilde{Z} : H \rightarrow 2^Z \setminus \{\emptyset\}$, where with each history $h \in H$, the set $\tilde{Z}(h) \subseteq Z$ is an admissible set of terminal histories considered by player $P(h)$ at history h . Denote by

$$\Gamma_{\tilde{Z}(h)} = (N_{\tilde{Z}(h)} \cup \{c\}, A_{\tilde{Z}(h)}, H_{\tilde{Z}(h)}, P_{\tilde{Z}(h)}, f_{\tilde{Z}(h)}^c, \mathcal{I}_{\tilde{Z}(h)}, \{v_{\tilde{Z}(h)}^i : i \in N_{\tilde{Z}(h)}\}),$$

the restricted game associated with the set of terminal histories $\tilde{Z}(h)$ that is imputed to player $P(h)$ at h .

We impose the following requirements on the perception mapping.

IN: (Information Neutrality): For all h, h' in H , $h' \in \mathcal{I}(h) \Rightarrow \tilde{Z}(h') = \tilde{Z}(h)$.

IA: (Increasing Awareness): If $h' \preceq h$ and $P(h) = P(h')$ then $\tilde{Z}(h') \subseteq \tilde{Z}(h)$.

NI: (No Impossibility): For all h , $\mathcal{I}_{\tilde{Z}(h)}(h) \neq \emptyset$.

The Information Neutrality property requires a player's level of awareness be congruent with the information structure of the full awareness game. More precisely, at any two histories in

² This is not true in the approach of Fagin and Halpern (1998) who take as their starting point the problem of logical omniscience which may be stated simply as follows. Supposing that an individual knows some proposition p to be true, and also knows $p \Rightarrow q$, does the individual also know q ? Fagin and Halpern address this question using the distinction between explicit and implicit knowledge. A proposition p is implicitly known if it is semantically equivalent to some q that is explicitly known. In the case when p is implicitly but not explicitly known by i , Fagin and Halpern say that i is unaware of p .

the same information set, the player's knowledge, beliefs and awareness are the same. This is consistent with the standard treatment of information sets in decision theory and game theory.

The Increasing Awareness property ensures that, once a player considers a (terminal) history, she does not forget it. We will say the game displays *non-trivial increasing awareness* for i at h if there exists some $h' \preceq h$ such that the inclusion is strict, that is, $\tilde{Z}(h') \subset \tilde{Z}(h)$.

The No Impossibility property ensures that the player who is to move always considers some history possible.³

For interactive awareness, we suppose for any pair of histories h and h' in H , such that there exists $h'' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ for which $h'' \preceq h'$, and with $i = P(h)$ and $j = P(h')$, we can define the *second-order imputation* in which i at h imputes to j at h' consideration of the set of terminal histories $\tilde{Z}(h) \cap \tilde{Z}(h')$. That is, player i at h cannot impute to j at h' consideration of histories which i at h herself has not also considered. On the other hand, there is no reason to suppose that player i at h should incorrectly impute to j at h' failure to consider histories that are in fact considered by both i at h and j at h' . In particular if $P(h) = P(h')$ then by properties **IN** (Information Neutrality) and **IA** (increasing awareness) of the perception mapping, it follows that $\tilde{Z}(h) \cap \tilde{Z}(h') = \tilde{Z}(h)$. That is, a player cannot anticipate her future increasing awareness.

Higher order imputations may be similarly constructed. However, given the properties of the perception mapping defined above, it turns out we need only consider second-order imputations.

To avoid the possibility that a player at the end of a second-order imputation is perceived to have reached an empty information set, we extend the No Impossibility condition as follows.

NI*: (No impossibility): For any for any pair of histories h and h' in H , such that there exists $h'' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ for which $h'' \preceq h'$, and with $i = P(h)$ and $j = P(h')$, the set of histories $\mathcal{I}_{\tilde{Z}(h)}(h') \cap \mathcal{I}_{\tilde{Z}(h')}(h')$ is not empty.

We now have all the elements needed to define a game with awareness.

Definition 8 A game with awareness \mathcal{G} is characterized by the tuple $(\Gamma, \tilde{Z}(\cdot))$ where $\Gamma = (N \cup \{c\}, A, H, P, f^c, \mathcal{I}, \{v^i : i \in N\})$ is the full (maximally aware) extensive form game the players are actually playing and $\tilde{Z}(\cdot)$ is the associated perception mapping to admissible sets of terminal histories, that satisfies **IN**, **IA** and **NI*** and encodes the evolution of interactive awareness of the players i in N . That is, for each history h in H , player $P(h)$ is aware of the game $\Gamma_{\tilde{Z}(h)}$. Furthermore, for any pair of histories h and h' in H , such that there exists $h'' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ for which $h'' \preceq h'$, and with $i = P(h)$ and $j = P(h')$, i at h imputes to j at h' consideration of the set of terminal histories $\tilde{Z}(h) \cap \tilde{Z}(h')$.

³ A special case of this property applies if awareness changes only when the information set would otherwise be empty. In this case, individuals apply Bayesian updating until new information rules out all the histories they had previously considered possible, at which point they must consider new conjectures about the world, associated with new prior probability distributions. This special case will not be considered in detail here.

In a game with awareness, the information set $\mathcal{I}_{\tilde{Z}(h)}(h)$ is determined by the set of terminal histories of which player $P(h)$ is aware, and which have not been ruled out by the information available to $P(h)$ at h . Whereas in a standard extensive form game, the set of terminal histories consistent with available information grows monotonically smaller until a unique terminal history is reached. By contrast, in a game with awareness, players may become aware of new possibilities. Nevertheless, the information set must ultimately contract as the game approaches the terminal history.

This treatment of differential awareness may be compared to the standard common knowledge assumptions in a game of full awareness. In this case, $\tilde{Z}(h) = Z$ for all h , and the associated game $\Gamma = \Gamma_{\tilde{Z}(h)}$, for all h in H , is a standard extensive form game. Obviously, this assumption greatly simplifies the analysis for the external modeller, and the computational problem facing the players. However, the assumption of common knowledge of the game is exceptionally strong. The approach adopted here represents the most limited possible modification of the full awareness case.

An alternative approach, adopted in the context of static decisions by Heifetz, Meier and Schipper (2006), is to consider bounded rationality as generating a ‘coarsening’ of the game structure. That is, one or more histories (or, in the static context, states) in the complete game are mapped to a single history in the coarsened game. The coarsening approach avoids the need to impute beliefs regarding propositions of which a player is unaware. As a result, the only false beliefs players can hold are beliefs about the knowledge and awareness available to other players (Galanis, 2011). By contrast, as can be seen from the discussion above, the restriction approach adopted here allows false beliefs even in the case of a single-player game with Nature. However, the coarsening approach has so far proved intractable in the context of extensive form games. Among other problems, we note the difficulty of specifying consequences when two terminal histories with different payoff vectors are merged. Partly as a result of these difficulties, most treatments of extensive-form games with unawareness, including those of Halpern and Rego (2006b) and Heifetz, Meier and Schipper (2009), have adopted an approach similar to that presented above.

Remark 3 *In a slight abuse of notation, for any proposition p , and history h , and where $i = P(h)$, we will write $h \models_{\Gamma} b_i p$ instead of $h \models_{\Gamma, \tilde{Z}(h)} b_i p$. This simplification reflects the fact that, since the restriction $\tilde{Z}(h)$ applicable at h in Γ is entirely determined by h and \mathcal{G} , it is redundant to spell it out.*

2.4 Behavior rules, strategies, subjective probabilities and equilibrium in games with awareness

In a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$, at each history h in the game, the player $P(h)$ who is called upon to play at that history selects a randomization defined over $A_{\tilde{Z}(h)}(h)$, the actions she

thinks are available. A collection of these randomizations constitutes a ‘rule’ for determining the ‘play’ of the game Γ .⁴

Definition 9 Let $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ be a game with awareness. A **behavioral rule** r for the players in the game Γ is a collection of randomizations over actions $(r(h) : h \in H, \text{ s.t. } P(h) \neq c)$, where each $r(h)$ is a probability distribution whose support is a subset of $A_{\tilde{Z}(h)}(h)$, the set of actions player $P(h)$ perceives to be available at her information set $\mathcal{I}_{\tilde{Z}(h)}(h)$. Denote by $\tau(r)$ the probability distribution over terminal histories induced by r and the chance assignment f^c of nature. The (ex ante) expected value for player i of the play of the game Γ associated with the rule r is given by

$$V_{\Gamma}^i(r) = \sum_{h \in \mathcal{Z}} \tau(r)[h] v^i(h). \quad (1)$$

To determine at each history h in the game Γ what the player $P(h)$ will select we shall impute to that player a ‘theory’ of how she thinks the game will be played, in the form of a strategy profile for the continuation of the game she perceives she is playing at that history. We assume her theory of how the game is being played at that history is the same for any player at any history that goes through that information set who perceives himself to be playing the same game. In particular, this implies that the strategy profile ascribed to a player remains unchanged if her level of awareness at a subsequent information set is unchanged. In addition, we require the strategy a player’s theory ascribes to another player must be consistent with the awareness imputed to that player. Formally we define a strategy profile for a game with awareness as follows.

Definition 10 Let $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ be a game with awareness. A **strategy profile** $\beta = (\beta_h : h \in H)$ for \mathcal{G} assigns to each history h , a behavioral strategy profile β_h for the continuation of the game $\Gamma_{\tilde{Z}(h)}$ from the information set $\mathcal{I}_{\tilde{Z}(h)}$ with the consistency properties: for any history $h' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ and any history h'' for which $h' \preceq h''$:

1. if $\tilde{Z}(h') = \tilde{Z}(h'')$ then the continuation of $\beta_{h'}$ from $\mathcal{I}_{\tilde{Z}(h')}(h'')$ coincides with $\beta_{h''}$;
2. for player $P(h'') = j$, the support of $\beta_h(\mathcal{I}_{\tilde{Z}(h)}(h''))$ is a subset of $A_{\tilde{Z}(h)}(h'') \cap A_{\tilde{Z}(h'')}(h'')$.

