

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

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

Help ensure our sustainability.

Give to AgEcon Search

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

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

Risk & Sustainable Management Group	
Risk & Uncertainty Program Working Paper: R05_7 (revised 27 July 2006)	
wship	Learning and Discovery
tion Fello tm	Simon Grant
Federat ation.h	Rice University
ouncil o	and
Research supported by an Australian Research Co http://www.arc.gov.au/grant_programs/discovery	John Quiggin Australian Research Council Federation Fellow, University of Queensland
Schools of Economics and Political Science University of Queensland Brisbane, 4072 rsmg@uq.edu.au http://www.uq.edu.au/economics/rsmg	

Learning and Discovery^{*}

Simon Grant Department of Economics, Rice University

John Quiggin School of Economics, University of Queensland

27 July 2006

Abstract

We formulate a dynamic framework for an individual decisionmaker within which discovery of previously unconsidered propositions is possible. Using a game-theoretic representation of the state space as a tree structure generated by the actions of agents (including acts of nature), we show how the existence of unconsidered propositions can be represented by a coarsening of the state space. Furthermore we develop a syntax rich enough to describe the individual's awareness that currently unconsidered propositions may be discovered in the future. We consider quantified beliefs derived as subjective probabilities conditional on implicit beliefs about unconsidered propositions. We derive conditions under which a Bayesian learning approach can be applied to a subset of known propositions. We show that our model of quantified beliefs encompasses the case when individuals observe events previously considered impossible, and discuss the implications for the endogenous discovery of new propositions.

JEL Classification: D80, D82

Key words: unawareness, bounded rationality

^{*}We thank, without implication, Michele Cohen, Larry Epstein, Itzhak Gilboa, Joseph Halpern, Edi Karni, Jeff Kline, Mamuro Kaneko, Peter Klibanoff, Benoît Menoni, Robert Nau, Klaus Nehring, Matthew Ryan, David Schmeidler and participants in the FUR06 and RUD06 conferences for helpful comments and criticism. The usual disclaimer applies.

Every solution of a problem raises new unsolved problems; the more so the deeper the original problem and the bolder its solution. The more we learn about the world, and the deeper our learning, the more conscious, specific, and articulate will be our knowledge of what we do not know, our knowledge of our ignorance. For this, indeed, is the main source of our ignorance - the fact that our knowledge can only be finite, while our ignorance must necessarily be infinite. *Popper, On the Sources of Knowledge and of Ignorance, 1963*

1 Introduction

Economists and decision theorists have developed a powerful dynamic theory of learning, based on Bayesian epistemology. Beginning with a prior probability distribution over a set of states, information, taking the form of a sequence of observations on partitions of the state space, induces a posterior conditional distribution. Associated with this theory of learning is a normative and positive theory of decision, including an analysis of the valuation of information, based on maximization of subjective expected utility theory. There is, however, no corresponding theory of discovery, concerning the way in which the state space itself may be refined and revised, or, in epistemological terms, how learning raises new unsolved problems.

The crucial feature of Bayesian learning is the steady reduction of uncertainty. Each signal realization eliminates possible states of nature, reducing ultimately, with sufficient information, either to a discrete realized state of nature or to an event of arbitrarily small prior measure. Yet, as Popper (1963) has noted, echoing proverbial wisdom, learning commonly expands the domain of our ignorance. The problem of 'unknown unknowns', famously referred to by US Defense Secretary Rumsfeld, plagues any attempt to give a complete formal account of our uncertainty about the world.

In decision-theoretic terms, the problem is that the set of states of nature available for consideration by any decision-maker is necessarily incomplete. This point may be made either by postulating the existence of a subset of states that are omitted from consideration, or by treating the subjective space as a coarsening of some more refined objective partition of the state space into possible events. Approaches to this problem include Cubitt and Sugden (2001), Dekel, Lipman, and Rusticchini (2001), Epstein and Marinacci (2006), Ghirardato (2001), Grant and Quiggin (2006), Heifetz, Meier, and Schipper (2006), Kreps (1979, 1992), Modica and Rusticchini (1994, 1999), Mukerji (1997) and Nehring (1999). Related work in a modal-logical context includes that of Halpern (2001), Halpern and Rêgo (2005, 2006a,b). Alternative approaches, dispensing completely with the state space, include Gilboa and Schmeidler (1995) and Karni (2005).¹

The significance of bounded rationality may be illustrated by the work of Aragones, Gilboa, Postlewaite and Schmeidler (2005), who make the point that the complexity of standard decision problems, such as the specification of a linear regression model, is, in general, so great that humans with bounded calculation capacity cannot possibly consider all hypotheses that might be relevant in the search for an optimal solution. It follows that learning is possible without new data. Someone working with a locally optimal model might, for example by talking with another person concerned with the same problem, be led to consider a different and superior model. This finding is in contrast with the Bayesian case applicable under unbounded rationality when the optimal model for a given data set is always known, and when learning can take place only as a result of new data and raises the question of when, if at all, Bayesian updating is reasonable for decision-makers with bounded rationality.

A closely related observation is that of Maskin and Tirole (1999), who show that incomplete contracts cannot be represented merely by making knowledge about the state space private and unverifiable. As long as agents can define probability distributions over outcomes and undertake dynamic programming, they can achieve first-best outcomes. This observation raises the question of how best to represent bounded rationality and incomplete knowledge of the state space. A consequent question, raised by Halpern and Rêgo (2005), is how to represent individuals' awareness of the limits imposed by their own bounded rationality.

¹Kaneko & Kline (2006) also develop a dynamic model of decision makers with limited cognitive abilities who only have a imprecise view of the environment in which they operate but the emphasis in their model is on how the raw experiences of an individual might develop from short-term memories to long-term memories and how these are used to construct a personal view of the world.

The aim of this paper is to address these questions. We claim that a complete theory of belief and choice under uncertainty must incorporate the discovery of new, previously unconsidered propositions and states of nature, and that this process can only be modelled adequately by taking account of the bounded rationality of economic agents.

The paper is divided into two parts. In Part I, we develop a dynamic model of learning and discovery, sufficiently powerful to encompass a formal presentation of statements such as that of Popper quoted above, as well as precise statements about awareness of existential propositions. In particular, we show how to describe an individual's awareness that currently unconsidered propositions may be discovered in the future, and show discovery of previously unconsidered propositions is possible in our framework. In Part II, we introduce quantified notions of beliefs. We derive conditions under which such beliefs may be defined and updated using the standard Bayesian model, and extend this characterisation to encompass more general models of belief such as the multiple priors model of Gilboa and Schmeidler (1989). Using the concept of implicit beliefs about unconsidered propositions, we show how individuals may have 'impossible beliefs', and how the falsification of implicit beliefs may lead to endogenous discovery of new propositions.

Part I: A dynamic model of learning and discovery

The idea of discovery may naturally be expressed either in terms of the discovery of new propositions expressed in the language available to us or in terms of a refinement of the state space we consider. The former approach, which we will term 'syntactic', is intuitively natural, and allows the use of the tools of modal logic, such as knowledge and awareness operators, as in the work of Halpern and Rêgo (2005, 2006). The latter approach fits more naturally with economic models of choice under uncertainty and permits the construction of an explicit dynamic representation. We therefore develop both approaches in parallel, with a framework that allows direct translation between semantic and syntactic constructs.

A crucial analytical device in achieving our representation of bounded rationality is the adoption of the perspective of an unboundedly rational, but not necessarily perfectly informed, external observer, as in Grant and Quiggin (2006). The external observer may be interpreted as having the beliefs and knowledge that would be available to the agent in the absence of bounded-rationality constraints. Thus, the language available to the external observer is rich enough to express propositions unavailable to the agent, and similarly, the state space is rich enough to allow modelling of the development of the agent's knowledge and awareness, including the discovery of new propositions. A similar approach is adopted by Halpern and Rêgo (2006b).

In section 2, we present an example based on the work of Aragones et al. (2005) and a second example showing how unexpected observations can generate new scientific discovery. In section 3 we present a tree-structure representation of dynamic uncertainty and notation for modal-logical representations of knowledge, expressibility and uncertainty of propositions. In Section 4, drawing on Heifetz, Meier and Schipper (2006), we represent expressibility of propositions using the notion of a lattice of tree structures ordered by refinement, or, alternatively by the scope of the set of propositions available to individuals. A crucial feature of the structure is the inclusion of a maximal element, corresponding to the external representation of the problem that would be available to an unboundedly rational observer with a given information set.

The main new contribution of Part I is presented in 5 and 6. In these sections, we show how the dynamics of learning and discovery can be represented within the lattice structure derived previously, and how the modallogical framework of knowledge can be extended to incorporate the notion that individuals may be aware that there exist unconsidered propositions which they might subsequently discover, or which might be known to others. The crucial innovations are the use of existence quantifiers over domains of propositions and the introduction of an awareness operator. We show that such an extension is necessary in that an individual cannot know (in the modal-logical sense) of the existence of unconsidered propositions.

2 Examples

2.1 Fact-free learning

We first consider the notion of 'fact-free' learning as described by Aragones et al. Consider the general case examined by Aragones et al. (2005) of a data set consisting of a variable of interest Y and m potential explanatory variables $(X_1, X_2...X_m)$ where the object is to derive an Ordinary Least Squares regression model having at most K explanatory variables, and maximizing (conditional on K) the coefficient of determination R_K^2 . Aragones et al show that this problem is, in general, (NP-)'hard' in the sense of complexity theory.

For concreteness, we may consider the problem discussed by Aragones et al, of deriving a model to predict economic growth. Take as a starting point, an econometrician who has estimated a model of economic growth with K = 3, where the explanatory variables are, say, initial income (negative and significant), initial stock of physical capital (positive and significant) and dominant religious affiliation (insignificant). In this position, the econometrician's beliefs might be summarized by a probability distribution over the parameter space for the relevant regression, with the distribution function being derived in the standard Bayesian fashion, assuming an initially diffuse prior.

In an attempt to improve the model, the econometrician might either consult the theoretical literature or examine the residuals. The theoretical literature might suggest the inclusion of a measure of initial human capital. Examination of the residuals might show that Hong Kong was a positive outlier and lead to estimation of a new model, incorporating measures of openness to trade. Since the data set is fixed, no new information is acquired in this process: rather new inferences are undertaken with existing information. In each case, the econometrician's beliefs may be summarized in terms of a probability distribution over two parameter spaces, one for the old model and one for the new one.

Now consider an external perspective, that of an unboundedly rational observer with access to the same data as the econometrician. The observer would be able to formulate all possible hypotheses regarding the data generating process (Hendry 1987) that produces the set (X, Y), including as a

subset of this class of hypotheses all linear models of the relationship between Y and the explanatory variables X. For such an observer, fact-free learning would not be possible. However, new information would allow Bayesian updating of the probability distribution of the parameters of the data generating process. A central idea of the present paper is to consider the knowledge available to a boundedly rational decision-maker relative to that available to an unboundedly rational external observer.

2.2 Experimental discovery

We now consider an example where new hypotheses arise from the occurrence of previously unconsidered events. The example is based on a highly stylized version of research into atomic and sub-atomic structure during the 20th century. The agent, a scientific researcher, is concerned to discover evidence supporting or rejecting the 'solar system' model of the atom, with a nucleus orbited by electrons. He considers undertaking an experiment, in which an electron beam is fired at a thin sheet of metal. Passage of electrons through the metallic sheet is taken as evidence for the solar system model and absorption of electrons by the sheet is taken as being inconclusive. Alternatively the agent may take a default course of action, such as pursuing some other line of research. In a standard decision model, we might represent this by the decision tree in Figure 1.

FIGURE 1 NEAR HERE

As illustrated there are three points in time, referred to as instants, t = 0, 1, 2. At t = 0 there is a decision node (labeled 0) for the researcher, who is denoted as player 1. At t = 1, conditional on the researcher's decision to undertake the experiment, there is a chance node (decision of Nature, denoted as player 0). At each node, we will denote a move to the left by 'L' and a move to the right by 'R'. Where there is no decision, the single default move is denoted R.

The situation as described is one of imperfect information. If the researcher undertakes the experiment, moving to the node R, he does not know whether Nature will choose R (passage of electrons) or L (absorption). However, he has available a well-developed theory of the value of information, allowing individuals to calculate the optimal choice in situations of this kind, where it can be supposed that subsequent decisions will be influenced by the information (if any) observed at t = 1.