Remark 4 Recall the property **NI*** (No impossibility) of the perception mapping ensures that a player at the end of any second-order imputation is not imputed to have reached an empty information set. Hence it is always possible for a player to ascribe how another player is playing in

⁴ If the players are all aware of the game Γ , then the behavioral rule corresponds to what we referred to in subsection 2.1 as a behavioral strategy profile. In the setting we consider below, however, one or more of the players may not be fully aware of the game Γ that they are actually playing. Although a less than fully aware player might adopt a ‘rule of play’ that will determine her choice at ever information set she may encounter in the game Γ , in general she will not have access to the behavioral rule generated by the way she and her opponents decide on their choice of actions during the course of play. Hence we feel it inappropriate to refer to such a behavioral rule as a strategy profile.

a way that is consistent with the game she imputes to her opponent. The property **IN** (Information Neutrality) in conjunction with the first part of the consistency property for a strategy profile ensures that the ‘theory’ a player holds about how the game is being played at a particular history is the same for all other histories in the same information set.

For a game of common awareness, the strategy profile conforming to the consistency property defines a strategy profile for a standard game. Conversely any standard behavioral strategy profile for a standard game defines a behavioral strategy profile for the associated game of common awareness.

Next consider a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$, in which, for each h in H , either $\tilde{Z}(h) = Z$ or $\tilde{Z}(h) = \tilde{Z} \subset Z$, and with $\Gamma_{\tilde{Z}}$ a game of common awareness. Then, whenever $\tilde{Z}(h) = Z$, the strategy profile defines at information set $\mathcal{I}(h)$ a probability measure on $A(h)$. On the other hand, if $\tilde{Z}(h) = \tilde{Z}$ with $\Gamma_{\tilde{Z}}$ a game of common awareness, then β_h is a strategy profile for the continuation of the standard game $\Gamma_{\tilde{Z}}$ and therefore for each h' in $H_{\tilde{Z}}$ that follows from $\mathcal{I}_{\tilde{Z}}(h)$, $\beta_h(\mathcal{I}(h'))$ defines a probability measure on $A_{\tilde{Z}}(h) \subseteq A(h)$. That is, the randomization over actions imputed by player $P(h)$ at history h to another player $P(h')$ at his information set $\mathcal{I}(h')$ must be consistent with the game $\Gamma_{\tilde{Z}}$ that the game with awareness at h imputes that player $P(h)$ imputes to player $P(h')$ at history h' .

More generally, in any game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$, any strategy profile $\beta = (\beta_h : h \in H)$ generates a behavioral rule r for the game Γ given by $r(h) = \beta_h(\mathcal{I}_{\tilde{Z}(h)}(h))$.

Analogous to standard games (that is, games of common awareness), given a strategy profile, in order for a player to be able to evaluate her expected continuation payoff from an information set that she would be called upon to play, we also need to specify her subjective beliefs about where in her information set she thinks she is. Thus we extend the definition of a subjective probability system for a game with awareness as follows.

Definition 11 Let $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ be a game with awareness. A **subjective probability system** $\mu = (\mu_h : h \in H)$ for \mathcal{G} assigns to each history h in H , a probability measure on the set of histories in $\mathcal{I}_{\tilde{Z}(h)}(h)$ with the consistency property: for any h', h'' in H , if $h'' \in \mathcal{I}(h')$ then $\mu_{h'} = \mu_{h''}$.

The interpretation of the subjective probability measure μ_h is that at history h , the probability that player $i = P(h)$ assigns to being at the history $h' \in \mathcal{I}_{\tilde{Z}(h)}(h)$, is $\mu_h(h')$. Thus the consistency condition ensures that subjective probabilities are common across all elements of an information set.

Definition 12 Let $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ be a game with 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} .

Given a strategy profile β for \mathcal{G} , for each h in H denote by $\tau_{\Gamma_{\tilde{Z}(h)}}(\beta_h)$ the probability distribution over terminal histories induced by the strategy profile β_h . Given the assessment (β, μ) , we thus have at each h in H , player $i = P(h)$ to whom the perception mapping imputes the game $\Gamma_{\tilde{Z}(h)}$ perceives the continuation value from her information set $\mathcal{I}_{\tilde{Z}(h)}(h)$ from the play of the game according to the continuation strategy profile β_h to be:

$$V_{\Gamma_{\tilde{Z}(h)}}^i(\beta_h, \mu_h) = \sum_{h' \in \mathcal{I}_{\tilde{Z}(h)}(h)} \mu_h[h'] \left(\sum_{h'' \in \tilde{Z}(h)} \tau_{\Gamma_{\tilde{Z}(h)}}(\beta_h)[h''] v^i(h'') \right) \quad (2)$$

Definition 13 Let \mathcal{G} be a game with awareness. An **assessment** (β, μ) is **sequentially rational** for player i , if i is playing a ‘best response’ at each of her information’s sets. That is, for every h in H , such that $i = P(h)$:

$$V_{\Gamma_{\tilde{Z}(h)}}^i(\beta_h, \mu_h) \geq V_{\Gamma_{\tilde{Z}(h)}}^i\left(\left(\hat{\beta}_h^i, \beta_h^{-i}\right), \mu_h\right),$$

for every continuation strategy $\hat{\beta}_h^i$ of i in the continuation of the game $\Gamma_{\tilde{Z}(h)}$ from the information set $\mathcal{I}_{\tilde{Z}(h)}(h)$.

The **assessment** (β, μ) is **sequentially rational** if it is sequentially rational for all i .

We shall refer to a behavioral strategy profile β for the game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ as being **completely mixed as can be done consistently** if for any three histories h, h' and h'' , such that $h' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ and $h' \preceq h''$, $\beta_h(\mathcal{I}_{\tilde{Z}(h)}(h'))$ assigns positive probability to every action in $A_{\tilde{Z}(h)}(h'') \cap A_{\tilde{Z}(h')}(h'')$. That is, if player $P(h)$ at h , perceives herself to be at the information set $\mathcal{I}_{\tilde{Z}(h)}(h)$ in the game $\Gamma_{\tilde{Z}(h)}$ and imputes to player $P(h'')$ at his information set $\mathcal{I}_{\tilde{Z}(h)}(h'')$ in that game that he will perceive himself to be at the information set $\mathcal{I}_{\tilde{Z}(h) \cap \tilde{Z}(h'')}(h'')$ in the game $\Gamma_{\tilde{Z}(h) \cap \tilde{Z}(h'')}$, then each action available in the game $\Gamma_{\tilde{Z}(h) \cap \tilde{Z}(h'')}$ at the information set $\mathcal{I}_{\tilde{Z}(h) \cap \tilde{Z}(h'')}(h'')$ (that is, each action in $A_{\tilde{Z}(h)}(h'') \cap A_{\tilde{Z}(h')}(h'')$) is assigned strictly positive weight by the strategy profile.

Definition 14 Let \mathcal{G} be a game with 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.

Thus we have all the elements to extend the concept of a sequential equilibrium to a game with awareness.

Definition 15 Let \mathcal{G} be a game with awareness. An **assessment** (β, μ) is a **sequential equilibrium** if it is sequentially rational and consistent.

We prove in the appendix the following.

Proposition 2 *A sequential equilibrium exists for any game with awareness.*

Sequential equilibrium is not the only equilibrium concept that might be considered for games with awareness. For example, Halpern and Rego consider subgame perfect equilibrium. We consider that sequential equilibrium is a sensible choice because, in a game with awareness, players will, in general, have the experience of reaching information sets that they previously considered as off-equilibrium. This experience can arise because other players were acting on the basis of a different perception of the game from the one imputed to them. It therefore makes sense for players to confine attention to strategies that prescribe reasonable behavior at every information set at which they may move, and not merely at those that occur with positive probability in the equilibrium of the game they perceive themselves to be playing.

Proposition 2 demonstrates existence, but not uniqueness of sequential equilibrium. This limitation is not specific to games with awareness. For example Rubinstein (1985) demonstrated the existence of large numbers of sequential equilibria in an asymmetric information version of a two-player bargaining game of alternating offers. However, he is able to establish a unique equilibrium for this game, under additional assumptions on what inferences the uninformed player may draw from an off-equilibrium move by the informed player.

This approach to refinement may be more problematic in games with awareness, where an unexpected off-equilibrium move may lead to a change in awareness. Moreover, if a player is unaware of the possibility of an action by another player, then the other player's decision not to take that action cannot inform her inference about equilibrium selection. Heifetz, Meier and Schipper (2011) provides an interesting example of this kind.

To the best of our knowledge, there is no general characterization of uniqueness conditions for sequential equilibrium (except under perfect information). We conjecture that any sufficient condition for uniqueness of sequential equilibrium in standard extensive-form games would also apply to games with awareness if the condition is preserved under the restriction operation.

However, a comprehensive treatment of the problem of equilibrium selection is beyond the scope of the present paper. We will confine our attention to the case of games in which there exists a unique equilibrium for the maximal game Γ and for all restrictions of Γ . While this is a restrictive condition, it admits a significant class of games, notably including (generically) all one-person games with Nature.

2.5 Example

To illustrate the structure of a game with awareness, 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 full or maximal awareness game. The initial history is $\langle \rangle$. Nature's first move is a choice from the set $\{\alpha_n, \alpha_0\}$, (innovation or null action). Let $p \in (0, 1)$ denote the probability Nature chooses α_n . Nature's second move is a choice from the set $\{\alpha_\ell, \alpha_0\}$ (lawsuit or null action). Let $q \in (0, 1)$ denote the probability Nature chooses α_ℓ . There are now four histories $\langle \alpha_n, \alpha_\ell \rangle, \langle \alpha_n, \alpha_0 \rangle, \langle \alpha_0, \alpha_\ell \rangle, \langle \alpha_0, \alpha_0 \rangle$, 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 $\mathbf{I}_2 = \{\langle \alpha_n, \alpha_\ell, \alpha_1 \rangle, \langle \alpha_n, \alpha_0, \alpha_1 \rangle, \langle \alpha_0, \alpha_\ell, \alpha_1 \rangle, \langle \alpha_0, \alpha_0, \alpha_1 \rangle\}$ 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 history h in \mathbf{I}_1 , the buyer only considers the terminal histories in which the lawsuit does not arise. That is,

$$\tilde{Z}(h) = \tilde{Z}_1 = \{\langle \alpha_n, \alpha_0, \alpha_0 \rangle, \langle \alpha_n, \alpha_0, \alpha_1, \alpha_A \rangle, \langle \alpha_n, \alpha_0, \alpha_1, \alpha_R \rangle, \langle \alpha_0, \alpha_0, \alpha_0 \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_A \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_R \rangle\},$$

for all h in \mathbf{I}_1 .