We now consider a more radical form of uncertainty, in which there is a possibility that is unforeseen by the researcher. This is the possibility that the electron beam will interact with particles in the nucleus, producing previously unobserved emissions such as gamma rays. For the purposes of the example, we shall assume that the emission of gamma rays means that no useful information is obtained regarding passage or absorption of electrons.

The possibility of gamma ray emission may be represented by adding an extra instant, three extra chance nodes and an extra terminal node, as in Figure 2. Now, Nature's move at t = 1 determines whether gamma rays will be emitted (*R*) or not (*L*). If gamma rays are not emitted, the experiment proceeds as before at t = 2, producing passage or absorption. In this case, the researcher does not become aware of the unrealized possibility of gamma ray emissions.

FIGURE 2 NEAR HERE

If gamma rays are emitted, the researcher is confronted with new, previously unconsidered, possibilities at t = 3, and may formulate new hypotheses to explain the unexpected observation. In particular, the researcher, or others may be led to consider the hypothesis that protons, neutrons and electrons are not fundamental building blocks of matter but are themselves made up of smaller subatomic particles.

Processes of this kind occur not only in scientific research but in day-today economic activity. For example, entrepreneurs discover and exploit new market opportunities that others have not previously considered. The aim of this paper is to provide a framework within which such activities may be modeled and analyzed.

3 Structure and notation

We use a dynamic tree structure based on that of an extensive form game between Nature and a boundedly rational individual, with additional structure intended to permit the use of semantic and modal logic operators. Notation is drawn from Belnap, Perloff and Xu (2001) and Halpern (2003), along with standard extensive-form game theory. The model is explicitly finite, as regards the number of possible nodes in the tree structure, semantically distinct propositions, time periods and so on. This turns out to be technically convenient in a number of respects. More significantly, it reflects the fact that the model deals with decisions made by a finite, boundedly rational individual, who can act only in discrete time.

3.1 Nodes and Trees.

The elementary units of this model are a set of nodes N, n = 0, ..., N, partially ordered under a relation \prec into a tree structure τ . The corresponding weak order is \preceq , where $n_1 \leq n_2$ is interpreted as $(n_1 \prec n_2) \lor (n_1 = n_2)$. Nodes are binary (two successors), unary (one successor) or terminal (no successors). Two distinct nodes n_1, n_2 are comparable if either $n_1 \prec n_2$ or $n_2 \prec n_1$.

The main interest is in binary nodes. All such nodes are treated as representing decisions, either by the individual decision-maker or by Nature.² As discussed in more detail, below, at each such node n, the agent controlling the node, either the individual or Nature, makes a decision, represented as setting the value of a proposition d_n either to *True* or *False*. A chain c in τ is a totally ordered subset of τ . A subtree is a subset of τ satisfying postulates P.1-P.5 listed below.

We now describe the tree structure in more detail. The postulates are as follows:

P.1. Nontriviality and Finiteness: N is a non-empty finite set of nodes, n = 0, 1...N.

P.2. **Partial causal ordering**: The binary relation \prec is transitive and antisymmetric.

P.3. No backward branching: Incomparable nodes never have a common upper bound. For distinct nodes n_1, n_2 ,

$$(n_1 \prec n_3 \land n_2 \prec n_3) \Rightarrow (n_1 \prec n_2 \lor n_2 \prec n_1)$$

²Note that decisions over arbitrary finite sets of alternatives can be represented using binary nodes. As an illustration, Figure 2 shows Nature choosing between three alternatives (gamma ray emission, electron passage and electron absorption). This is represented using two binary nodes and one unary node.

Forward branching represents the openness or indeterminacy of the future and the absence of backward branching represents the determinacy of the past. By virtue of P.3, any n that has predecessors has exactly one immediate predecessor, which we will denote n^- . This asymmetry is the characteristic feature distinguishing tree structures from general partially ordered networks.

P.4. **Historical connection**: Every two nodes have a common lower bound

$$\forall n_1 \forall n_2 \exists \hat{n} | \hat{n} \preceq n_1 \land \hat{n} \preceq n_2.$$

By virtue of P.1 and P.4, there must exist a unique initial node, and it will be denoted by 0.

For any n, n' we say that n' is an immediate successor for n if $n \prec n'$ and there exists no $n'', n \prec n'' \prec n'$. Similarly, for any n, n' we say that n' is an immediate predecessor for n if $n' \prec n$ and there exists no n'', such that $n' \prec n'' \prec n$.

P.5 Binary: Each node has zero, one or two immediate successors.

A node with two or one immediate successors is called a decision node. For a decision node with two immediate successors n, the immediate successors are denoted n^R and n^L . A trivial decision node n has exactly one successor n^R . A terminal node has no successors. The existence of terminal nodes is implied by P.1, but terminal nodes will play no substantive role in the analysis to follow, since all relevant bounds will be derived from the bounded rationality of individuals. The indicator function $ind(\cdot)$ is defined on the set of non-terminal nodes and takes on the value 1 (respectively, 0) if it is the individual (respectively, Nature) who makes the choice at that node.

Thus a tree is characterized by the triple $\tau = (N, \prec, ind(\cdot))$.

3.2 Occurrences, histories, events and instants

An occurrence O is a collection or set of nodes $O \subseteq N$. A history H is a maximal chain in τ . For any n, the partial history for n, denoted H_n^- , is the maximal chain for which n is an upper bound. Let $|H_n^-|$ denote the height of node n (that is, the 'length' of the partial history H_n^-). We add the convention that the past history of the initial node 0 is \emptyset and hence its height is zero. An event E is a union of histories. $E_{(n)}$ is the union of the set of histories passing through n, that is, $H \subseteq E_{(n)}$ iff $n \in H$.

We assume all histories considered in the model can be traced to a common origin (this is essentially P.4) and that the individual has a finite time horizon. The terminal period can be taken to extend past this time horizon. So without significant loss of generality we can assume that the length of all histories in τ is T.³

Given this convention, we shall refer to the occurrence consisting of all nodes, for which the past history has some given length t as the *instant* t. That is, instants are 'horizontal slices' of the decision trees. An instantaneous occurrence O_t is a collection of nodes all occurring at the same instant t. For any event E, and instant t there is a unique instantaneous occurrence $O_t = E \cap t$.

The ordering \prec induces an ordering on instantaneous occurrences, also denoted \prec . Define $O_t \prec O'_{t'}$ if for each $n' \in O'_{t'}$ the predecessor of n' at time t, given by

$$n = H_{n-}^- \cap t$$

satisfies $n \in O_t$.

3.3 Syntax and semantics

For a given tree τ , the syntactic structure of the model begins with a set of propositions \mathbf{P}_{τ} , defined as sentences in a formal language. The language is built up from a set of primitive propositions \mathbf{B}_{τ} , the standard logical operators \vee, \wedge and \neg , and a tense logic operator w_{τ} derived from the tree structure τ . The elements of \mathbf{P}_{τ} are well-formed formulae built up from these elements: the syntax is standard and is not described in detail.

A crucial element of the model, reflecting the bounded rationality of individuals that is of central concern here is that, in general, not all sentences derivable in a formal language that is rich enough to describe the objective world will be expressible in \mathbf{P}_{τ} . The idea of expressibility, developed in more detail below, is that a proposition should be characterized by a set of nodes at which it is true. For $p \in \mathbf{P}_{\tau}$, truth is relativized to nodes in the given tree

 $^{^{3}\}mathrm{In}$ Figure 1 (respectively, 2) the length of all histories in the tree is two (respectively, three).

 τ : for $n \in N$, the statement 'p is true at n' is written $(\tau, n) \models p$. Some, but, in general, not all, of the propositions in \mathbf{P}_{τ} can be evaluated directly from the topology of the tree structure.

Of particular importance is the subset of decision propositions $\mathbf{D}_{\tau} \subseteq \mathbf{B}_{\tau}$, $\mathbf{D}_{\tau} = \{d_{0\tau}, d_{1\tau}, ..., d_{N\tau}\}$ indexed by the set of nodes, where for a node $n_{\tau}, d_{n\tau}$ is interpreted as 'The relevant agent (the individual or Nature) chose R at n_{τ} .'⁴ A central feature of the framework adopted here is that the truth value for decision propositions is derived directly from the tree structure. Proposition $d_{n\tau}$ is true at n', written $(\tau, n') \models d_{n\tau}$, only if $n^R \preceq n'$. Notice the convention (adopted in general) that a proposition dependent on a decision at node n is false at n', where $n \not\prec n'$. So, for example, in the tree in Figure 2, the propositions 'nature moves R at 2' and 'nature moves L at 2' are both false if the individual moves L at node 0, since $2 \not\prec n'$, for n = 1, 3 and 6.

In addition, we can derive, as in Belnap, Perloff and Xu (2001), the tense logic operator w_{τ} , where $w_{\tau}p$ is interpreted as 'it was true that p' directly from the tree structure τ , assuming that the truth value of the proposition p is known.

Definition 1 $(\tau, n) \models w_{\tau}p$ iff there exists $n' \prec n$ and $(\tau, n') \models p$.⁵

Proposition 1 For any *n* there exists p_n in the logical closure of \mathbf{D}_{τ} such that p_n is true precisely at *n*

Proof: Assume $n \neq 0$. Let n^- be the immediate predecessor of n, and suppose wlog that $n = (n^-)^R$. Then the proposition $p_n \equiv d_{n^-\tau} \wedge \neg (w_\tau d_{n^-\tau})$ that is 'It is true, and has not earlier been true that the decision maker at n^- chose R' is true precisely at n.

Now, for 0, let $p_0 = \neg (w_\tau p)$, where p is any valid formula.

That is, any node can be characterized by the unique partial history leading to that node, and this characterization has an expression in the logical

closure of \mathbf{D}_{τ} .

⁴The choice of R is arbitrary. The analysis could equally well be undertaken in terms of $\hat{d}_{n\tau}$, where for node n, $\hat{d}_{n\tau}$ is interpreted as 'the relevant agent chose L at n.' Note that it is not true that $\hat{d}_{n\tau} \equiv \neg d_{n\tau}$. If node n is never reached then both $d_{n\tau}$ and $\hat{d}_{n\tau}$ are false.

⁵Noting that the set of nodes is finite, this definition can be expressed without quantifiers, as $\neg (\bigvee_{n' \prec n} (\tau, n') \nvDash p)$.

Corollary 2 For any occurrence $O \subseteq N$ there exists p(O) in the logical closure of \mathbf{D}_{τ} such that p is true precisely on O.

Proof: The required proposition is just $\bigvee_{n \in O} p_n$.

In general, the set \mathbf{P}_{τ} will include many sentences that are not syntactically expressed in terms of decision propositions d_n , that is, are not in the logical closure of \mathbf{D}_{τ} . The semantic content of the model lies in the assignment of truth values to such propositions for given node n. In the example illustrated in Figure 2 we may consider the propositions p_{γ} 'gamma rays have been emitted' and p_e 'electron passage has taken place'. Proposition p_{γ} is true at nodes 5 and 9, and false at all other nodes. Proposition p_e is true at node 8, and false at all other nodes.

Formally, truth is determined by a valuation function $V_{\tau}(\bullet; n) : \mathbf{P}_{\tau} \to \{True, False\}$ and truth values satisfy the usual logical properties. As above, we have the associated notation $(\tau, n) \models p$ for $V_{\tau}(p; n) = True$. Conversely, we can define an interpretation function $\Im(\bullet; \tau)$ mapping propositions p into the occurrence (that is, the set of nodes) for which p is true.

Propositions, p, p' are semantically equivalent if for all nodes $n, V_{\tau}(p; n) = V_{\tau}(p'; n)$. That is, two propositions p, p' are semantically equivalent if $\Im(p; \tau) = \Im(p'; \tau)$, and so \Im partitions \mathbf{P}_{τ} into a set of equivalence classes, which may be denoted by \mathbf{P}_{τ}/\Im and referred to as the set of semantically distinct propositions.

As shown in Corollary 2, each occurrence O is characterized by a proposition p(O) in the logical closure of \mathbf{D}_{τ} , and so each equivalence class in \mathbf{P}_{τ}/\Im contains a canonical element of the logical closure of \mathbf{D}_{τ} .

Following Belnap, Perloff and Xu (2001), we also introduce the notion that a proposition $p \in \mathbf{P}_{\tau}$ is settled true at n if for all $n' \succeq n$, $V_{\tau}(p; n') = True$.