Denoting the game of which the buyer is aware at any history h in \mathbf{I}_1 by Γ^1 , we see that it is obtained from the fully aware game Γ by deleting all histories containing nature's second move α_ℓ .

Similarly, at each history h in \mathbf{I}_2 , the owner only considers the terminal histories in which the business opportunity does not arise. That is,

$$\tilde{Z}(h) = \tilde{Z}_2 = \{\langle \alpha_0, \alpha_\ell, \alpha_0 \rangle, \langle \alpha_0, \alpha_\ell, \alpha_1, \alpha_A \rangle, \langle \alpha_0, \alpha_\ell, \alpha_1, \alpha_R \rangle, \langle \alpha_0, \alpha_0, \alpha_0 \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_A \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_R \rangle\},$$

for all h in \mathbf{I}_2 .

Denoting the game of which the owner is aware at any history h in \mathbf{I}_2 by Γ^2 , we see that it is obtained from the fully aware game Γ by deleting all histories containing nature's first move α_n . The games are illustrated in figures 2 and 3, respectively.

<INSERT FIGURES 2 and 3 here>

Both parties impute to the other a restriction of their own game. The buyer is unaware of the possible lawsuit, and assumes the owner to be unaware of the possible innovation (at all non-terminal histories), while the converse is true for the owner. Hence each imputes the other only considers the terminal histories

$$\tilde{Z}_1 \cap \tilde{Z}_2 = \{\langle \alpha_0, \alpha_0, \alpha_0 \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_A \rangle, \langle \alpha_0, \alpha_0, \alpha_1, \alpha_R \rangle\}.$$

These lead to the same game Γ^3 , as illustrated in figure 4 and which is a game of common awareness.

<INSERT FIGURE 4 around here>

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 (β^*, μ^*) , in which

$$\begin{aligned} \beta_{<>}^* (\mathbf{I}_1) [\alpha_1] &= 1, \beta_{<>}^* (\mathbf{I}_2) [\alpha_A] = 1 \\ \beta_h^* (\mathbf{I}_1^1) [\alpha_1] &= 1, \beta_h^* (\mathbf{I}_2^1) [\alpha_A] = 1, \text{ for all } h \text{ in } \mathbf{I}_1 \\ \beta_h^* (\mathbf{I}_2^2) [\alpha_A] &= 1, \text{ for all } h \text{ in } \mathbf{I}_2. \end{aligned}$$

$$\mu_{<>}^* [<>] = 1$$

$$\mu_h^* [\langle \alpha_n, \alpha_0 \rangle] = p, \mu_h^* [\langle \alpha_0, \alpha_0 \rangle] = 1 - p, \text{ for all } h \text{ in } \mathbf{I}_1$$

$$\mu_h^* [\langle \alpha_0, \alpha_\ell \rangle] = q, \mu_h^* [\langle \alpha_0, \alpha_0 \rangle] = 1 - q, \text{ for all } h \text{ in } \mathbf{I}_2$$

To check sequential rationality, notice that at $\mathbf{I}_1^1 = \mathbf{I}_1 \cap H_{\Gamma^1}$ (in Γ^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 $\mathbf{I}_2^2 = \mathbf{I}_2 \cap H_{\Gamma^2}$ (in Γ^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 $\langle \alpha_0, \alpha_0, \alpha_1, \alpha_R \rangle$.⁵

3 Inductive reasoning about unawareness

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

⁵ Note that, at a price of 1, both parties strictly prefer to trade, and each imputes to the other a game in which they are indifferent between trading and no-trading. Further, all of this is common knowledge. This example does not, however, allow for common knowledge of a strict preference for trade. Heifetz, Meier and Schipper (2012) show that, in general, unawareness cannot produce common knowledge of mutual strict preference for speculative trade.

We begin by extending the language to include existential propositions of the general form ‘there exists a proposition p with property θ ’. Our primary interest is on existential propositions related to awareness, most notably, ‘there exists a proposition p of which I am unaware’. Our first result, Proposition 3, is negative. We show that, within the modal-logical representation of knowledge developed above, an individual can never believe that there exist propositions of which she is unaware. In view of Proposition 3, we must 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.

As we have argued above, an individual’s understanding of their own unawareness cannot be represented within the context of a semantic-syntactic game representation, even when the game itself is extended to allow for differential awareness. The kind of reasoning that can be represented in such a context may broadly be described as ‘deductive’. That is, an initial set of premises (the game tree and prior probabilities in the semantic rendition, the set of known propositions, tautologies and implications in the syntactic rendition) is combined with new information (signals in the semantic rendition, learning about the truth values of propositions in the syntactic rendition) to yield a new and improved model. Given sufficient information, the realised history of the game and the truth value of all propositions in the associated language may be determined.

The deductive mode of reasoning associated with games of common awareness and the associated modal logic of knowledge appear to offer the logical certainty of conclusions derived, in accordance with stated axiomatic properties, from known premises. In a game with awareness, however, this certainty is spurious. As we have seen, a proposition may be true in all states an individual considers possible, but nevertheless false in reality. Decision procedures that presume that logical certainty can be attained are likely to yield poor outcomes.

To address this problem, we need to answer two questions.

* First, how can, or should, individuals reason about their own unawareness?

* Second, if this reasoning supports the conclusion that the individual is unaware of some relevant contingencies, how should she act?

To address the first question, we need to consider how an individual might reach either a positive or a negative answer to the question ‘Do there exist relevant propositions of which I am unaware?’. One answer (though not the only one) is to consider inductive reasoning, based on generalization of past experience.

Consider first the situation of an individual who is to move in an extensive form game with awareness. Under the assumption of increasing awareness (and assuming it holds non-trivially), the individual’s previous experience includes a number of ‘surprises’, that is, discoveries of possible terminal histories of the game she had previously not considered. Inductive reasoning supports the judgement ‘if I have been surprised in the past, I may be surprised in the future’.

Taken to an extreme, such reasoning could be paralyzing (a point made by critics of the precautionary principle, discussed below). Since decisions cannot be avoided, decisionmakers must adopt some combination of heuristics and formal rules to guide their choices.

In particular, individuals may use inductive criteria for identifying a ‘small world’, in which it is reasonable for players to disregard the possible existence of relevant unconsidered contingencies, and act on the basis of Bayesian decision theory applied to the game they perceive.

Inductive reasoning about games with awareness may therefore lead individuals to conjecture either that they are or they are not, aware of all relevant contingencies. As will be shown by Proposition 3 below, such reasoning cannot be encompassed by the usual modal logic of awareness, if unawareness is represented, as it is here, by a failure to consider some possible histories of the game.

However, it can be represented using a more general syntactic-semantic framework, such as that put forward by Walker (2011), with precisely this kind of reasoning in mind. Walker suggests a two-stage evaluation framework for propositions, with the first ‘subjective’ stage incorporating the decisionmaker’s inductively derived conjectures about their awareness, and the second ‘objective’ stage, incorporating the objective evaluation of an external observer.

Walker develops this approach with reference to the Fagin-Halpern (1988) approach in which unawareness is modelled in terms of a distinction between implicit and explicit awareness, so that a decisionmaker may not be (explicitly) aware of a proposition even though they are aware of semantically equivalent proposition. However, no changes to the central idea are required to apply the idea in the present context. The main difference is that whereas Walker’s syntax, developed in a static setting, includes only the single string corresponding to the sentence ‘I am aware of all propositions’, our dynamic, extensive-form, approach would require subjective evaluation of a larger class of sentences of the general form ‘at my current position (in the game) I am aware of all relevant propositions.’

4 The existential quantifier, existential propositions and unawareness

The representation of differential awareness developed thus far is fairly standard, and is consistent with Halpern and Rego (2007) and endorsed by Heifetz, Meier and Schipper (2008). The central unresolved issue in the literature is how to deal with the fact that individuals may be conscious of their own bounded awareness, and of the possibility that others may be aware of histories (or propositions) of which they themselves are unaware. We note that from the definition of the belief operator for the restricted game an individual has access to, that individual cannot believe that there exist propositions of which they are unaware. *On the other hand, individuals may believe*

(in fact, know) that other players are unaware of some propositions. Thus, we must consider interpretations of the statement ‘there exist propositions/histories of which I am unaware’ that lie outside the syntactic-semantic framework developed so far.

In a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ although the language $\mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$ available to individual $P(h)$ at the history h is sufficiently expressive to describe the restricted game $\Gamma_{\tilde{Z}(h)}$ she perceives to be playing, it is inadequate to describe propositions she might reasonably entertain about the full awareness game Γ and its associated language \mathcal{L}_{Γ} . The approach adopted here begins by extending the languages $\{\mathcal{L}_{\Gamma_{\tilde{Z}(h)}} : h \in H\}$ to allow for reasoning about the existence of propositions as follows. We include an existential quantifier \exists , used in conjunction with a formula for substitution to produce propositions of the form

$$\exists q \in \mathcal{L}_{\Gamma} \text{ such that } (\theta(q)), \quad (3)$$

where $\theta(q)$ is a Boolean combination of the free proposition q and propositions in \mathcal{L}_{Γ} . We denote by $\theta(q|p)$ the proposition obtained by replacing all instances of q with p . For compactness, in the formal analysis to follow, the existential proposition in (3) will be denoted $q_{\exists}\theta$. Then $h \models_{\Gamma} q_{\exists}\theta$ if and only if there is some $p \in \mathcal{L}_{\Gamma}$ such that $h \models_{\Gamma} \theta(q|p)$.

As is standard, we will define the derived universal operator \forall by

$$p_{\forall}\theta \Leftrightarrow \neg q_{\exists}\neg\theta$$

That is, property θ holds for all p if there does not exist q such that $\neg\theta q$ holds.