3.4 Knowledge and information

Fix a tree structure τ . We specify an information correspondence $\Pi_{\tau} : N \to 2^N$, which describes the set of nodes the individual considers possible at the instant in which node n arises, that is, her information set at n. Note that we are following conventions of modal logic of knowledge, where information

sets are defined for every possible state of the world. By contrast the usual convention of game theory is that information sets for an individual are defined only for decision nodes controlled by that individual, and not for 'chance' nodes controlled by Nature (or, in a multi-player game, by other players).

We require that $n \in \Pi_{\tau}(n)$, and also that if $n \in t$, $\Pi_{\tau}(n) \subseteq t$, that is, the individual knows what instant she is at.⁶ Unless otherwise noted, we assume that Π_{τ} satisfies the standard properties, namely:

I.1 Π_{τ} induces a partition of each instant t.

I.2 if $n' \in \Pi_{\tau}(n)$, then $(n')^{-} \in \Pi_{\tau}(n^{-})$.

I.3 For any instant t and any pair of nodes n, n' in t, if there exists a pair of nodes \hat{n} and \hat{n}' , and an instant \hat{t} , such that $(\hat{n}, \hat{n}') \subseteq \hat{t}, \hat{n} \preceq n, \hat{n}' \preceq n',$ $ind(\hat{n}^{-}) = 1, \hat{n} = (\hat{n}^{-})^{R}$ and $\hat{n}' = ((\hat{n}')^{-})^{L}$, then $\Pi^{i}_{\tau}(n) \cap \Pi^{i}_{\tau}(n') = \emptyset$.

I.2 says an individual never forgets any information she had in previous instants. I.3 entails that she knows her own moves. Notice that the hypothesis of I.3 may be interpreted as saying that the pair of nodes \hat{n} and \hat{n}' in instant \hat{t} arise from different choices by the individual in the previous instant (that is, \hat{n} follows from a choice of R by the individual at node \hat{n}^- , and \hat{n}' follows from a choice of L by the individual at node $(\hat{n}')^-$). I.3 then requires that this pair of nodes and any pair of their respective successor nodes cannot reside in the same information set. Together I.2 and I.3 imply the individual has perfect recall within the tree structure.

The information correspondence associated with the tree illustrated in Figure 2 may be taken to be: $\Pi_{\tau}(n) = \{n\}$ for $n \neq 4, 5$ and $\Pi_{\tau}(4) = \Pi_{\tau}(5) = \{4, 5\}$. That is, the researcher does not learn the result of the experiment (if undertaken) until t = 3.

We can now define the knowledge (modal) operator k_{τ} which will form the basis of an extension of the formal language available to the individual. The proposition $k_{\tau}p$ is stated as 'the individual knows p'.

Definition 2 For $p \in \mathbf{P}_{\tau}$, $(\tau, n) \models k_{\tau}p$ if for all n' such that $n' \in \Pi_{\tau}(n)$, $(\tau, n') \models p$.

⁶The latter assumption is not essential to the analysis, but simplifies the formulation of the lattice structure in subsequent sections, and also the characterisation of probability.

We define a set of propositions \mathbf{Q}_{τ} , available to the individual at τ , and containing the set of sentences \mathbf{P}_{τ} extended to incorporate the knowledge operator k_{τ} .

A useful derived operator (Hart, Heifetz and Samet 1996) is

the knowing whether operator $j_{\tau}p \equiv k_{\tau}p \lor k_{\tau}\neg p$;

Note that in the structure defined here, the set of nodes is fixed and finite, and knowledge operators are defined relative to that set of nodes. By contrast, a common approach is to begin with a set of propositions and derive a state space for which any two propositions that are not logically equivalent are semantically distinct. When combined with knowledge operators, this approach generates an uncountably large state space (Hart, Heifetz and Samet 1996).

We can derive

the considered operator $c_{\tau}p \equiv j_{\tau}p \lor k_{\tau}\neg j_{\tau}p$; and

the unconsidered operator $u_{\tau}p \equiv \neg c_{\tau}p$.

In Halpern and Rêgo's (2005) syntactic rendition of the semantic model of Heifetz, Meier and Schipper (2006), the operator c_{τ} is referred to as the awareness operator. In view of Proposition 3 below, and implications developed thereafter, we prefer the term 'considered'. We follow Halpern and Rêgo (2006) in proposing a broader notion of awareness, developed below.

As already observed, the state space in this model is finite, and the finite set \mathbf{P}_{τ} is rich enough to describe all distinct events. We similarly assume that \mathbf{Q}_{τ} is finite, but require that it be rich enough that for any $p \in \mathbf{P}_{\tau}$, the propositions $k_{\tau}p, j_{\tau}p, c_{\tau}p$ and $u_{\tau}p$ are all elements of \mathbf{Q}_{τ} . Thus, the individual 'knows what she knows' as regards propositions in \mathbf{P}_{τ} , but does not have access to propositions containing chains of modal operators of arbitrary length.⁷

Proposition 3 Under the stated conditions, if $p \in \mathbf{P}_{\tau}$, then, for all n, $(\tau, n) \models c_{\tau}p$.

⁷We can formalize this, and the extension of the knowledge operator to Q_{τ} by stating that for any, p such that $k_{\tau}p \in Q_{\tau}, (\tau, n) \models k_{\tau}p$ if for all n' such that $n' \in \Pi_{\tau}(n), (\tau, n') \models p$.

Proof: Recall that $\Im(p;\tau)$ is the set of nodes in τ at which p is true. If $\Pi_{\tau}(n) \subseteq \Im(p;\tau)$, then $k_{\tau}p$. If not, there exists $n^* \in \Pi_{\tau}(n), (\tau, n^*) \models \neg p$. By I.1, for any $n' \in \Pi_{\tau}(n), n^* \in \Pi_{\tau}(n') = \Pi_{\tau}(n)$. Hence, for any $n' \in \Pi_{\tau}(n), (\tau, n') \models \neg k_{\tau}p$. Hence, $(\tau, n) \models k_{\tau} \neg k_{\tau}p$.

That is, all propositions $p \in \mathbf{P}_{\tau}$ for which truth value can be assigned relative to the given tree τ are considered.

3.5 The existential quantifier

A critical feature of the model to be developed here is that the individual decisionmaker has the experience of discovering new possibilities, and individual may reasonably anticipate discovering new possibilities in the future. However, this anticipation cannot be expressed in \mathbf{P}_{τ} or even in the richer set \mathbf{Q}_{τ} augmented by the modal operators of knowledge. To capture it, we introduce an existence quantifier over propositions.

The existence quantifier \exists is used in conjunction with a formula for substitution, for example:

$$\exists p \in \mathbf{Q}_{\tau} : (p \Rightarrow p') \land \neg (p' \Rightarrow p).$$

That is, there is some (non-equivalent) proposition p that implies p'. We will write a generic existential proposition as

$$\exists p \in \hat{\mathbf{Q}}, \theta(p)$$

where $\hat{\mathbf{Q}}$ is the domain of quantification, and $\theta(p)$ is a syntactically valid sentence in which p appears as a predicate variable (Jeffrey 1990).

To refer to the existence of undiscovered possibilities we need an appropriate domain of quantification $\hat{\mathbf{Q}}$ encompassing propositions that may be considered in the future. We now turn to the task of constructing a framework in which such domains may be defined.

4 The Lattice structure

Proposition 3 shows that a useful notion of unawareness requires the existence of more than one tree structure for the world, such that propositions may be expressible relative to some trees but not relative to others. To represent this, we adopt an approach similar to that of Heifetz, Meier, and Schipper (2006), but with additional structure required to capture notions of learning and discovery, and to allow for higher-order propositions representing awareness of unconsidered propositions.

We consider a lattice \mathcal{T} of trees, each denoted τ (and each satisfying P.1–P.5) and an ordering \sqsubseteq , for which there exists a maximal tree τ_0 , that is $\tau \sqsubseteq \tau_0$ for all τ . Associated with τ_0 , are a valuation function V_0 : $\mathbf{P}_0 \to \{True, False\}$, a possibility correspondence $\Pi_0 : N_0 \to 2^{N_0}$, and an assignment giving for each n_0 in N_0 , a tree $\tau(n_0) \sqsubseteq \tau_0$ in \mathcal{T} . We define sets of propositions \mathbf{P}, \mathbf{Q} where \mathbf{P} is the closure of $\cup_{\tau \in \mathcal{T}} \mathbf{P}_{\tau}$ under logical operators (and therefore infinite) \mathbf{Q} is the closure of \mathbf{P} under the knowledge operator and \mathbf{R} is the logical closure of \mathbf{Q} under the existence quantifier and the awareness operator to be defined below.

Here the tree $\tau(n_0)$ is the subjective representation of the world that arises if the individual is at the objective node n_0 . We assume that \mathcal{T} is the image of N_0 under the given assignment. That is, for any τ' in \mathcal{T} , there exists $n'_0, \tau' = \tau(n'_0)$.

The possibility correspondence Π_0 represents the information sets that would apply in a standard decision tree in which the individual is fully aware of the structure of the external world as represented by τ_0 . The lattice structure captures the fact that, because of bounded rationality, not all propositions expressible in τ_0 are available to the individual at node n_0 and therefore the information set actually available to her must be represented in terms of a coarser information structure. The proposition sets **P**, **Q** and **R** incorporate propositions in successively richer languages, derived from a base set **B** which includes a subset **D**, the disjoint union of the decision proposition sets \mathbf{D}_{τ} for each τ .

4.1 Mappings between trees

If for any pair of trees τ and τ' , we have $\tau \sqsubseteq \tau'$, then there exists a surjective mapping $r_{\tau}^{\tau'}$: $\{0, \ldots, N_{\tau'}'\} \rightarrow \{0, \ldots, N_{\tau}\}$ exhibiting the property that for any pair of nodes n', n'' in $\{0, \ldots, N_{\tau'}'\}, r_{\tau}^{\tau'}(n') = r_{\tau}^{\tau'}(n'')$ implies $|H_{n'}^{-}| = |H_{n''}^{-}|$. That is, if two nodes in the finer tree are mapped to the same node in the coarser tree, then these nodes must come from the same instant in the finer tree. Further, we assume that, if $\tau \sqsubseteq \tau'$, then the associated mapping $r_{\tau}^{\tau'}(\cdot)$ is order-preserving, that is, for all $n', n'' \in \tau', n' \prec_{\tau'} n''$ implies $r_{\tau}^{\tau'}(n') \prec_{\tau} r_{\tau}^{\tau'}(n'')$ in τ .

Since both τ and τ' are trees each satisfying P.1–P.5, it readily follows that if two nodes in the finer tree are mapped to the same node in the coarser tree, then their immediate predecessors, if distinct, are also mapped by $r_{\tau}^{\tau'}(\cdot)$ to a common node in the coarser tree. Since τ_0 is maximal with respect to the ordering \sqsubseteq , it follows that all histories in all trees have the same length and the following relations hold between any two trees ordered by \sqsubseteq . That is, $r_{\tau}^{\tau'}$ induces a surjective mapping of occurrences in τ' onto occurrences in τ such that histories map to histories, events to events, instantaneous occurrences to instantaneous occurrences.

In particular, we can now distinguish the objective and subjective perspectives. Consider an individual who is at node n_0 in the maximal tree τ_0 . Associated with this node is a subjective tree $\tau(n_0)$ and a mapping $r_{\tau(n_0)}^0$. The image of n_0 under $r_{\tau(n_0)}^0$ may be denoted by $n_0^* = r_{\tau(n_0)}^0(n_0)$. This is the node in the subjective tree at which the individual is located, given the objective node n_0 .

Returning to our example of the scientific researcher in Section 2, let the tree in Figure 2 be τ_0 and consider the following tree $\hat{\tau}$, illustrated in Figure 3.

FIGURE 3 NEAR HERE

We have $\hat{\tau} \sqsubseteq \tau_0$ with an associated surjective mapping given by

$$\begin{aligned} r^{0}_{\hat{\tau}}(0) &= 0, \, r^{0}_{\hat{\tau}}(1) = i, \, r^{0}_{\hat{\tau}}(2) = ii, \, r^{0}_{\hat{\tau}}(3) = iii, \\ r^{0}_{\hat{\tau}}(4) &= r^{0}_{\hat{\tau}}(5) = iv, \, r^{0}_{\hat{\tau}}(6) = v, \, r^{0}_{\hat{\tau}}(7) = r^{0}_{\hat{\tau}}(9) = vi, \, r^{0}_{\hat{\tau}}(8) = vii. \end{aligned}$$

That is, the tree $\hat{\tau}$ is defined by identifying nodes 4 and 5 and nodes 7 and 9. An alternative coarsening could be obtained by identifying nodes 4 and 5 and nodes 8 and 9. As shown below, this choice will not make a difference to the expressibility of propositions in $\hat{\tau}$.

We take the lattice for the world of this stylized example as $\mathcal{T} = (\tau_0, \hat{\tau})$. The associated assignment of trees is $\tau(n_0) = \hat{\tau}$ if $n_0 \neq 9$, and $\tau(9) = \tau_0$. That is, the researcher only becomes aware of the possibility of gamma ray emission at t = 3 in the history where the experiment is undertaken and nature chooses to emit gamma rays. More generally, observe that, whenever $\tau \sqsubseteq \tau' \sqsubseteq \tau''$, for any n'' in τ''

$$r_{\tau}^{\tau'}\left(r_{\tau'}^{\tau''}\left(n''\right)\right) = r_{\tau}^{\tau''}\left(n''\right).$$
(1)

That is, the diagram in Figure 4 commutes.

FIGURE 4 NEAR HERE

We denote the inverse correspondence $(r_{\tau}^{\tau'})^{-1}$ by $\rho_{\tau}^{\tau'}$, and for any n in $N(\tau)$, refer to $\rho_{\tau}^{\tau'}(n)$ as the preimage of n in τ' . Similarly, for any occurrence $O \subseteq N(\tau)$, write

$$\rho_{\tau}^{\tau'}(O) = \bigcup_{n \in O} \rho_{\tau}^{\tau'}(n)$$

and refer to $\rho_{\tau}^{\tau'}(O)$ as the pre-image of O.

By an argument similar to that yielding 1, for any n' in τ'

$$\rho_{\tau''}^{\tau} \left(r_{\tau}^{\tau'} \left(n' \right) \right) = \rho_{\tau''}^{\tau'} \left(n' \right).$$
(2)

Suppose $\tau \sqsubseteq \tau'$. Then there exists a 1–1 mapping $\hat{\rho}_{\tau'}^{\tau} : \mathbf{P}_{\tau} / \mathfrak{T} \to \mathbf{P}_{\tau'} / \mathfrak{T}$ induced by $\rho_{\tau}^{\tau'}$. For given $p \in \mathbf{P}_{\tau}$, and $p' \in \mathbf{P}_{\tau'}$, we have that $p' \in \hat{\rho}_{\tau'}^{\tau}(p)$ if and only if $\mathfrak{T}(p'; \tau') = \rho_{\tau}^{\tau'}(\mathfrak{T}(p; \tau))$. Hence, a truth valuation function $V_{\tau'}$ induces a truth valuation function V_{τ} and therefore, for all τ , truth valuation functions may be derived from V_0 . For any $p \in \mathbf{P}_{\tau}$ this construction implies the existence of a corresponding equivalence class in \mathbf{P} . The canonical element of this class is referred to as the corresponding proposition in \mathbf{P} for p.

With this setup, we can formalize the notion that a proposition $p \in \mathbf{P}_{\tau}$ is expressible in τ' .

Definition 3 A proposition $p \in \mathbf{P}$ is expressible in τ if the corresponding occurrence in τ_0 is the pre-image of an occurrence in τ . A proposition $p \in \mathbf{P}_{\tau}$ is expressible in τ' if the corresponding proposition in \mathbf{P} is expressible in τ' .

If $\tau \sqsubseteq \tau'$, then any $p \in \mathbf{Q}_{\tau}$ is expressible in τ' . In particular, any $d_{n\tau} \in \mathbf{D}_{\tau}$ is expressible in τ' .

Consider the example illustrated in Figures 2 and 3. We observe that the proposition $p_{\gamma} \in \mathbf{P}$ corresponding to the occurrence $\{5,9\}$ in τ_0 is inexpressible in $\hat{\tau}$, since this is not the pre-image of any occurrence in $\hat{\tau}$. The proposition $p_e \in \mathbf{P}$ corresponding to the occurrence $\{8\}$ in τ_0 is expressible, since it is the pre-image of $\{vii\}$ in $\hat{\tau}$. Since p_e , defined previously in τ_0 , is expressible in $\hat{\tau}$, we will use the same label to refer to this element of $\mathbf{P}_{\hat{\tau}}$. With this convention, $\mathbf{P}_{\hat{\tau}} \subseteq \mathbf{P}$.

In general, a given sentence in natural language will not correspond, in a semantic sense, to the same proposition in different trees. For example, the sentence 'electron absorption is observed' corresponds to occurrence $\{7\}$ in τ_0 and to occurrence $\{vi\}$ in $\hat{\tau}$. However, $\{7\} \neq \rho_{\tau}^0(\{vi\})$. This mismatch is an inevitable consequence of the coarsening of the tree structure, which implies that, relative to the more refined language, distinctions are blurred, qualifying conditions are omitted and so on. As discussed below, the concept of ambiguity, in its ordinary language sense, must be understood in the light of the observation that sentences cannot, in general, have the same meaning for individuals (or a given individual at different nodes) with different tree-structure representations of the world.

For given $\tau', \tau \sqsubseteq \tau'$, and $\Pi_{\tau'}$, the mapping $r_{\tau}^{\tau'}$ induces a possibility correspondence $\Pi_{\tau} : N_{\tau} \to 2^{N_{\tau}}$ which makes the diagram in Figure 5 commute.

FIGURE 5 NEAR HERE

That is, let $\hat{r}_{\tau}^0: 2^{N_0} \to 2^{N_{\tau}}$ be the set mapping induced by r_{τ}^0 , giving for any $O \subseteq N_0$,

$$\hat{r}^{0}_{\tau}(O) = \left\{ r^{0}_{\tau}(n) : n \in O \right\}$$

Now for any $n \in \tau$, define

$$\Pi_{\tau}(n) = \hat{r}_{\tau}^{0} \left(\Pi_{0} \left(\rho_{\tau}^{0}(n) \right) \right)$$

In the construction of the world, we require that, for any node n_0 the possibility correspondence $\Pi_{\tau(n_o)}$ be that induced from Π_0 in this way. Now consider the corresponding diagram for τ, τ' such that $\tau \sqsubseteq \tau'$, shown in Figure 6.

FIGURE 6 NEAR HERE

We have:

Lemma 4 Given the construction above, the diagram in Figure 6 commutes.

Proof: Consider τ, τ' such that $\tau \sqsubseteq \tau'$. We want to show that for n' in τ' ,

$$\Pi_{\tau}\left(r_{\tau}^{\tau'}\left(n'\right)\right) = \hat{r}_{\tau}^{\tau'}\left(\Pi_{\tau'}\left(n'\right)\right).$$

That is,

$$\hat{r}_{\tau}^{0}\left(\Pi_{0}\left(\rho_{0}^{\tau}\left(r_{\tau}^{\tau'}\left(n'\right)\right)\right)\right) = \hat{r}_{\tau}^{\tau'}\left(\hat{r}_{\tau'}^{0}\left(\Pi_{0}\left(\rho_{0}^{\tau'}\left(n'\right)\right)\right)\right).$$

By applying 1 and 2 to the RHS and LHS respectively, both sides of this equation are equal to

$$\hat{r}_{\tau}^{0}\left(\Pi_{0}\left(\rho_{0}^{\tau'}\left(n'\right)\right)\right).$$

This set-up allows us to give content to the notion of unconsidered propositions developed above. First note that, given Lemma 4, knowledge operators for τ are expressible in τ' , where $\tau \sqsubseteq \tau'$. Take a proposition p', expressible in τ' , and node $n' \in \tau'$ with corresponding node $n = r_{\tau}^{\tau'}(n')$ in τ The occurrence $\rho_{\tau}^{\tau'}(\Pi_{\tau}(n))$ is the preimage of $\Pi_{\tau}(n)$ in τ' . Then $k_{\tau}p'$ is true at n'in τ' if $\rho_{\tau}^{\tau'}(\Pi_{\tau}(n)) \subseteq \Im(p'; \tau')$. If p' is not expressible in τ , $(\tau', n') \models \neg k_{\tau}p'$, and similarly, $(\tau', n') \models \neg k_{\tau}(\neg k_{\tau}p')$. Conversely, if $(\tau, n) \models \neg k_{\tau}p$, then also $(\tau', n') \models \neg k_{\tau}p$. Hence, we have:

Proposition 5 For an individual with tree τ , $u_{\tau}p$ if and only if p is not expressible in τ .

Proof: 'If' follows from the argument above. 'Only if' follows from Proposition 3.

4.2 Expressibility and knowledge of existential propositions

We now consider what kinds of existential propositions are expressible in $\mathbf{Q}_{\tau(n_0)}$. We have:

Proposition 6 Suppose $\hat{\mathbf{Q}} \subseteq \mathbf{Q}_{\tau(n_0)}$. Any existential proposition of the form $\exists p \in \hat{\mathbf{Q}}, \theta(p)$, where $\theta(p)$ is a well-formed formula incorporating p is contained in the logical closure of $\mathbf{Q}_{\tau(n_0)}$. Conversely, if $\hat{\mathbf{Q}} \not\subseteq \mathbf{Q}_{\tau(n_0)}$, then propositions of the form $\exists p \in \hat{\mathbf{Q}}, \theta(p)$ are inexpressible in $\tau(n_0)$.

Proof: Noting that $\mathbf{Q}_{\tau(n_0)}$ is finite, the condition $\hat{\mathbf{Q}} \subseteq \mathbf{Q}_{\tau(n_0)}$ implies that propositions of the form $\exists p \in \hat{\mathbf{Q}}, \theta(p)$ may be expressed in $\mathbf{Q}_{\tau(n_0)}$, without

use of the existence quantifier, as:

$$\bigvee_{p \in \hat{\mathbf{Q}} \subseteq \mathbf{Q}_{\tau(n_0)}} \theta\left(p\right),\tag{3}$$

so that for $\hat{\mathbf{Q}} \subseteq \mathbf{Q}_{\tau(n_0)}$, the existence quantifier \exists does not generate any propositions not already contained in the logical closure of $\mathbf{Q}_{\tau(n_0)}$. For the converse observe that, if $\hat{\mathbf{Q}} \not\subseteq \mathbf{Q}_{\tau(n_0)}$, $\hat{\mathbf{Q}}$ contains propositions inexpressible in $\mathbf{Q}_{\tau(n_0)}$ and it follows that $\exists p \in \hat{\mathbf{Q}}, \theta(p)$ is also inexpressible in $\tau(n_0)$.

Observe in particular that if $\hat{Q} = Q_{\tau(n'_0)}$ $n'_0 \leq n_0$, $\hat{Q} \subseteq Q_{\tau(n_0)}$. That is, since there is no 'forgetting' in this model, existential propositions regarding past sets $Q_{\tau(n'_0)}$ are always expressible without use of the existence quantifier. This will not be true, in general, in the case $n'_0 \geq n_0$ or when n_0 and n'_0 are on different histories and therefore unrelated by \succeq . Nevertheless, given past experience of discovery, it seems reasonable to suppose that the individual will be aware [in a sense to be made precise] of the possibility that there will exist a set of propositions $\mathbf{Q}_{\tau(n'_0)}$ that will be expressible if she reaches $n'_0 \succeq n_0$, (or, in the subjective viewpoint available at $n^*_0 = r^0_{\tau(n_0)}(n_0)$, the node $n' = r^0_{\tau(n_0)}(n'_0)$) and therefore that she is aware at n_0 of propositions of the form $\exists p \in \mathbf{Q}_{\tau(n'_n)}, \theta(p)$.

An important implication of Proposition 6 is that, in the finite languages considered here, no increase in expressive power is required to incorporate the existence quantifier \exists .

On the other hand, the following result shows that an individual can never know (in the modal-logical sense formalized above) that there exists an unconsidered proposition.