Example 1 *As an illustration, fix a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ and consider a history h , for which $\tilde{Z}(h) \subset \tilde{Z}$. For a given $p \in \mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$, the extended language contains such propositions as ,*

$$\exists q \in \mathcal{L}_{\Gamma} \text{ such that } ((q \Rightarrow p) \wedge \neg(p \Rightarrow q))$$

which we may interpret as saying that there is some (non-equivalent) proposition q in the richer language \mathcal{L}_{Γ} 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 for. The evidence could be either propositional (X was at the scene of the crime) or epistemological (X knew that the gun was loaded).

Given that the language \mathcal{L}_{Γ} , as defined, contains an infinite number of propositions, evaluation of propositions involving existential quantifiers is, in general, problematic. But in the framework developed here, all propositions refer to the extensive form game Γ , which contains only finitely many histories and awareness preserves semantic equivalence. So, we may consider the finite set

of propositions derived from the power set 2^H of H which consists of all finite disjunctions of the form $p_h \vee p_{h'} \vee \dots$. We will abuse notation by referring to this set of propositions simply as 2^H .

Now for any proposition $p \in \mathcal{L}_\Gamma$ let $\mathfrak{S}(p)$ denote the set of histories at which p is true. Then there exists a unique element $\hat{p} \in 2^H$ such that \hat{p} is semantically equivalent to p , defined as

$$\hat{p} = \bigvee_{h \in \mathfrak{S}(p)} p_h$$

that is, the disjunction of the elementary propositions p_h taken over $\mathfrak{S}(p)$. It is immediate that, for any $\theta(q)$, $\theta(q|p) \Leftrightarrow \theta(q|\hat{p})$. But now it follows that

$$q_{\exists}\theta \Leftrightarrow \bigvee_{p \in 2^H} \theta(q|p)$$

That is, the truth value of $q_{\exists}\theta$ may be evaluated by testing the finite disjunction of all propositions of the form $\theta(q|p)$ taken over $p \in 2^H$.⁶

The implication of this in a game with awareness is stated in the next proposition.

Proposition 3 Fix a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$. For all $h \in H$, and $i = P(h)$,

$$h \models_{\Gamma} b_i(\neg(q_{\exists}u_iq)) \text{ or, equivalently, } h \models_{\Gamma} b_i(q_{\forall}a_iq).$$

That is, individuals can never believe they are unaware of anything.

Proof. For any $h' \notin H_{\tilde{Z}(h)}$, the corresponding proposition $p_{h'}$ is not an element of the set $2^{H_{\tilde{Z}(h)}}$ over which the unawareness operator in $\mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$ is defined. On the other hand, for any $h'' \in H_{\tilde{Z}(h)}$, individual i is aware of $p_{h''}$. Taking the logical closure, for any $\tilde{h} \in \mathcal{I}_{\tilde{Z}(h)}(h)$, and any $p \in \mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$, $\tilde{h} \models_{\Gamma_{\tilde{Z}(h)}} a_i p$. Since this result holds when h is replaced by any $\tilde{h} \in \mathcal{I}_{\tilde{Z}(h)}(h)$, we have $b_i(q_{\forall}a_iq)$ which is logically equivalent to $b_i(\neg(q_{\exists}u_iq))$, as required. ■

Proposition 3 shows that, given the specification of $\mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$, including the existential quantifier \exists which generates the set of existential propositions, the richer language \mathcal{L}_Γ is not expressive enough to allow valid statements of the form ‘individual i believes that there exists some proposition $q \in \mathcal{L}_\Gamma$ of which he is currently unaware’. This is not surprising. To say that a player believes that there exist events of which he is unaware suggests, in some sense, that he is aware of those events, which might be seen as violating the spirit of what it means to be unaware of something.⁷ On the other hand, to the extent that individuals understand the structure of a game with awareness, that understanding must encompass the possibility that their own awareness is incomplete. This apparent contradiction suggests the need to consider modes of reasoning going beyond the

⁶ Similarly, the truth value of $q_{\forall}\theta$ may be evaluated by testing the finite conjunction of all proposition of the form $\theta(q|p)$ taken over $p \in 2^H$.

⁷ We owe this characterization of the result to an anonymous referee.

semantic-syntactic model considered thus far. We argue below that the deductive reasoning characteristic of the semantic-syntactic model must be combined with inductive reasoning about the structure of the model itself.

Proposition 3 does not hold if unawareness is modelled in terms of a distinction between implicit and explicit knowledge, as in Fagin and Halpern (1988). With this representation, as noted by Halpern and Rego (2011), unawareness may arise even in a static model with a single state of the world. In such a model, individuals know any proposition that is in fact true and of which they are explicitly aware, including the proposition $q \exists u_i q$.

4.1 Historical induction and induction over individuals

The most commonly used alternative to deductive reasoning is reasoning based on induction from experience or observation. 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 ϕ , provide inductive support for the proposition ‘All members of set S satisfy property ϕ ’. 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’.

As the famous example of black swans shows, inductive reasoning is never conclusive. It is easy to define propositions that have always been true, but will cease to be true at some point in the future, either because they inherently involve time dating (person i , now aged 20, has always been younger than 21) or because the properties to which they refer change over time (US population has always been less than 320 million). Moreover, it is possible to derive inductive support for two or more propositions that may be logically inconsistent. For example, it is common to use inductive arguments to predict the outcomes of Presidential elections (for example, that no incumbent president has been re-elected if his approval rating is below x , or that incumbents are always re-elected if the economy has improved during their term of office). It will often be the case that two such inductive arguments will point in opposite directions.

Under conditions of bounded rationality, however, no system of reasoning is absolutely reliable. Moreover, a judgement that it is appropriate to use deductive reasoning in some particular context must be based on some prior process of reasoning that is not itself deductive. So, boundedly rational individuals may find it appropriate to employ a mixture of inductive and deductive reasoning.

We begin by considering reasoning based on induction from experience (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.

Formally, we state it as follows.

Definition 16 (support by historical induction) Fix a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ and a history h , and let $i = P(h)$. Suppose that for some proposition $p \in \mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$,

$$h \models_{\Gamma} b_i w p \wedge \neg b_i w \neg p$$

then $h \models_{\Gamma} t_i p$ [read as ‘at h individual i regards p as supported by historical induction’].

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

In interactive games, the principle of induction may also be applied to reasoning about other individuals. In the application here, the sets to which induction is to be applied will consist of the set of individuals N . If individual i believes that some proposition is true for all other players, then induction suggests that the same proposition is true for i .

To formalize this idea in our notation we define the following.

Definition 17 (support by induction over others) Fix a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ and a history h , and let $i = P(h)$. Let $\theta(\tilde{j})$ be a Boolean combination of propositions in $\mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$ and a free variable \tilde{j} , defined as an element of the set of players N , such that $\theta(\tilde{j})$ is the proposition in which each instance of \tilde{j} in $\theta(\cdot)$ is replaced by j . Suppose that, for all $j \neq i$

$$h \models_{\Gamma} b_i q(j)$$

Then, $h \models_{\Gamma} n_i \theta(i)$ [read as ‘at h player i regards $q(i)$ as supported by induction over the set of players $N - \{i\}$ ’]

Thus players can use their observations of others to draw conclusions about themselves that cannot necessarily be reached by deductive reasoning.

4.2 Inductive reasoning, awareness and unawareness

Our account of historical inductive reasoning about unawareness has a structure similar to that of more familiar examples of historical induction. Over time, players become aware of propositions that, previously, they have not considered. Thus, at any history h , players know that there exist propositions of which they were unaware at some previous information set. That is, they know that the existential proposition ‘there exists a proposition of which I am unaware’ was true at a previous information set. 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.⁸

On the other hand, players also learn positive lessons about their ability to understand particular subproblems in larger decision problems. So, they may become confident, when they perceive a particular problem to be of limited complexity, that they will not be surprised by becoming aware of unconsidered propositions relevant to their decision. The most important case, for our purposes, will be the choice of a strategy in the continuation of a game from some information set.

Suppose that for $h' \prec h$, player i becomes aware of previously unconsidered propositions in the course of the partial history from h' to h , so that $\mathcal{L}_{\Gamma_{\tilde{Z}(h')}} \subset \mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$ where the inclusion \subset is strict. Then for propositions $q \in \mathcal{L}_{\Gamma_{\tilde{Z}(h)}} - \mathcal{L}_{\Gamma_{\tilde{Z}(h')}}$ we have $h' \models_{\Gamma} u_i q$ and $h \models_{\Gamma} a_i q$, so that $h \models_{\Gamma} w u_i q$. Further, these evaluations hold for any $\tilde{h} \in \mathcal{I}_{\tilde{Z}(h)}(h)$, so that $h \models_{\Gamma} b_i a_i q$ and $h \models_{\Gamma} b_i w u_i q$. That is, at history h , the player (correctly) believes that there exist propositions of which she is now aware, but was unaware in the past.

We have established that player i at h cannot express and therefore cannot believe (in the modal logical sense of b_i) propositions of the form $\exists q \in \mathcal{L}_{\Gamma} : u_i q$. Nevertheless, given past experience of discovery, it seems reasonable to suppose that the player 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 4 *Fix a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ that displays non-trivial increasing awareness for i at h . Then $h \models_{\Gamma} t_i(q \exists u_i q)$.*

Proof. By non-trivial increasing awareness for i at h , there exists $h' \prec h$, in which $\mathcal{L}_{\Gamma_{\tilde{Z}(h')}} \subset \mathcal{L}_{\Gamma_{\tilde{Z}(h)}}$. Hence for a proposition $p \in \mathcal{L}_{\Gamma_{\tilde{Z}(h)}} - \mathcal{L}_{\Gamma_{\tilde{Z}(h')}}$ we have $h \models_{\Gamma} a_i p$ and $h' \models_{\Gamma} u_i p$, so that $h \models_{\Gamma} w u_i p$. By IN (information neutrality) these evaluations hold for all $\tilde{h} \in \mathcal{I}_{\tilde{Z}(h)}(h)$, so that $h \models_{\Gamma} b_i(q \exists u_i q)$ holds. From proposition 3 it follows that $q \exists u_i q$ can never known to be false at any $h' \prec h$, hence $h \models_{\Gamma} \neg b_i w(\neg(q \exists u_i q))$ holds as well, implying $h \models_{\Gamma} t_i(q \exists u_i q)$ as required. ■

Informally, given non-trivial increasing awareness, the player believes at h that, for at least some past history, the proposition $q \exists u_i q$ was true. On the other hand, this proposition can never be known false. Hence, the player judges the proposition that her awareness is incomplete is supported by historical induction.

⁸ 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.

Proposition 5 Fix a two-player game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$ such that the restricted game $\Gamma_{\tilde{Z}(h)}$ is not a game of common awareness. Then $h \models_{\Gamma} n_i(q_{\exists}u_iq)$

Proof. Let $i = P(h)$, and denote the other player by j . Since $\Gamma_{\tilde{Z}(h)}$ is not a game of common awareness, $\exists h' \in H_{\tilde{Z}(h)}$, with $P(h') = j$, such that $\Gamma_{\tilde{Z}(h') \cap \tilde{Z}(h)} \sqsubset \Gamma_{\tilde{Z}(h)}$. Hence, $h \models_{\Gamma} n_i(q_{\exists}u_jq)$. But as there are only two players, this means that $h \models_{\Gamma} n_i(q_{\exists}u_kq)$ holds for (all) $k \neq i$ which in turn implies $h \models_{\Gamma} n_i(q_{\exists}u_iq)$ as required. ■

Proposition 5 is similar in its structure to Proposition 4. Given the belief on the part of i that j is unaware of some propositions, inductive reasoning supports the proposition that i is also 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 games perceived by the player $P(h)$ in the past were restrictions of the game she currently perceives to be playing, and this is known to her at h . Similarly, other players from player $P(h)$'s perspective at h , must be imputed to perceive games that are restrictions of $\Gamma_{\tilde{Z}(h)}$.

In summary, although the representation of a game with awareness presented here does not allow for knowledge of unawareness in the standard modal logical, inductive reasoning and symmetry arguments can provide a basis for a judgement that there exist unconsidered possibilities of which other players may be aware.

On the other hand, in a finite extensive game, players eventually reach a terminal history, at which point they are aware of the full history of the game, at least insofar as it is relevant to the payoff they receive. The fact that the game includes unrealised histories of which the player may remain unaware is no longer relevant.

This point may be extended to suggest an inductive basis for identifying a ‘small world’, in which it is reasonable for players to disregard the possible existence of unconsidered contingencies, and act on the basis of Bayesian decision theory applied to the game they perceive. The crucial requirement for a small world is not that the decisionmaker should be aware of all possible contingencies (this requirement can never be met) but that she should be aware of all contingencies relevant to the outcome of a particular decision. This condition is trivially (but unhelpfully) satisfied at any terminal history of a game, since this is the point at which the outcome is actually realised and players receive their payoffs.

More generally, given that unawareness arises from bounded rationality, it seems reasonable to suppose that the more complex the game, the more likely it is that there exist unconsidered contingencies. Further, for a history h , the more complex the continuation of the game from the information set the player perceives to be in at h , the more likely it is that there exist unconsidered contingencies. Since games that have reached a terminal history are minimally complex in this sense, it seems plausible to argue that, as a game (as perceived by the player) ‘approaches the

end', it becomes easier for the player to consider all payoff-relevant contingencies.

Based on their past experience of the game (or, not modelled here, drawing on experience of similar games) players may reason inductively as follows:

“In the past, I have found myself able to consider all relevant possibilities that might arise in the next n moves. That is, previously unforeseen contingencies have come to my attention at least n moves before they were actually realised. Hence, if the continuation of the game I now perceive as relevant to my choices will end in no more than n moves, induction supports the conclusion that I am aware of all relevant contingencies.”

We begin by defining a proposition p as relevant (denoted $\rho_i p$) at some information set $\mathcal{I}(h)$ if the truth or falsity of p affects the continuation value of the game at $\mathcal{I}(h)$. We can now formalize the proposition that player i is aware of all relevant propositions at h as

$$\forall p (\rho_i p \Rightarrow a_i p) \tag{4}$$

If (4) holds at $\mathcal{I}(h)$, the continuation of the game at $\mathcal{I}(h)$ may be regarded as a ‘small world’ in the sense of Savage.

Next we need a measure of the complexity of the continuation of the game. There is a large literature on the complexity of games, and a variety of measures have been proposed (state-space complexity, game tree size, decision complexity and so on). These measures are developed for games of full awareness (and mostly for games of perfect information). Fortunately for our purposes, we require only that the complexity measure should satisfy some elementary properties.

We will use $\xi_\Gamma(\mathcal{I})$ to describe a complexity measure of the continuation of the game Γ from the information set \mathcal{I} , and impose the properties:

C.1. Zero property: the measure should take the value zero at terminal histories.

C.2. Declining with learning: For any pair of information sets \mathcal{I} and \mathcal{I}' in Γ , $\mathcal{I} \neq \mathcal{I}'$, and $P(h) = P(h') = i$ for all $h \in \mathcal{I}$, all $h' \in \mathcal{I}'$ and some $i \in N$, if there exists $h \in \mathcal{I}$ and $h' \in \mathcal{I}'$ such that $h \preceq h'$ then $\xi_\Gamma(\mathcal{I}) > \xi_\Gamma(\mathcal{I}')$.

C.3. Increasing with awareness. If $\Gamma \sqsubseteq \Gamma'$, then $\xi_{\Gamma'}(\mathcal{I}') \geq \xi_\Gamma(\mathcal{I})$ for any $\mathcal{I} \subseteq \mathcal{I}'$.

Property C.1 requires that all terminal histories are of equal and minimal complexity. Within a given game, property C.2 entails that the complexity of the continuation game must decline whenever a player makes a move (thereby eliminating some possible histories for that player). Property C.3, on the other hand, requires that increasing awareness will, in general, increase the complexity of the game.

We propose the following rule, which may be justified by historical induction over previous games

Consider the following inductive principle

Small World Principle: Treat the proposition $\forall p (\rho_i p \Rightarrow a_i p)$ as inductively justified whenever $\xi_{\Gamma_{\bar{Z}(h)}} (\mathcal{I}_{\bar{Z}(h)}(h)) \leq \xi^*$ for some constant ξ^* .

The Small World Principle cannot be guaranteed to avoid surprises. Sophisticated players will understand this by virtue of Proposition 4 on inductive justification of unawareness. Nevertheless, use of the Small World Principle for appropriate choices of ξ^* may lead to the adoption of ecologically rational heuristics, in a sense that will be made precise below.

A related approach may be applied to reasoning about the awareness of other players. In general, considerations of symmetry justify the conclusion that, if i is aware of possibilities of which j is unaware, the converse is also true. In the language of complexity, we have no general reason to suppose that ξ_i^* is either greater than or less than ξ_j^* . But in specific cases, i may indeed be more capable of reasoning about the problem at hand, so that $\xi_i^* > \xi_j^*$. In this case, i may conclude that she is aware of all relevant possibilities even though j is not.

4.3 Syntax and semantics of inductive reasoning about awareness

Space does not permit a full syntactic rendition of the model described above, in which players hold inductively justified beliefs about their own awareness or lack of it. Such a rendition has been developed by Walker (2011). In Walker’s rendition, the agent’s beliefs about whether or not she is fully aware in any state (or history) depend on a conjecture about awareness that does not require ‘a comparison between her actual awareness and the domain of ‘things to be aware of’. As Walker observes, a process of inductive reasoning such as that developed here provides a natural basis for beliefs about whether or not the agent is fully aware of all relevant propositions, leaving beliefs about particular propositions to be determined by the standard process of deduction from what is true in states that are deemed possible given the available information set.

Walker develops this two-stage structure in detail and derives a semantic structure corresponding to a language in which the string (in our notation) $q \forall a_i q$ stands for the claim ‘the agent is fully aware’. Walker shows that appropriate axiomatizations for this language are sound and complete for semantic structures in which agents’ beliefs about their awareness may be supposed to be derived from inductively justified conjectures.

Walker’s approach may be adapted to the extensive-form setting of the present problem. In the context of an extensive-form game, what matters is not awareness of all propositions, but awareness of all relevant propositions, as defined above. As shown above, player i can reach a subjective judgement, based on inductive reasoning, as to whether they are, or are not, aware of all relevant contingencies at any history h where $P(h) = i$, that is, whether $h \models \forall p (\rho_i p \Rightarrow a_i p)$. This enables us to define a semantic structure incorporating reasoning about awareness.

Now following the approach of Walker (2012) (but with a change of notation to fit our extensive form game model) we define the subjective interpretation relation \models_{Γ}^* which coincides

with the objective relation \models_{Γ} except for propositions regarding beliefs about awareness. The critical change is that the interpretation relation $h \models_{\Gamma}^* b_i (\forall p (\rho_i p \Rightarrow a_i p))$ holds if and only if the proposition $\forall p (\rho_i p \Rightarrow a_i p)$ is supported by inductive judgement for i at h , and more importantly, $h \models_{\Gamma}^* b_i \neg (\forall p (\rho_i p \Rightarrow a_i p))$ if the proposition $\neg (\forall p (\rho_i p \Rightarrow a_i p))$ is supported by inductive judgement for i at h .

Following Walker (2012) it would be possible to construct a characterization of the semantic structures consistent with this interpretation relation. However, this is beyond the scope of the present paper. Instead, our focus will be on the way in which inductively derived subjective beliefs about awareness and unawareness may be used to guide individual decisions and strategic choices. The key idea is that of heuristics.

5 Games Subject to Heuristic Constraints

We now develop a basis for making decisions that depend upon judgements about propositions that cannot be expressed explicitly in the language of the game available to that player but are nonetheless supported either on historical inductive grounds or by induction over individuals. We refer to these to as heuristic constraints. Conversely, we consider heuristics that may be applied to justify the adoption of a best response for the continuation of the game from an information set. Combining these heuristics, we obtain a definition for a game subject to heuristic constraints.