Proposition 7 For any $n_0 \in \tau_0$, and any $n'_0 \succeq n_0$,

$$(\tau_0, n_0) \vDash \neg k_{\tau(n_0)} \left(\exists p \in \mathbf{Q}_{\tau(n'_0)}, u_{\tau(n_0)} p \right).$$

Proof: Either $Q_{\tau(n'_0)} = Q_{\tau(n_0)}$ in which case every $p \in Q_{\tau(n'_0)}$ is expressible in $\tau(n_0)$, in which case by Proposition 3, $c_{\tau(n_0)}p$ for all p, so $\neg \exists p \in Q_{\tau(n'_0)}$, $u_{\tau(n_0)}p$, or, alternatively, $Q_{\tau(n_0)} \subset Q_{\tau(n'_0)}$, in which case $\exists p \in Q_{\mathbf{Q}_{\tau(n'_0)}}, u_{\tau(n_0)}p$ (considered as a finite disjunction in τ_0) is inexpressible in $\tau(n_0)$. Hence, by Proposition 5

$$(\tau_0, n_0) \vDash u_{\tau(n_0)} \left(\exists p \in \mathbf{Q}_{\tau(n'_0)}, u_{\tau(n_0)} p \right),$$

which implies

$$(\tau_0, n_0) \vDash \neg k_{\tau(n_0)} \left(\exists p \in \mathbf{Q}_{\tau(n'_0)}, u_{\tau(n_0)} p \right).$$

Proposition 7 is central to our analysis and to the concerns of Halpern and Rêgo (2005), who raise the question of whether, in the interactive unawareness framework of Heifetz, Meier and Schipper (2006), there exists an extension of the logic of awareness such that it would be possible to say, for example, that there exists a fact that agent 1 is unaware of but agent 1 knows that agent 2 is aware of it. Although interaction between agents is not considered here, we show that, given the interpretations of k, c and u adopted here agents cannot know that they may in the future consider currently unconsidered propositions. If the future self is considered as agent 2, and awareness is defined in terms of the operator c, the question raised by Halpern and Rêgo (2005) may be answered in the negative. Proposition 7 shows that the knowledge operator k which generates the set of propositions Q_t is not powerful enough to permit statements of the form 'I know that there exists some currently unconsidered proposition p'.

5 The awareness operator

We have argued that, while individuals cannot consider (in the formal sense derived above) propositions that are inexpressible in their subjective treestructure representation, they may nonetheless be aware of the possibility that they will in the future discover previously unconsidered propositions. To formalize this idea, we propose an awareness operator, constructed in a manner similar to that used for knowledge operators. The idea is to define a set of nodes, analogous to the information set used in the construction of the knowledge operator.

The motivation for developing an awareness operator is similar to that of Halpern and Rego (2006), but there are differences in the terminology and formalization. In particular, whereas the notion of awareness developed here is tied to the hierarchy of state spaces, Halpern and Rego allow differential awareness in an example with only a single state of nature.

5.1 Awareness of existential propositions

Fix a node n_0 in the objective tree and associate with the corresponding subjective node $n_0^* = r_{\tau(n_0)}^0(n_0)$ a set $\mathcal{A}(n_0^*) \subseteq N(\tau(n_0))$ such that

- (i) $n_0^* \in \mathcal{A}(n_0^*)$
- (ii) if $n' \in \mathcal{A}(n_0^*), \Pi_{\tau(n_0)}(n') \subseteq \mathcal{A}(n_0^*)$

Property (i) will imply that the individual can engage in quantified modal logical reasoning with regard to the propositions considered at n_0^* , using the existence quantifier \exists . Property (ii) will imply that $\mathcal{A}(n_0^*)$ cannot yield information that would allow a refinement of $\Pi_{\tau(n_0)}$ either at n_0^* or any $n' \succeq n_0^*$.

With each $n' \in A(n_0^*)$, we associate a set of admissible predicates $\Theta(n'; n_0)$. Admissible predicates $\theta(\hat{p})$ are those the individual can entertain at n_0 when quantified over the propositions expressible at n'. We assume that, for all $n' \in A(n_0^*), \theta(n'; n_0)$ includes all predicates of the form $\hat{p} \Rightarrow p$, for $p \in Q_{\tau(n_0)}$

Now we introduce the logical operator $a_{\tau(n_0)}(p)$ stated as 'is aware of p at n_0 ' and say that $a_{\tau(n_0)}(p)$ is true whenever p is equivalent to a proposition of the form

$$\exists \hat{p} \in \hat{\mathbf{Q}}, \, \theta\left(\hat{p}\right),$$

where $n' \in A(n_0^*)$, $\hat{Q} = \bigcup_{n'_0 \in \rho^0_{\tau(n_0)}(n')} Q_{\tau(n'_0)}$ and $\theta \in \Theta(n'; n_0)$ is an admissible predicate. That is, \hat{Q} is the set of propositions of which the individual may be aware at n'_0 for some n'_0 in the pre-image of n' in the objective tree τ_0 .

We first show that, if an individual considers a proposition, they are aware of that proposition.

Lemma 8 For each $n_0 \in N_0$, $p \in Q_{\tau(n_0)}$, $(\tau_0, n_0) \models c_{\tau(n_0)}(p) \Rightarrow (\tau_0, n_0) \models a_{\tau(n_0)}(p)$.

Proof: Set $\hat{\mathbf{Q}} := \mathbf{Q}_{\tau(n_0)}$, and consider the predicate $\theta(\hat{p}) = \hat{p} \Leftrightarrow p$.

Hence we define $\mathbf{R}_{\tau(n_0)}$ as an extension of $\mathbf{Q}_{\tau(n_0)}$ obtained by adding propositions of the form $a_{\tau(n_0)}(p)$ as described above. As before, we do not require that every possible sentence be included in $\mathbf{R}_{\tau(n_0)}$. As foreshadowed above, for all $n_0, \mathbf{R}_{\tau(n_0)} \subseteq \mathbf{R}$, where **R** is the closure of **Q** under the existence quantifier and the awareness operator. As an example, consider the propositions of the form

$$\exists p' \in \mathbf{Q}_{\tau(n'_0)} : u_{\tau(n_0)}p',$$

for $n'_0 \succeq n_0$. That is, the individual is aware of the possibility that there exist propositions which she will discover in the future, if she reaches node n'_0 .

Recall the example illustrated in Figures 2 and 3 and the position of the researcher at t = 0. The researcher can consider the domains $\mathbf{Q}_3^{\mathbf{L}}$ and $\mathbf{Q}_3^{\mathbf{R}}$, containing propositions of which she will be aware at t = 3 conditional on choosing \mathbf{L} (not to conduct the experiment) or \mathbf{R} (to conduct the experiment) at t = 0. From the external perspective the proposition $(\exists p' \in \mathbf{Q}_3^{\mathbf{L}} : u_{\tau(n_0)}p')$ is settled False at t = 0 while the truth value of $(\exists p' \in \mathbf{Q}_3^{\mathbf{R}} : u_{\tau(n_0)}p')$ depends on Nature's choice at node 2 in instant t = 1.

As another example, for any proposition p, such that $k_{\tau(n_0)} \left(\neg j_{\tau(n_0)} p\right)$, consider the proposition

$$\exists p' \in \mathbf{Q}_{\tau(n'_0)} : u_{\tau(n_0)} p' \land (p' \Rightarrow p) .$$

$$\tag{4}$$

That is, there is a currently unconsidered proposition in $\mathbf{Q}_{\tau(n'_0)}$ which, if true, would imply 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).

Of equal importance is awareness about future choice nodes. In standard models where all contingencies are considered, individuals currently at some node n_0 can, in general correctly anticipate their choices contingent on arriving at any given future choice node, with an associated information set. This is true whether preferences are dynamically consistent, as in the standard model of Bayesian subjective expected utility, and some non-EU models (Machina 1989), or merely behaviorally consistent, as in non-EU representations of dynamic choice such as that of Karni and Safra (1990).

Such self-knowledge is not assumed here for two reasons. First, some future choice nodes are not represented in the subjective tree $\tau(n_0)$. Individuals are generally supposed to be aware that they will be confronted with choices in the future that they have not yet considered. We can write this as

$$a_{\tau(n_0)} \left(\exists d \in \mathbf{Q}_{\tau(n'_0)} : u_{\tau(n_0)} d \wedge (ind(d) = 1) \right).$$
(5)

Alternatively, the individual may consider a future choice, but may be aware that there exists an unconsidered proposition which, if true, would lead them to choose one way, and if false, would lead them to choose the other way. For a considered decision proposition d, ind(d) = 1 this may be written as $a_{\tau(n_0)} (\exists p' \in \mathbf{Q}_{\tau(n'_0)} : u_{\tau(n_0)}p' \land (p' \Leftrightarrow d))$. With this setup, a naive decisionmaker is one for whom $A(n^*_0) = \prod_{\tau(n_0)} (n^*_0)$.

With this setup, a naive decisionmaker is one for whom $A(n_0^*) = \prod_{\tau(n_0)} (n_0^*)$ That is, at node n_0^* in the subjective tree $\tau(n_0)$ available to them at n_0 , they can apply modal logic to the set of propositions Q_{τ} , but they are not aware of the possibility that they may discover new propositions in the future.

6 Dynamics of learning and discovery

We are now in a position to describe the dynamics of learning and discovery, from the external perspective given by τ_0 . Consider a move from n_0^- to n_0 , arising from a decision $d_{n_0^-}$. Associated with this move is an assignment of the tree τ (n_0) and possibility correspondence $\Pi_{\tau(n_0)}$. We adopt the following assumption.

Assumption (Increasing refinement) For all non-initial n_0 , $\tau(n_0^-) \sqsubseteq \tau(n_0)$.

Thus, the individual's representation of the world, represented by the tree τ (n_0^-) at node n_0^- , and by the tree τ (n_0) at node n_0 , grows progressively finer. On the other hand, given the assumptions on the associated possibility correspondence $\Pi_{\tau(n_0)}$, the individual's knowledge about the world, relative to any fixed $\tau \sqsubseteq \tau (n_0^-)$ grows more accurate. Thus subjective uncertainty may either increase or decrease over time.

This distinction may be usefully expressed in terms of propositions.

Definition 4 For given n_0 , consider a proposition $p \in \mathbf{Q}$ that is settled true at $n.^8$ The individual learns p at n_0 , written $(\tau_0, n_0) \models L_{\tau(n_0)}p$, if

⁸Note that the requirement that p be settled true at n involves no real loss of generality. If p is not settled true, either $\neg p$ is settled true, or a proposition of the form $p \lor w^- p$ (or else $\neg p \lor w^- \neg p$) is settled true at n.

 $(\tau_0, n_0) \models k_{\tau(n_0)}p \text{ and } (\tau_0, n_0^-) \models \neg k_{\tau(n_0^-)}p$, that is, if she knows p at n_0 but not at n_0^- .

Definition 5 Consider a proposition $p \in \mathbf{Q}$. The individual discovers the possibility p at n_0 , written $(\tau_0, n_0) \models D_{\tau(n_0)}p$, if $(\tau_0, n_0) \models c_{\tau(n_0)}p$ and $(\tau_0, n_0^-) \models u_{\tau(n_0^-)}p$, that is, if she considers p at n_0 but not at n_0^- .

In the example set out above, if the researcher arrives at node vii in $\hat{\tau}$ (corresponding to node 8 in the external tree τ_0) she learns the proposition p_e 'electron passage has taken place'. If she arrives at node 9 in tree τ_0 , she both discovers and learns p_{γ} 'gamma ray emission has taken place'.

In the case of discovery, the new proposition was not expressible in the subjective tree τ (n_0^-) which is why the process must be described with respect to τ (n_0) . In the case of learning, the newly-learned proposition may be, but need not be expressible in the subjective tree τ (n_0^-) . In particular, in the case of fact-free learning, the newly-learned proposition was not expressible in the subjective tree τ (n_0^-) . Thus, fact free learning involves both discovery and learning. On the other hand, Bayesian learning involves no discovery, but does require the observation of information, so that $\Pi_{\tau(n_0)}$ induces a refinement of the set of histories relative to $\Pi_{\tau(n_0)}$.

Definition 6 Continuing discovery: We say that the world is characterized by continuing discovery at n_0 if for every history h passing through the information set $\Pi_0(n_0)$, there exists $n'_0 \succ \Pi_0(n_0)$, $n'_0 \in h$ and $p \notin \mathbf{Q}_{\tau(n_0)}$ such that $(\tau_0, n'_0) \models D_{\tau(n'_0)}p$

The example illustrated in Figures 2 and 3 exhibits this property at node 5 but not elsewhere. A realistic extension allowing for multiple experiments would do so at most nodes in subtrees where the researcher decides to carry out experiments.

6.1 Discussion

We first observe that, from the external perspective, the existence quantifier can be used to describe unconsidered propositions.

Proposition 9 For any $n_0 \in N_0$, $(\tau_0, n_0) \models (\exists p \in \mathbf{Q} : u_{\tau(n_0)}p)$ if and only if $\tau(n_0) \neq \tau_0$.

Proof: If $\tau(n_0) = \tau_0$, then, for all $p \in \mathbf{Q}$, $(\tau_0, n_0) \models c_{\tau(n_0)}p$ by Proposition 3. If $\tau(n_0) \neq \tau_0$, then there exist at least two nodes $n'_0, n''_0 \in N_0$ such that $r_0^{\tau(n_0)}(n'_0) = r_0^{\tau(n_0)}(n''_0)$. Hence, the proposition $d_{n'_0}$ is inexpressible in $\tau(n_0)$ and, by Proposition 3 $(\tau_0, n_0) \models u_{\tau(n_0)}d_{n'_0}$, so that $(\tau_0, n_0) \models (\exists p \in \mathbf{Q} : u_{\tau(n_0)}p)$ as required.

This proposition gives an exact characterization from the external viewpoint, but not one that is very useful from a subjective viewpoint, since it is not clear what it would mean for an individual to recognize that her tree $\tau(n_0)$ was or was not equal to τ_0 , and it is not true in general that the individual will achieve the full awareness implied by $\mathbf{Q}_{\tau(n_0)} = \mathbf{Q}_{\tau_0}$ for some n_0 .

A more useful characterization may be derived from the concept of continuing discovery:

Proposition 10 Fix a node n_0 . If the world is characterized by continuing discovery at n_0 then, for any $\hat{n} \succeq n_0$, $\exists n'_0 \succeq \hat{n}$, such that

$$(\tau_0, n_0) \vDash \left(\exists p \in \mathbf{Q}_{\tau(n'_0)} : u_{\tau(n_0)} p \right).$$

Proof: Consider a history passing through $\Pi_0(n_0)$ and \hat{n} . By the definition of continuing discovery, each such history must contain a node \hat{n}'_0 such that $\left(\exists p \in \mathbf{Q}_{\tau(\hat{n}'_0)} : u_{\tau(n_0)}p\right)$. If $\hat{n}'_0 \succeq \hat{n}$, set $n'_0 = \hat{n}'_0$. If not, then since no propositions are forgotten, we can set $n'_0 = \hat{n}$.

Again, there is no way for the individual to determine, reasoning with the set of propositions in $\mathbf{Q}_{\tau(n_0)}$ that the world is characterized by continuing discovery. But an individual willing to rely on philosophical induction may reasonably conclude that, since the world has been characterized by continuing discovery in the past, it will continue to be so in the future.

Observe that an unboundedly rational agent, with access to the external tree τ_0 , does not gain any benefit from awareness of propositions involving future discovery. For such an agent, all affirmative propositions of this kind are false: the unboundedly rational agent never makes discoveries.

Part II: Beliefs, probabilities and choice

In Part I, we used a modal-logical definition of knowledge as the basis for our characterization of considered propositions and awareness. For many purposes in economics, a quantified concept of belief, such as subjective probability, is more valuable. The purpose of this part of the paper is to develop such a concept and relate it to the lattice structure developed above. Note that, as is common in the literature, knowledge in our framework is a stronger concept than 'belief with probability 1'. Whereas our conception of knowledge is partitional, and therefore excludes 'false knowledge', our representation of beliefs allows for 'impossible beliefs'.

In Section 7 we introduce probability concepts and argue that subjective probabilities may be induced from the probabilities that would be held by an unboundedly rational observer, conditional on implicit beliefs about unconsidered propositions. We define the special case of restricted Bayesianism, which corresponds to the observer's unconditional marginal probabilities on events considered by the agent. Next, we show how the model allows for impossible beliefs, modelled as the falsification of implicit beliefs, leading to the occurrence of events with zero subjective probability.

In Section 8 we consider, given the existence of unconsidered propositions, under what conditions a standard Bayesian learning approach can be applied to a subset of known propositions. We derive sufficient conditions for a framework in which prior subjective probabilities are induced by an 'underlying' probability distribution on the objectively given state space. In essence, these conditions amount to a requirement that the restricted domain in which Bayesian learning is applied should be orthogonal to currently unconsidered propositions that may be discovered during the period of Bayesian learning. This point is of particular relevance in considering the contribution of Aragones et al. (2005). We show that in the regression context they consider, the conditions we derive amount to requirements for orthogonality restrictions separating the block of variables under consideration from the rest of the data set. Under appropriate orthogonality conditions, it is possible to show that the regression problem posed by Aragones et al is computationally tractable.

In section 9 we show that with our lattice structure of different-tree representations of the world, ambiguity in a semantic sense naturally arises in the agent's subjective view of the world. Furthermore, a multiple prior quantification of the uncertainty embodied in the agent's subjective view of the world can be obtained by conditioning on the possible truth values for one or more implicit beliefs

In section 10 we discuss the way in which the agent's learning about previously-considered propositions leads endogenously to the discovery of new propositions. In particular, we claim that new discovery must occur if the agent observes events believed 'impossible' (that is, assigned zero prior probability) as a result of the falsification of implicit beliefs.

Finally, in section 11 we take some preliminary steps towards the development of a theory of choice for boundedly rational agents aware of the incompleteness of their state space. We argue that non-EU models of choice under uncertainty may be interpreted as heuristic adaptions of EU preferences to the context of an incomplete state space. Some possible behavioral implications are discussed.

7 Probabilities and beliefs

Up to this point we have eschewed any consideration of the individual (at n_0) having quantifiable beliefs over the nodes she perceives as possible for a given tree τ (n_0) that represents her current perception of the world. In this section we shall introduce such beliefs. The central idea is to posit that quantified beliefs can be represented by probabilities induced from an underlying probability distribution on the objectively given state space, possibly conditional on implicit beliefs about unconsidered propositions. We consider conditions under which the standard Bayesian learning approach can be applied to the subset of known propositions.

7.1 The structure of probabilistic beliefs

More formally, for a given objective node n_0 and the associated subjective tree $\tau = \tau (n_0)$, we shall consider the case where for each node n in the tree τ , given the unconditional likelihood of reaching that node, we can quantify by a well-defined probability the likelihood the move from that node will be R or L. That is, associated with the tree τ is a probability assignment $\pi_{\tau} : N \to [0, 1]$ with the properties

$$\pi_{\tau} (0) = 1$$

$$\pi_{\tau} (n) = \pi_{\tau} (n^{R}) + \pi_{\tau} (n^{L}).$$

Also, we have a conditional probability assignment on nodes. For any n, such that $\pi(n) > 0, \pi$ induces an assignment $\pi(\bullet; n)$ such that, for $n' \succ n^9$

$$\pi\left(n'|n\right) = \frac{\pi\left(n'\right)}{\pi\left(n\right)}.$$

This assignment can be extended to yield probability distributions for either horizontal slices (instantaneous occurrences) or vertical slices (histories and events) in the tree.

For any t, the assignment π_{τ} induces a probability distribution for instantaneous occurrences O_t , given by

$$\pi_{\tau}\left(O_{t}\right) = \sum_{n \in O_{t}} \pi_{\tau}\left(n\right)$$

with the usual convention that $\pi_{\tau}(\emptyset) = 0$.

For any partition of an instant t into occurrences $O_{1t}, ..., O_{Kt}$, we have

$$\sum_{k=1}^{K} \pi_{\tau} \left(O_{kt} \right) = 1.$$

For O'_t such that $\pi_{\tau t}(O'_t) > 0$, we obtain, for any O_t , the usual conditional probability:

$$\pi_{\tau}\left(O_{t}|O_{t}'\right) = \frac{\pi_{\tau}\left(O_{t}\cap O_{t}'\right)}{\pi_{\tau}\left(O_{t}'\right)}.$$

Now for any history H, with terminal node n(H,T) let

$$\tilde{\pi}_{\tau}(H) = \pi_{\tau}(n(H,T)).$$

Then

$$\sum_{H} \tilde{\pi}_{\tau} \left(H \right) = 1$$

For any event E let

$$\tilde{\pi}_{\tau}\left(E\right) = \sum_{H \in \mathbf{H}(E)} \tilde{\pi}_{\tau}\left(H\right)$$

so that $\tilde{\pi}_{\tau}$ is a probability distribution over events in the tree τ , given by the total probability of histories in which a node in the event occurs at some time

⁹In the cases $n' \not\succ n$, the conditional probability assignment is trivial.

For $n' \leq n$, $\pi(n'|n) \equiv 1$.

For $n', n' \leq n \& n \leq n', \pi(n'|n) \equiv 0$.

t, and with a slight abuse of notation, we may write $\tilde{\pi}_{\tau}(E;n)$ for $\tilde{\pi}_{\tau}(E|\{n\})$. Again, we adopt the usual convention that $\tilde{\pi}_{\tau}(\emptyset) = 0$.

Finally we may define a probability distribution, with associated conditional probabilities, for general occurrences as the sum of probabilities for histories intersecting the occurrence, that is, the probability of the event 'the occurrence takes place for some t'. These need to be defined carefully to allow for the fact that different occurrences may occur at different time periods. Such occurrences may be disjoint, considered as sets of nodes, even though the histories on which they occur overlap; in this case we want the conditional probability to be non-zero. We therefore define:

$$\tilde{\pi}_{\tau}\left(O\right) = \sum_{H \in \mathbf{H}(O)} \tilde{\pi}_{\tau}\left(H\right) \tag{6}$$

$$\tilde{\pi}_{\tau}\left(O|O'\right) = \frac{\sum_{H \in \mathbf{H}(O) \cap \mathbf{H}(O')} \tilde{\pi}_{\tau}\left(H\right)}{\sum_{H \in \mathbf{H}(O')} \tilde{\pi}_{\tau}\left(O'\right)}.$$
(7)

We can also define absolute and conditional probabilities for propositions $p \in \mathbf{Q}_{\tau}$, writing:

$$\widetilde{\pi}_{\tau}(p) = \widetilde{\pi}_{\tau}(\Im(p;\tau))$$

$$\widetilde{\pi}_{\tau}(p;n) = \widetilde{\pi}_{\tau}(\Im(p;\tau);n)$$

7.2 Derivation of probabilistic beliefs with bounded rationality

Now consider the case of unconsidered propositions. We have a lattice of trees, each of which may be associated with a probability assignment. The problem now is to define a probability π_{τ} in such a way that, for $\tau \sqsubseteq \tau'$, the triangular diagram involving π_0, π_{τ} and $\pi_{\tau'}$ commutes. The approach proposed here is to begin with a set of beliefs $\pi_0(n_0)$ from the external perspective, contingent on the available information. These are interpreted as the probability beliefs the individual would hold in the absence of the constraints associated with bounded rationality.

Next, we introduce the idea of an *implicit belief*. Let $p \in P$ be an unconsidered proposition, which may in general be written as a compound proposition of the form $(p_1 \vee p_2 ... \vee p_K) \vee (\neg p_{1'} \vee \neg p_{2'} ... \vee \neg p_{K'})$. That is, p asserts the truth of $p_1, p_2...p_K$ and the falsity of $p_{1'} \vee p_{2'}... \vee p_{K'}$ where the p_i and $p_{i'}$ are unconsidered propositions. An assignment of probability 1 to an unconsidered proposition p (simple or compound) will be referred to as an implicit belief.

Now consider the conditional probability distribution $(\pi_0(n_0)|p)$, and the associated marginal distribution over τ given by

$$\pi_{\tau}(n|p) = \sum_{n_0 \in \rho_{\tau}^0(n)} (\pi_0(n_0)|p).$$

That is, the beliefs of the individual with subjective tree τ are those that the unboundedly rational individual would hold given the information contained in p, and restricted to the events and propositions expressible in τ .

We have:

Lemma 11 If $\tau \sqsubseteq \tau'$, and p is a proposition unconsidered in τ' (and hence also in τ)

$$\pi_{\tau}(n|p) = \sum_{n' \in \rho_{\tau}^{\tau'}(n)} \pi_{\tau'}(n'|p) \,.$$

Proof:

$$\sum_{n' \in \rho_{\tau}^{\tau'}(n)} \pi_{\tau'}(n'|p) = \sum_{n' \in \rho_{\tau}^{\tau'}(n)} \sum_{n_0 \in \rho_{\tau}^0(n)} \pi_0(n_0|p) \,.$$

So, by equation (2), characterizing the commutativity properties of ρ , the result holds.