A heuristic constraint for a particular player associated with the game they perceive to be playing is an admissibility rule precluding the adoption by that player and her opponents of certain strategies whenever a certain proposition that involves the player or her opponents 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 game with awareness subject to heuristic constraints as follows.

Definition 18 *Fix a game with awareness \mathcal{G} . An awareness-based heuristic \mathcal{H} for the game \mathcal{G} is an admissibility rule such that, for each h in H in which $P(h) \neq c$, $\mathcal{H}(h) \subseteq \tilde{Z}(h)$ is a subset of the terminal histories available in the game $\Gamma_{\tilde{Z}(h)}$ from the information set $\mathcal{I}_{\tilde{Z}(h)}(h)$. The restriction $\mathcal{H}(h)$ depends only on the perceived game $\Gamma_{\tilde{Z}(h)}$, and on the judgement of the player at h as to whether she and/or her opponents are aware of all relevant possibilities and generates a restricted game $\Gamma_{\mathcal{H}(h)} \sqsubseteq \Gamma_{\tilde{Z}(h)}$. We denote by $(\mathcal{G}, \mathcal{H})$ the game with awareness \mathcal{G} that is subject to the heuristic constraints \mathcal{H} . We denote by \mathcal{H}_i the restriction of \mathcal{H} to the set of*

histories $\{h : P(h) = i\}$.

Imposing heuristic constraints on a game with differential awareness means that the perceived game in which the heuristic constraints may apply, is modified by removing 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.

We revise the definition of a strategy profile, so that in the game with awareness \mathcal{G} subject to the heuristic \mathcal{H} , at each h the continuation behavioral strategy profile is defined for the game $\Gamma_{\mathcal{H}(h)}$.

The definitions of subjective probabilities, assessments and sequential rationality from subsection 2.4 can be modified accordingly with each instance of the game $\Gamma_{\bar{Z}(h)}$ for history h , being replaced by the (in general even more restricted) game $\Gamma_{\mathcal{H}(h)}$ 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 awareness subject to heuristic constraints. This definition requires that players, at each information set where they are called on to move, adopt a best response to the equilibrium strategies of the continuation game they perceive at that information set, excluding histories that are precluded by the adoption of their own heuristic constraints and those of other players. This is a natural generalization of the standard assumption of common knowledge of rationality, which applies to the game \mathcal{G} .

As an immediate corollary to Proposition 2 we have:

Corollary 6 *A sequential equilibrium exists for any game with differential awareness subject to heuristic constraints.*

As discussed above, we will confine attention here to the case when the equilibrium is unique.

5.1 Ecological rationality

A variety of criteria have been suggested for the adoption and evaluation of decision heuristics. Goldstein and Gigerenzer (2002) propose a concept of ecological rationality for heuristics, explained as ‘the capacity of the heuristic to exploit the structure of the information in natural environments’. As an example, they consider the ‘recognition heuristic’ illustrated by the idea that a decision-maker asked to estimate which of two cities has a larger population should choose one they have heard of in preference to one they have not heard of. In an environment where mentions of cities are positively correlated with their population, this heuristic is ecologically rational.

In a game with awareness, the concept of ecological rationality may be seen as the way a heuristic adopted by players with bounded awareness would be evaluated from the perspective of a more aware player, or an external modeller aware of the maximal game. In a game with increasing awareness, a player may evaluate the ecological rationality of their own heuristics, applied early in the game, from the more aware perspective available later in the game.

It is not necessary to confine attention to the actual outcome realised subsequent to the adoption of a heuristic. From the perspective of a more aware player (or from the perspective of a player considering their own earlier state of awareness) it is possible to compare the decision recommended by a heuristic to the optimal decision recommended by Bayesian decision theory for the full game, and also to the ‘naive’ Bayesian decision derived for the restricted game. Even if, by virtue of an unlucky draw by nature, the heuristic yields a bad outcome in a particular play of the game, it may nonetheless be judged to be superior, in expected payoff terms, to the naïve Bayesian decision.

Consider a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$, subject to heuristic constraints \mathcal{H} . Let \mathcal{H}^{-i} be the heuristic obtained from \mathcal{H} by removing any constraints for information sets controlled by player i . That is, $\mathcal{H}^{-i}(h) = \mathcal{H}(h)$, if $P(h) \neq i$ and $\mathcal{H}^{-i}(h) = \tilde{Z}(h)$, if $P(h) = i$. We adopt the following definition.

Definition 19 Consider a game with awareness $\mathcal{G} = (\Gamma, \tilde{Z}(\cdot))$, subject to heuristic constraints \mathcal{H} . The heuristic constraints \mathcal{H} are ecologically rational for player i if

$$V_{\Gamma}^i(r') \geq V_{\Gamma}^i(r)$$

where r (respectively, r') is the behavioral rule generated by the sequential equilibrium (β, μ) for $(\mathcal{G}, \mathcal{H})$ (respectively, sequential equilibrium (β', μ') for $(\mathcal{G}, \mathcal{H}^{-i})$.)

It is important to note that the ecological rationality of \mathcal{H} cannot be determined from within the $\Gamma_{\tilde{Z}(h)}$ available to player $i = P(h)$ at a particular history h . Ecological rationality can be determined, however, by an unboundedly rational external modeller with access to the full awareness game Γ . More relevantly from the perspective of actual players, the belief that a particular heuristic is ecologically rational in a given setting may be justified inductively on the basis of its past performance in similar settings.

5.2 Example Part 2:

We now return to the speculative trade example of Heifetz, Meier and Schipper (2006). By Proposition 5, for each party $i = 1, 2$, the proposition $\exists q \in \mathcal{L}_{\Gamma} : (u_j q \wedge a_i q)$ is true for $j \neq i$.

Moreover, in the given example, the proposition can be extended to

$$\theta(q_{\alpha^j}) \equiv \exists q \in \mathcal{L}_\Gamma \left(u_j q \wedge a_i q \wedge (q \Rightarrow q_{loss}^j) \right), \quad (5)$$

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 .⁹

Now let party i consider the proposition $\exists q \in \mathcal{L}_\Gamma : (u_i q \wedge a_j q)$, 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 $\exists q \in \mathcal{L}_\Gamma : (u_i q \wedge a_j q \wedge (q \Rightarrow q_{loss}^i))$, so we have for both $h^1 = \langle \alpha_0, \alpha_0 \rangle$ and $h^2 = \langle \alpha_0, \alpha_0, \alpha_1 \rangle$,

$$h^i \models_\Gamma n_i \theta(q_{\alpha^i}).$$

Let us take $\mathcal{R}(h) \subseteq \tilde{Z}(h)$ to be the set of terminal histories in the game $\Gamma_{\tilde{Z}(h)}$ that passes through $\mathcal{I}_{\tilde{Z}(h)}(h)$ and that includes an action $\hat{\alpha}^i \in A_{\tilde{Z}(h)}$ that leads to a loss in the event that some unforeseen proposition q holds. We can then define the heuristic $\mathcal{H}^i(h) = \tilde{Z}(h) - \mathcal{R}(h)$ that precludes the adoption of $\hat{\alpha}^i$ by party i . If either the buyer adopts the heuristic \mathcal{H}^1 or the owner adopts the heuristic \mathcal{H}^2 , then 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 better outcome, then the agent should choose α . As shown above, this principle leads the players to engage in trade, even though, in the minimal game of common awareness the trade generates zero surplus. Each party's own perceived game with awareness suggests a strictly positive surplus from trade for the player concerned, coming at the expense of the other party.

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 expression (5). Noticee that, evaluated in the fully aware game, (5) is in fact true for both individuals.

A potential problem with this analysis is that it might lead to the conclusion that individuals should never trade in the presence of differential awareness. However, this conclusion only arises in the case where there are no gains from trade, as in the given example. If both parties evaluate the transaction as mutually beneficial on the basis of their own awareness, there is, in general, no reason for them to conclude that this mutual benefit would not persist under full awareness.

⁹ 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).

In particular, the fact that party 2 is willing to engage in trade does not imply that she must be aware of a contingency that would make the trade less appealing to party 1.

6 Applications

In this section, we sketch two applications, in which heuristics may be used as a way of responding to unawareness. Both applications deal with problems that may be addressed in terms of individual decisions (that is, one-person games with Nature) or in a multi-player context. For ease of exposition, we will focus here on the case of individual decision.

6.1 Research and Discovery

In some circumstances, unforeseen possibilities are both desirable and essential. The allocation of funds for research and development (R&D) provides an example. It is inherent in the concept of pure research that investigators cannot predict what they may discover. By contrast, the uncertainties associated with the development phase of R&D are well understood in most cases, and may reasonably be modelled in a Bayesian framework. The contrast between research and development raises obvious difficulties in the allocation of funds. A common response is to require research proposals which purport to contain accurate predictions of the path of research for which funding is supported.¹⁰ However, this does not appear to be a satisfactory solution.

This problem may be addressed by modelling the situation as a game with awareness. For simplicity, we will abstract from the interactive aspects of the problem and consider a decision-maker playing a one-person game against nature with awareness. Let (β^0, μ^0) denote her (unique) sequentially rational and consistent assessment and let r_0 denote the associated behavioral rule generated by β^0 .

At some history h she is considering whether to allocate funds to one of a set of research projects, and if so which one. The alternative is to allocate the funds to a development project, where we assume there are no unconsidered outcomes. For ease of exposition (and without essential loss of generality) assume that $\mathcal{I}(h) = \mathcal{I}_{\bar{Z}(h)}(h) = \{h\}$. That is, the decisionmaker believes she is at history h (and from the perspective of the fully aware game she actually is at history h).

At h , the decisionmaker faces the set of available actions $A(h) = \{a_0, a_1, \dots, a_K\}$, leading to the set of histories $\{h_0, h_1, \dots, h_K\}$, where $h_k = \langle h \cdot a_k \rangle$. Here a_0 denotes the choice of the development project, while a_k , $k = 1, \dots, K$ are research projects. We assume that the complexity of the subgame starting at h_0 is less than ξ , while for $k > 0$, the complexity of the subgame starting at h_k is greater than ξ .

¹⁰ One common way of meeting this requirement is to seek funding for research that is effectively complete, but has not yet been published.

In the absence of unconsidered contingencies, let us suppose her unique best continuation strategy β_h^0 from the sub-game starting at h entails selecting the development option at h , that is, $\beta_h^0(h) = a_0$. We normalize by setting $V_{\Gamma_{Z(h)}}(\beta_h^0) = V_{\Gamma}(r_0) = 0$. That is, resources used for any project are evaluated in terms of the opportunity cost associated with the development option.

For each $k = 1, \dots, K$, let β_k denote the continuation strategy that agrees with β_h^0 at every information set in the subgame starting at h , except initially where $\beta_k(h)$ is equal to a_k . And let (β_{-h}^0, β_k) denote the corresponding modified strategy profile for the entire game with awareness.

Since β_h^0 is the unique best option for the subgame starting at h , it follows that $V_{\Gamma_{Z(h)}}^i(\beta_k) < 0$, for all $k = 1, \dots, K$. That is, based on current awareness, the development project yields higher expected returns than any of the research projects. Hence, the unconstrained heuristic \mathcal{H}^{-i} (which may also be denoted \mathcal{H}^0 in the context of individual decision) will lead to the choice of a_0 at h .

Now suppose she has sufficient experience to provide inductive support for some proposition of the form:

‘For problems of complexity greater than ξ , there may exist contingencies of which I am unaware.’

Hence the decisionmaker believes herself to be aware of all contingencies relevant after a choice of a_0 but not for any choice a_k , $k > 0$.

Then, given a ‘research budget’ of B , we may consider the heuristic \mathcal{H}^B :

“Exclude a_0 from consideration unless $V_{\Gamma_{Z(h)}}^i(\beta_k) < -B$, for all $k = 1, \dots, K$.”

This leads to the behavioral rule that at h prescribes the following: ‘Choose a_k (that is, undertake the research project k) that maximizes $V_{\Gamma_{Z(h)}}^i(\beta_k)$, provided $V_{\Gamma_{Z(h)}}^i(\beta_k) \geq -B$. Otherwise choose a_0 .

The heuristic allows for the choice of a research project rather than the ‘safe’ development option whenever the expected net cost, based on the possibilities already under consideration, is less than the research budget B . The quantity B might be considered as a value at risk constraint.

In addition to the benefits that may be evaluated in advance, the research project may generate valuable discoveries, associated with awareness of previously unconsidered possibilities. If we denote by r_k the behavioral rule generated by the strategy profile (β_{-h}^0, β_k) , then the full value of a project, which may be estimated by a fully aware outside observer, but not by the decisionmaker, is $V_{\Gamma}(r_k)$. In this context, we assume that, for all k , the benefits of discovery are non-negative, so that, $V_{\Gamma}^i(r_k) \geq V_{\Gamma_{Z(h)}}^i(\beta_k)$.

Proposition 7 *The following conditions are sufficient for \mathcal{H}^B to be ecologically rational:*

- (a) $\frac{1}{K} \sum_{k=1}^K (V_{\Gamma}^i(r_k) - V_{\Gamma_{Z(h)}}^i(\beta_k)) \geq B$; and
- (b) $(V_{\Gamma}^i(r_k) - V_{\Gamma_{Z(h)}}^i(\beta_k))$ is positively correlated with $V_{\Gamma_{Z(h)}}^i(\beta_k)$.

The first condition states that on average, the value of unanticipated discoveries exceeds the research budget B . Given this condition, the second condition ensures that the expected return from choosing the project with maximal expected net value (as perceived by the decisionmaker) will exceed B , and therefore be preferable to the development option.

As already noted, a crucial feature of ecological rationality is that it cannot be evaluated within the perspective available to a decisionmaker at h who perceives the game $\Gamma_{\bar{z}(h)}$. Conversely, the heuristic \mathcal{H}^B which may be judged as ecologically rational by an outside observer under the given conditions, can be implemented by the decisionmaker without requiring awareness of the full game Γ .

Thus, one interpretation of the heuristic is that it represents the advice that would be given by an unboundedly aware expert to a decisionmaker who must make decisions under conditions of bounded awareness.

6.2 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.

Grant and Quiggin (2012) consider the case of a decisionmaker, faced with an uncertain choice,

who has available a ‘safe option’, yielding a return that can be normalized to zero, and conditional on which the decisionmaker judges (by induction) that there are no relevant propositions of which she is unaware. Grant and Quiggin say that a behavioral rule other than the safe option is ‘subject to unfavorable surprises’ ‘there exist unconsidered moves by Nature, against which the payoff from the behavioral rule r is less than zero. They define the Strong Form of the Precautionary Principle as the heuristic that excludes behavioral rules subject to unfavorable surprises, and show that the Strong Version of the Precautionary Principle is ecologically rational if and only if $V(r_0) < 0$, where r_0 is the behavioral rule associated with the unconstrained heuristic \mathcal{H}^{-i} (which, in the context of individual decision may be denoted \mathcal{H}^0). Grant and Quiggin also define a weaker form of the Precautionary Principle and derive conditions for ecological rationality.

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.

7 Concluding comments

The model presented in this paper has incorporated the minimal deviations from the standard case of an extensive form game necessary to allow a representation of inductive reasoning about differential awareness and to model possible behavioral responses. The analysis presented here could be extended in a number of ways.

First, the model presented here allows players to become aware of previously unconsidered possibilities, but not to forget about possibilities they have previously considered. To the extent that unawareness reflects bounded rationality, this seems unreasonable, since the model accessible to players becomes steadily more complex. A desirable property for a representation of bounded awareness is that the bound should be determined by limits on reasoning capacity which should be constant over time or at least should not increase monotonically

It would be desirable, therefore, to extend the analysis to allow for imperfect recall. The simplest version of imperfect recall is to suppose that individual i at (perceived) information set $\mathcal{I}_{\bar{Z}(h)}(h)$ forgets about some or all histories that do not pass through $\mathcal{I}_{\bar{Z}(h)}(h)$, since these histories involve actions (by player i , some other player j , or Nature) known by i not to have been taken. In the context of individual decision, such histories are strictly irrelevant, but in a game-theoretic context, the fact that they have not been chosen may inform reasoning about the other players. More generally, individuals may forget the details of their past history.

In the present paper, the perception mapping is taken to be exogenous, subject only to the restriction that players must always consider some history possible. A natural extension would be to consider games in which awareness is derived naturally from the tree structure of the game. For example, we might impose the requirement that players are always aware of actions when they become available.

Finally, the relationship between equilibrium and awareness could be explored further. As we have argued, differential awareness means that players will, in general, arrive at information sets they previously regarded as being off-equilibrium. This would appear to rule out equilibrium concepts that admit arbitrary actions at off-equilibrium information sets, and therefore to mitigate the problem of multiplicity of equilibria. On the other hand, since changes in awareness imply, in general, changes in the perceived set of equilibria, difficulties may arise for notions such as forward induction, to the extent that they rely on the use of apparent deviations from equilibrium to make inferences about other players.

Much work remains to be done. However, the model presented in this paper allows for a coherent account of behavior in games where players are not, in general, aware of all possibilities and understand this fact. This seems likely to yield a more plausible account of actual behavior than a framework based on the assumption that all players are unboundedly rational and that this is a matter of common awareness.

A Appendix:

Proof of Proposition 2. By standard arguments (for example Osborne & Rubinstein [1994, p227]) it follows that an assessment is sequentially rational (β, μ) if and only if it satisfies the one-shot deviation property. That is, at each h , for the behavioral strategy $\beta_h^{P(h)}$ of player $P(h)$ in the continuation of the game $\Gamma_{\tilde{Z}(h)}$, there is no subsequent information set $\mathcal{I}_{\tilde{Z}(h)}(h')$, with $P(h') = P(h)$ in the continuation of the game at which a change in $\beta_h^{P(h)}(h')$ increases his payoff conditional on reaching $\mathcal{I}_{\tilde{Z}(h)}(h')$.

Therefore we will first establish the existence of a trembling-hand equilibrium for the agent-normal form of \mathcal{G} , which by the one-shot deviation property also constitutes a trembling hand equilibrium of \mathcal{G} . It will then suffice to show that for any trembling hand equilibrium strategy profile β , there exists a subjective probability system μ such that (β, μ) is a sequential equilibrium.

We take the *agent-normal form* of the game with awareness $(\Gamma, \tilde{Z}(\cdot))$, to be the game with awareness $(\Gamma^{an}, \tilde{Z}(\cdot))$, where Γ^{an} is the agent normal form Γ , in which there is one player for each information set in the extensive form game and where player h is imputed to be playing $\Gamma_{\tilde{Z}(h)}^{an}$ the agent normal form of $\Gamma_{\tilde{Z}(h)}$. For each h , denote the perturbation of the game $\Gamma_{\tilde{Z}(h)}^{an}$ by first fixing the strategies of all players h'' , such that there exists $h' \in \mathcal{I}_{\tilde{Z}(h)}(h)$ and $h'' \preceq h'$ and

$h'' \neq h'$ (i.e. player h'' is a player has already moved by the time the game reaches the information set $\mathcal{I}_{\bar{Z}(h)}(h)$) and then letting the set of actions of each player h'' , such that there exists history $h' \in \mathcal{I}_{\bar{Z}(h)}(h)$ and $h' \preceq h''$ (i.e. player h'' is a player who moves in the continuation of the game $\Gamma_{\bar{Z}(h)}^{an}$ after information set $\mathcal{I}_{\bar{Z}(h)}(h)$) be the set of mixed strategies in $A_{\bar{Z}(h)}(h'') \cap A_{\bar{Z}(h'')}(h'')$ that assign probability of at least $\varepsilon_{h''}^a(h)$ to each action that player h (at h) imputes to player h'' at her information set $\mathcal{I}_{\bar{Z}(h)}(h'')$. That is, this constrains h and every player who follows h to use every strategy h imputes that they have available with some minimal probability. Consider a sequence of such perturbed games in which $\varepsilon_{h''}^a(h) \rightarrow 0$, for all h , for all h'' and a ; by the compactness of the set of strategy profiles, some sequence of selections (β^k) from the sets of strategy profiles that are sequentially rational for all i , of the games $\{\Gamma_{\bar{Z}(h)}^{an} : h \in H\}$ converges to say β .