Corollary 12 : Consistency similarly applies for the probability distributions over events and propositions $\tilde{\pi}$.

7.3 Restricted Bayesianism

In the absence of any implicit beliefs, we begin with π_0 defined on τ_0 and define, for any n:

$$\pi_{\tau}\left(n\right) = \sum_{n_{0} \in \rho_{\tau}^{0}(n)} \pi_{0}\left(n_{0}\right)$$

We will call this choice for π_{τ} restricted Bayesian. It is the subjective probability distribution over events expressible in τ that would be held by a decision-maker with access to the maximal tree τ_0 , prior probability distribution π and information set $\Pi(n_0)$.

In general, restricted Bayesian beliefs will be reasonable within 'small worlds' in the sense described by Savage (1954). That is, a boundedly rational Bayesian will define particular subproblems for which she judges that a welldefined prior over relevant states (the projections of events in the larger world) is available, and will then apply Bayesian decision theory to these subproblems. One of the concerns of this paper is to specify conditions under which such a procedure may be applied in a consistent fashion.

7.4 Impossible beliefs

We next consider what it means for an individual to have 'impossible' beliefs. Consider implicit beliefs that assign probability zero to an unconsidered proposition p for which $\Pr(p \mid \pi_0(n_0)) > 0$. This means that at some node $n' \succ n_0$, p will be true. At or before this point, the individual's implicit beliefs must be revised if coherent subjective probabilities are to be derived conditional on those beliefs.

An obvious question is whether implicit impossible beliefs as we have defined them lead to explicit impossible beliefs, that is, assigning probability zero to some considered proposition $p \in P_{\tau}$ for which $\Pr(p \mid \pi_0(n_0)) > 0$. As a counterexample, consider the case where the individual's subjective tree is trivial, consisting of a single history passing through the node at which they are located. The associated proposition set is similarly trivial.

On the other hand, as long as the individual's subjective representation of the world is reasonably rich, and the individual's information sets are non-trivial, impossible implicit beliefs will be reflected in the occurrence of events explicitly considered as having probability zero. Whenever a considered proposition p corresponds to an event that is a subset of the truth set for the unconsidered proposition p' and the implicit beliefs assign probability zero to p', the induced explicit beliefs must assign probability zero to p.

Example 1 In the research example, the researcher implicitly assigns zero probability to gamma ray emission. Obviously, these probabilities must be revised if gamma ray emission is observed.

8 Bayesian updating

Bayesian learning procedures are not, in general, reliable for boundedly rational agents, since the possibility of discovery cannot be excluded. Given the power of Bayesian inference in many practical applications, however, it is obviously of interest to consider special cases where Bayesian learning is reliable.

8.1 Bayesian learning with unbounded rationality

We first show that the definitions above yield the standard characterization of learning. That is, in the unbounded rationality case with imperfect information and learning, Bayesian updating works in the usual fashion and so the law of iterated expectations holds.

In the standard unbounded rationality case with tree τ , the knowledge of a partially-informed individual at time t, given the occurrence of node nis summarized by an instantaneous occurrence $\Pi_{\tau}(n)$. Hence, for any O_t the probability for an individual with prior assignment π_{τ} and information $\Pi_{\tau}(n)$ is

$$\tilde{\pi}_{\tau}\left(O_{t};n\right) = \pi_{\tau t}\left(O_{t}|\Pi_{\tau}\left(n\right)\right)$$

and this definition can be extended to general occurrences as in equation (7).

Now suppose node $n' \succ n$ is realized at t' > t, so the information set is $\Pi_{\tau}(n')$. For any instantaneous event $O'_{t'}$, we can now derive $\pi_{\tau t'}(O'_{t'};n')$ in two ways. First, we can repeat the definition above, yielding

$$\tilde{\pi}_{\tau}\left(O_{t'}';n'\right) = \pi_{\tau t}\left(O_{t'}'|\Pi_{\tau}\left(n'\right)\right).$$

Alternatively, we can apply Bayesian updating to $\tilde{\pi}_{\tau}(O'_{t'}; n)$, and define

$$\tilde{\pi}_{\tau}\left(O_{t'}';n,n'\right) = \frac{\tilde{\pi}_{\tau}\left(O_{t'}';n\right)}{\tilde{\pi}_{\tau}\left(\Pi_{\tau}\left(n'\right);n\right)}.$$

Say that Bayesian updating is consistent if for all $n, n' \succeq n, O'_{t'}$,

$$\tilde{\pi}_{\tau}\left(O_{t'}';n,n'\right) = \tilde{\pi}_{\tau}\left(O_{t'}';n'\right)$$

Now we observe that, since $n' \succeq n$:

$$\tilde{\pi}_{\tau}\left(\Pi_{\tau}\left(n'\right);n\right) = \frac{\tilde{\pi}_{\tau}\left(\Pi_{\tau}\left(n'\right)\right)}{\tilde{\pi}_{\tau}\left(\Pi_{\tau}\left(n\right)\right)}$$

Also, by construction

$$\tilde{\pi}_{\tau}\left(O_{t'}';n\right) = \frac{\tilde{\pi}_{\tau}\left(O_{t'}'\right)}{\tilde{\pi}_{\tau}\left(\Pi_{\tau}\left(n\right)\right)},$$

and hence,

$$\begin{aligned} \tilde{\pi}_{\tau} \left(O_{t'}'; n, n' \right) &= \left(\frac{\tilde{\pi}_{\tau} \left(O_{t'}' \right)}{\tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n \right) \right)} \right) \middle/ \left(\frac{\tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n' \right) \right)}{\tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n \right) \right)} \right) \\ &= \frac{\tilde{\pi}_{\tau} \left(O_{t'}' \right)}{\tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n' \right) \right)} \\ &= \tilde{\pi}_{\tau} \left(O_{t'}'; n' \right). \end{aligned}$$

Hence, we have, as expected:

Lemma 13 In the case of unbounded rationality with incomplete information, Bayesian updating is consistent.

8.2 Unconsidered propositions and consistency

We next consider the more difficult question. Suppose we consider two trees $\tau \sqsubseteq \tau'$, nodes n in τ and n' in τ' , both occurring at instant t, and such that $n = r_{\tau}^{\tau'}(n')$, and information sets $\Pi_{\tau}(n)$ and $\Pi_{\tau'}(n')$. Under what conditions are the resulting conditional probabilities for τ consistent with those for τ' , that is, for any instantaneous O_t

$$\tilde{\pi}_{\tau}\left(O_{t}|\Pi_{\tau}\left(n\right)\right) = \tilde{\pi}_{\tau'}\left(\rho_{\tau}^{\tau'}\left(O_{t}\right)|\Pi_{\tau'}\left(n'\right)\right).$$

Our first sufficient condition deals with the absence of learning in the move from $\langle \tau, \Pi_{\tau}(n) \rangle$ to $\langle \tau', \Pi_{\tau'}(n') \rangle$. Suppose that, in the move from $\langle \tau, \Pi_{\tau}(n) \rangle$ to $\langle \tau', \Pi_{\tau'}(n') \rangle$ no new information is acquired. There is only the discovery of new propositions and the associated priors, derived ultimately from π_0 . That is, although the individual discovers new propositions inexpressible in τ , she does not learn the truth value of any of these propositions.

Proposition 14 Let n be a node in τ at instant t. If $\Pi_{\tau}(n) = r_{\tau}^{\tau'}(\Pi_{\tau'}(n'))$, then conditional probabilities are consistent.

Proof: Under these conditions

$$\begin{aligned} \tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n \right) \right) &= \sum_{n_{\tau} \in \Pi_{\tau}(n)} \pi_{\tau} \left(n_{\tau} \right) = \sum_{n_{\tau} \in \Pi_{\tau}(n)} \sum_{n' \in \rho_{\tau}^{\tau'}(n_{\tau})} \pi_{\tau'} \left(n' \right) \\ &= \sum_{n_{\tau'} \in \Pi_{\tau'}(n)} \pi_{\tau'} \left(n_{\tau'} \right) = \tilde{\pi}_{\tau} \left(\Pi_{\tau} \left(n' \right) \right) \end{aligned}$$

Since, for any O_t ,

$$\pi_{\tau t}\left(O_{t}\right) = \pi_{\tau' t}\left(\rho_{\tau}^{\tau'}\left(O_{t}\right)\right)$$

and this is also true for $O_t \cap \Pi_{\tau}(n)$, $\rho_{\tau}^{\tau'}(O_t) \cap \Pi_{\tau'}(n')$, the result holds.

Our next sufficient condition is based on (probabilistic) independence. If the only information that is learned in the move from $\langle \tau, \Pi_{\tau}(n) \rangle$ to $\langle \tau', \Pi_{\tau'}(n') \rangle$ refers to occurrences that are independent, with respect to $\pi_{\tau'}$ and the derived probability distributions, of any occurrence expressible in τ , then this information does not affect the induced probabilities for such events. More precisely, we say that O'_t (an instantaneous occurrence in τ') is independent of instantaneous occurrences in τ if for any O_t in τ

$$\pi_{\tau't}\left(\rho_{\tau}^{\tau'}\left(O_{t}\right)\cap O_{t}'\right) = \pi_{\tau't}\left(\rho_{\tau}^{\tau'}\left(O_{t}\right)\right)\pi_{\tau't}\left(O_{t}'\right).$$
(8)

Proposition 15 Let n be a node in τ at instant t. If

$$\Pi_{\tau'}(n') = \rho_{\tau}^{\tau'}(\Pi_{\tau}(n)) \cap O_t'$$
(9)

where O'_t in τ' is independent of instantaneous occurrences in τ , then conditional probabilities are consistent.

Proof: Under the stated conditions, for any O_t in τ , $\Pi_{\tau}(n) \cap O_t$ is an instantaneous occurrences in τ , as is $\Pi_{\tau}(n)$, so:

$$\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right) \cap O_{t}^{\prime} \right) = \pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right) \right) \pi_{\tau't} \left(O_{t}^{\prime} \right)$$
$$\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \right) \cap O_{t}^{\prime} \right) = \pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \right) \right) \pi_{\tau't} \left(O_{t}^{\prime} \right)$$

by independence. Now:

$$\begin{aligned} \tilde{\pi}_{\tau'} \left(\rho_{\tau}^{\tau'} \left(O_{t} \right) | \Pi_{\tau'} \left(n' \right) \right) \\ &= \frac{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(O_{t} \right) \cap \Pi_{\tau'} \left(n' \right) \right)}{\pi_{\tau't} \left(\Pi_{\tau'} \left(n' \right) \right)} \\ &= \frac{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(O_{t} \right) \cap \rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \right) \cap O_{t}' \right)}{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \right) \cap O_{t}' \right)} \text{ (by (9))} \\ &= \frac{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right) \right) \pi_{\tau't} \left(O_{t}' \right)}{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right) \right)} \\ &= \frac{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right) \right)}{\pi_{\tau't} \left(\rho_{\tau}^{\tau'} \left(\Pi_{\tau} \left(n \right) \right) \right)} \\ &= \frac{\pi_{\tau t} \left(\Pi_{\tau} \left(n \right) \cap O_{t} \right)}{\pi_{\tau t} \left(\Pi_{\tau} \left(n \right) \right)} \text{ (by results above),} \end{aligned}$$

as required.

8.3 Restricted Bayesianism and updating

A particularly important application of Propositions 14 and 15 arises in the context of learning and discovery. Let $n_0, n'_0 \in \tau_0$ be such that $n_0 \prec n'_0$ and let $\tau = \tau (n_0), \tau' = \tau (n'_0)$. Assume continuing discovery, so $\tau \sqsubseteq \tau'$. Consider an individual who moves from n_0 to n_0 replacing tree τ by τ' and possibility set $\Pi_{\tau} (n_0)$ by $\Pi_{\tau'} (n'_0)$. Under what conditions is restricted Bayesian updating consistent? That is, under what conditions can the individual simply update probabilities for occurrences O expressible in τ using the new information in $\Pi_{\tau'} (n'_0)$ expressed with respect to the original tree τ , and obtain the same subjective probabilities as would be yielded by deriving subjective probabilities relative to the new tree τ' with the associated possibility correspondence, so that

$$\tilde{\pi}_{\tau} \left(O | \Pi_{\tau} \left(n'_{0} \right) \right) = \tilde{\pi}_{\tau'} \left(\rho_{\tau}^{\tau'} \left(O_{t} \right) | \Pi_{\tau'} \left(n'_{0} \right) \right).$$
(10)

The results above give two cases where this is true. The first is where no learning takes place between n_0 and n'_0 , so that no updating is required. The second, more interesting case is where any new information expressible in τ' is independent of any event in τ that is, any newly discovered proposition expressible in τ' .