Take the sequence (β^k) . At each information set $\mathcal{I}(h)$ define the belief $\mu_{h'}$ for each h' in $\mathcal{I}(h)$, to be the limit of the beliefs defined from (β^k) using Bayes rule, (β, μ) is then by construction a consistent assessment. Since the strategies are completely mixed as can be done consistently, each information set consistent with each player's level of awareness is reached with positive probability and each agent's strategy is a best response when the beliefs at each information set are defined by μ . Thus (β, μ) is a sequential equilibrium. ■

References

- Fagin, R. and Halpern, J. (1988), "Belief, awareness and limited reasoning," *Artificial Intelligence* 34, 39–76.
- Fagin, R., Halpern, J., Moses, Y. and Vardi, M. (1985), *Reasoning About Knowledge*, MIT Press, Cambridge, MA.
- Galanis, S. (2011), "Syntactic foundations for unawareness of theorems," *Theory and Decision*, 71(4), 593-614.
- Gilboa, I. and Schmeidler, D. (2000), 'Case-based knowledge and induction', *IEEE Transactions on Systems*, 173-90.
- Goldstein, D. and Gigerenzer, G. (2002), "Models of Ecological Rationality: The Recognition Heuristic," *Psychological Review*, 109(1), 75-90.
- 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.
- 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.
- Grant, S. and Quiggin, J. (2012), "Bounded awareness, heuristics and the Precautionary Principle", Working Paper WP R12_3, Risk and Sustainable Management Group, University of Queensland.

- Halpern, J. (2003), *Reasoning About Uncertainty*, The MIT Press, Cambridge, Massachusetts
 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 Rêgo, L. (2008), ‘Interactive unawareness revisited’, *Games and Economic Behavior*, 62(1), 232–262.
- Halpern, J. and Rêgo, L. (2012), “Generalized solution concepts in games with possibly unaware players”, forthcoming *International Journal of Game Theory* 41(1), 131-155.
- 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.
- Heifetz, A., Meier, M. and Schipper, B. (2009), “Prudent Rationalizability in Generalized Extensive-Form Games”, Working Paper, UC Davis.
- Heifetz, A., Meier, M. and Schipper, B. (2012), “Unawareness, Beliefs, and Speculative Trade”, Working Paper, UC Davis.
- Kripke, S. (1963), ‘A semantical analysis of modal logic I: normal modal propositional calculi,’ *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9, 67–96.
- Laudan, L. (1981), ‘A confutation of convergent realism’, *Philosophy of Science*, 48(1), 19–49.
- Li, J., A Note on Unawareness and Zero Probability (January 2, 2008). PIER Working Paper No. 08-022. Available at SSRN: <http://ssrn.com/abstract=1152281>
- 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>.
- Modica, S. and Rustichini A. (1994), ‘Awareness and partitional information structures,’ *Theory and Decision* 37, 107–124.
- Osborne, M.J. and Rubinstein, A. (1994), *A Course in Game Theory*, MIT Press, Cambridge, Mass.
- Rumsfeld, D. (2002), Statement to Press, Washington, February.
- 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.




Figure 3: Γ^2 , the game the owner perceives she is playing, which starts at the chance node on the left.




Figure 4: Γ^3 , the game of common awareness for the two players, which starts at the chance node on the left.

PREVIOUS WORKING PAPERS IN THE SERIES

RISK AND UNCERTAINTY PROGRAM

- R03_1 Bargaining power and efficiency in principal-agent relationships, Robert G Chambers and John Quiggin (August 2003).
- R03_2 Information and the risk-averse firm, Robert G. Chambers and John Quiggin (2003).
- R03_3 Narrowing the No-Arbitrage Bounds, Robert G. Chambers and John Quiggin (2003).
- R03_4 Separability of stochastic production decisions from producer risk preferences in the presence of financial markets, Robert G Chambers and John Quiggin (2003).
- R03_5 Comparative statics for state-contingent technologies, Robert G Chambers and John Quiggin (2003).
- R03_6 Dual structures for the sole-proprietorship firm, Robert G Chambers and John Quiggin (December, 2003).
- R04_1 Output price subsidies in a stochastic world, Robert G Chambers and John Quiggin (2004).
- R04_2 Supermodularity and risk aversion, John Quiggin and Robert G Chambers (March 2004).
- R04_3 Linear-risk-tolerant, invariant risk preferences, Robert G Chambers and John Quiggin (April, 2004).
- R04_4 Increasing uncertainty: a definition, Simon Grant and John Quiggin (May 2004)
- R04_5 Supermodularity and the comparative statics of risk, John Quiggin and Robert G Chambers (June, 2004).
- R04_6 Fixed wages and bonuses in agency contracts: the case of a continuous state space, Maria Racionero and John Quiggin (June, 2004).
- R04_7 Games without rules, Flavio Menezes and John Quiggin (July, 2004).
- R04_8 The risk premium for equity: implications for resource allocation, welfare and policy, Simon Grant and John Quiggin (August, 2004).
- R04_9 Capital market equilibrium with moral hazard and flexible technology, John Quiggin and Robert G Chambers (September, 2004).
- R05_1 The state-contingent approach to production under uncertainty, John Quiggin and Robert G Chambers (January, 2005).
- R05_2 Economists and uncertainty, John Quiggin (April, 2005).
- R05_3 Cost minimization and asset pricing, Robert G Chambers and John Quiggin (2005).
- R05_4 Consistent Bayesian updating with unconsidered propositions, Simon Grant and John Quiggin (February, 2005).

- R05_5 Comparative risk aversion for state-dependent preferences, John Quiggin and Robert G Chambers (May, 2005).
- R05_6 Outcomes and strategy choices in Tullock contests, Flavio Menezes and John Quiggin (March, 2005).
- R05_7 Learning and Discovery, Simon Grant and John Quiggin (July 2005).
- R06_1 Dual approaches to the analysis of risk aversion, Robert G Chambers and John Quiggin (June, 2006).
- R06_2 Efficiency analysis in the presence of uncertainty, Chris O'Donnell, Robert G Chambers and John Quiggin (May, 2006).
- R06_3 Lost in translation: honest misunderstandings and ex post disputes, Simon Grant, Jeff Kline and John Quiggin (August, 2006).
- R07_1 Estimating complex production functions: The importance of starting values, Mark Neal (January, 2007).
- R07_2 Risk and derivative price, Yusuke Osaki (2007).
- R07_3 A risk-neutral characterization of optimism and pessimism, and its applications, Yusuke Osaki and John Quiggin (2007).
- R07_4 Can game theory be saved? Flavio Menezes and John Quiggin (August, 2007).
- R07_5 Bargaining power and efficiency in insurance contracts, John Quiggin and Robert G Chambers (April 2007).
- R07_6 Sharp and diffuse incentives in contracting, Flavio Menezes and John Quiggin (2007).
- R07_7 Markets for influence, Flavio Menezes and John Quiggin (July, 2007).
- R07_8 Event-specific data envelopment models and efficiency analysis, Robert G. Chambers, Atakelty Hailu and John Quiggin (2007).
- R08_1 Generalized invariant preferences: two-parameter representations of preferences by Robert G Chambers and John Quiggin (February, 2008).
- R08_2 Bounded rationality and small worlds, Simon Grant and John Quiggin (June, 2008).
- R09_1 Inductive reasoning about unawareness, Simon Grant and John Quiggin (June 2009).
- R09_2 Markets for influence, Flavio M. Menezes and John Quiggin (September, 2009).
- R09_3 A matter of interpretation: bargaining over ambiguous contracts, Simon Grant, Jeff Kline and John Quiggin (November, 2009).
- R10_1 The computation of perfect and proper equilibrium for finite games via simulated annealing, Stuart MacDonald and Liam Wagner (April, 2010).
- R10_2 Uncertainty and technical efficiency in Finnish agriculture: a state-contingent approach, Céline Nauges, Christopher O'Donnell and John Quiggin (2010).
- R10_3 Production under uncertainty: a simulation study, Sriram Shankar, Chris O'Donnell and John Quiggin, (2010).
- R10_4 Economics as a social science: financial regulation after the crisis, John Quiggin (2010).

- R11_1 More competitors or more competition? Market concentration and the intensity of competition, Flavio M. Menezes and John Quiggin (August 2011).
- R11_2 A two-parameter model of dispersion aversion, Robert G Chambers, Simon Grant, Ben Polak and John Quiggin (August 2011).
- R11_3 Intensity of competition and the number of competitors, Flavio M Menezes and John Quiggin (June, 2011).
- R11_4 A matter of interpretation: ambiguous contracts and liquidated damages, Simon Grant, Jeff Kline and John Quiggin (September, 2011).
- R11_5 Capabilities as menus: A non-welfarist basis for QALY evaluation, Han Bleichrodt and John Quiggin (October, 2011).
- R12_1 Differential awareness ambiguity, and incomplete contracts: a model of contractual disputes, Simon Grant, J Jude Kline and John Quiggin (December, 2012).
- R12_2 Optimal access regulation with downstream competition, Tina Kao, Flavio Menezes and John Quiggin (March, 2012).
- R12_3 Inductive reasoning about unawareness, Simon Grant and John Quiggin (April, 2012).
- R12_4 Bounded awareness, heuristics and the precautionary principle, Simon Grant, John Quiggin (April, 2012).
- R12_5 A matter of interpretation: ambiguous contracts and liquidated damages, Simon Grant, Jeff Kline and John Quiggin (May, 2012).
- R12_6 Meeting the competition: commitment and competitive behaviour, Tina Kao, Flavio Menezes and John Quiggin (December 2012).
- R13_1 Inferring the strategy space from market outcomes, Flavio Menezes and John Quiggin (January, 2013).
- R13_2 The value of information and the value of awareness, John Quiggin (January, 2013).
- R13_3 A theory of strategic interaction with purely subjective uncertainty, Simon Grant, Idione Meneghel and Rabee Tourky (June, 2013).
- R13_4 Ordinal, nonlinear context-dependence, Patrick O'Callaghan (May, 2013).