More generally, if equation (10) holds for all n'_0 such that $n_0 \prec n'_0$, we may say that Bayesian updating is consistent at n_0 .

8.4 The Aragones et al. example

The work of Aragones et al. shows that, in general, the problem of determining the optimal statistical model relating a set of m explanatory variables to a dependent variable is NP-hard. The analysis of restricted Bayesianism developed here may be used to characterize special cases when the problem is relatively easy. These special cases rest on orthogonality conditions of one kind or another, where the independence requirements for consistent Bayesian updating are satisfied.

The simplest case is where the explanatory variables are known to be mutually orthogonal, for example, because they arise from an experimental design. Then the optimal regression can be obtained by calculating and ranking covariances with the dependent variable, then running a stepwise regression. More generally, we may consider cases where the data set may be partitioned into blocks which are mutually orthogonal. In this case, standard Bayesian reasoning may be applied within blocks, provided they are sufficiently small to be tractable.

The analysis presented above shows that the orthogonality restrictions required to make the regression problem tractable are an instance of more general orthogonality conditions required to permit consistent use of Bayesian updating in the presence of unconsidered propositions.

9 Ambiguity and multiple priors

Discussion of 'ambiguity' in decision theory is most commonly associated with a state-space approach where there are no unconsidered proposition. The central idea is that an event (a measurable subset of the state space) is ambiguous if it is not associated with a well-defined probability number. Given a semantic interpretation of the state space (that is, a mapping between events and propositions such that each measurable event is the truth set for some proposition), this approach naturally yields a probabilistic concept of ambiguity, in which it makes sense to say that 'a proposition is ambiguous if its probability is unknown'.

But, in ordinary language, ambiguity is a semantic notion, and its meaning is different from (though intuitively related to) that defined above. An ambiguous statement is one for which the meaning is ill-defined. The framework developed above allows us to say that a proposition p, expressible in τ_0 is semantically ambiguous in τ if it is not expressible in τ .

The analysis of restricted Bayesianism suggests a probabilistic characterization of ambiguity, most simply expressed in terms of propositions. Suppose we are given a prior distribution π_0 , with induced distribution over propositions, evaluated at n_0 given by $\tilde{\pi}_0(\bullet; n_0)$. Call an uncertain proposition p, (that is, one such that $\neg j_0 p$), probabilistically ambiguous at n_0 for given $\tau(n_0)$, if there exist p' such that

 $(i)(\tau(n_0), n_0) \models u_{\tau(n_0)}p'$, that is, p' is unconsidered at n_0 ;

(ii) $\tilde{\pi}_0 (p \wedge p'; n_0) \neq \tilde{\pi}_0 (p; n_0) \tilde{\pi}_0 (p'; n_0)$, that is, p and p' are not independent.

Thus, there exists a proposition p' not considered at n_0 but which, if known to be true or false, would change the probability of p.

Lemma 16 Semantic ambiguity implies probabilistic ambiguity, but not vice versa.

Proof: For the implication, let p' = p. For the second, consider the case where there are four states of nature, generated by two propositions p and p', and only p is expressible in $\tau(n_0)$. As long as condition (ii) holds, p is probabilistically ambiguous but not semantically ambiguous.

Observe that the definitions of semantic and probabilistic ambiguity are based on the external viewpoint given by τ_0 and not on the subjective viewpoint, given by $\tau(n_0)$.

A couple of examples may be useful. Consider the case of an individual faced with the choice between investing in stocks or in bonds. Both investments involve some risk, since the real return on bonds will depend on the rate of inflation, and the rate of return on stocks depends on a range of systematic and idiosyncratic factors. We can simplify the discussion a little by supposing that the stock investment is diversified using an index such as the Standard & Poors 500. The individual may be supposed to have a well-defined probability distribution over possible future paths for the price level,

derived from a tree structure where Nature can move the rate of inflation up or down at regular intervals. This distribution will be subject to learning over time, but, it may be supposed, no significant element of discovery.

By contrast, in observing past stock returns, and the explanations commonly offered for them, the individual observes that large movements in stock prices are commonly explained with reference to propositions unconsidered by most participants in the market, such as the invention of a new product. It follows that the stock market investment is subject to ambiguity. For any feasible model of stock market returns, and prior distributions over relevant parameters, the individual knows that it will be necessary, sooner or later, to replace the model with an alternative, incorporating previously unconsidered variables. Faced with such ambiguity, an investor might reasonably require a premium in returns for equity (that is, from the external viewpoint τ_0 , the expected return to equity should exceed the expected return to debt).

Now consider the Ellsberg two-urn problem, where an individual chooses between an urn with a known distribution of black and white balls and a second urn with an unknown distribution, each urn containing 100 balls. It seems reasonable to suppose that the individual has access to a representation of the world sufficiently expressive to include the proposition 'The second urn has n black balls' where n ranges from 0 to 100, but this is unhelpful in formulating a probability distribution. The ambiguity in the problem arises from the existence of a range of unconsidered propositions of the general form 'Rule x was used to determine n'. If the individual had a probability distribution over such rules, incorporating the true rule with non-zero probability, she could formulate a prior distribution over the number of balls and update it in the standard Bayesian fashion. In practice, however, the class of possible rules is so large that no-one can reasonably hope to consider all its elements.

9.1 Multiple priors

Thus far, we have considered cases where the prior distribution π_{τ} on the restricted tree τ generated by the considered propositions is induced by a unique prior π_0 on the full tree τ_0 , because all unconsidered propositions are independent of propositions expressible with respect to τ . To generate a mul-

tiple priors model, it is natural to suppose that there may be more than one such measure. An obvious way to do this is to look at the measures induced conditional on alternative implicit beliefs about unconsidered propositions.

Considering any $p' \notin P_{\tau}$, there are two induced probabilities for $p \in P_{\tau}$, namely

$$\tilde{\pi}_{\tau}\left(p;n_{0}|p'\right) = \frac{\tilde{\pi}_{0}\left(p \wedge p';n_{0}\right)}{\tilde{\pi}_{0}\left(p';n_{0}\right)}, p \in P_{\tau}$$

and

$$\tilde{\pi}_{\tau}\left(p;n_{0}|\neg p'\right) = \frac{\tilde{\pi}_{0}\left(p \land \neg p';n_{0}\right)}{\tilde{\pi}_{0}\left(\neg p';n_{0}\right)}, p \in P_{\tau}.$$

For a proposition $p' \notin P_{\tau}$ that is independent of P_{τ} in the sense that, for all $p \in P_{\tau}$

$$\tilde{\pi}_0\left(p \wedge p'; n_0\right) = \tilde{\pi}_0\left(p; n_0\right) \tilde{\pi}_0\left(p'; n_0\right),$$

we have $\tilde{\pi}(\cdot|p') = \tilde{\pi}(\cdot|\neg p')$ since, for all $p \in P_{\tau}$,

$$\begin{aligned} \tilde{\pi}_{\tau} \left(p; n_0 | p' \right) &= \frac{\tilde{\pi}_0 \left(p \land p'; n_0 \right)}{\tilde{\pi}_0 \left(p'; n_0 \right)} = \frac{\tilde{\pi}_0 \left(p; n_0 \right) \tilde{\pi}_0 \left(p'; n_0 \right)}{\tilde{\pi}_0 \left(p'; n_0 \right)} \\ &= \frac{\tilde{\pi}_0 \left(p; n_0 \right) \tilde{\pi}_0 \left(\neg p'; n_0 \right)}{\tilde{\pi}_0 \left(\neg p'; n_0 \right)} = \frac{\tilde{\pi}_0 \left(p \land \neg p'; n_0 \right)}{\tilde{\pi}_0 \left(\neg p'; n_0 \right)} = \tilde{\pi}_{\tau} \left(p; n_0 | \neg p' \right) . \end{aligned}$$

In general, however, $\tilde{\pi}_{\tau}(p; n_0 | p') \neq \tilde{\pi}_{\tau}(p; n_0 | \neg p')$, and consideration of probability values for p' in the range [0, 1] gives rise to probabilities for p in the range bounded by $\tilde{\pi}_{\tau}(p; n_0 | \neg p')$ and $\tilde{\pi}_{\tau}(p; n_0 | p')$. Thus, we can define a set of priors:

$$M(p') = \{\lambda \tilde{\pi}(\cdot | p') + (1 - \lambda) \tilde{\pi}(\cdot | \neg p') : 0 \le \lambda \le 1\}$$

The natural interpretation here is that each element of the set of multiple priors may be derived as a conditional probability measure, given a probability number for the unconsidered proposition p'. Thus p' has a status intermediate between propositions in P_{τ} that are under active consideration, and unconsidered propositions in the case of restricted Bayesianism. Although the decision-maker does not explicitly consider p', the range of multiple priors corresponds to the probability measure that would arise if p' were a considered proposition with probability λ .

For a more general version of the multiple priors model, let P^* be a finite set of unconsidered propositions, closed under \neg and \land , and let Δ be

the unit simplex with dimension equal to $K = card(P^*)$, having typical element $\boldsymbol{\delta} = (\delta_1, ... \delta_K)$ such that $\sum_k \lambda_k = 1$. For each $p_k \in P^*$, considered as an implicit belief, we have a conditional subjective probability distribution $\tilde{\pi}(\cdot | p_k)$ and we define the set of priors

$$M(P^*) = \left\{ \sum_k \delta_k \tilde{\pi}_k : \boldsymbol{\delta} \in \Delta \right\}.$$

This definition agrees with that given above for the case $P^* = \{p', \neg p'\}$.

10 Endogenous discovery

In the model presented here, the discovery of new propositions and possible events has been treated as exogenous. When the individual arrives at a given node in the objectively given tree, there is an associated more refined subjective tree, but there has been no account of how this process of refinement takes place. However, the developments presented above indicate some possible approaches to the problem.

10.1 Impossible beliefs and discovery

First, the idea of impossible beliefs leads to some minimal conditions under which discovery must take place. Suppose that the individual explicitly believes some event to have zero probability and the event occurs; this can only arise in association with the falsification of some implicit belief. In this context, Bayesian updating on a restricted domain clearly breaks down, as does any system of updating relying on conditional probabilities. Furthermore, it is impossible to 'start afresh' and derive subjective probabilities as marginal probabilities from the external probability distribution, conditional on implicit beliefs. Hence, some revision of implicit beliefs is clearly necessary, and this will normally generate a new and more refined tree structure.

Example 2 Consider further the example of observation of the emission of gamma rays, making it necessary to refine the subjective tree structure to accommodate the previously unconsidered event. This might be done in an ad hoc fashion, treating the emission of rays as an unexplained natural phenomenon. However, a more satisfactory adjustment would involve revision

of implicit and explicit beliefs about atomic structure, generating one or more models that predict the observed outcome with positive probability.

Possible requirements on the new beliefs are that they should assign positive probability to previously observed events, and that the new subjective tree should be 'close' to the old one in the topology defined by the lattice structure, and that the set of implicit beliefs should change in a 'natural' way.

There is a trade-off between these requirements. A minimal adjustment to beliefs, sufficient to accommodate evidence inconsistent with previouslyheld beliefs, may be defined by the requirement that there should be no tree intermediate in refinement between the new and old subjective trees that is sufficiently refined to assign positive probability to the observed (and previously believed impossible) event. Note that such a minimal adjustment will not, in general be unique. A more radical change in implicit beliefs may yield a better account of the observed evidence (that is, one that assigns a higher probability to the observed outcomes).

The case when tree structures and implicit beliefs are adjusted only in response to the breakdown of existing implicit beliefs corresponds reasonably closely to Kuhn's (1962) idea of a paradigm shift. In particular, it will not be true, in general, that the new set of beliefs account better for all the evidence than the old ones - evidence that confirmed the old paradigm may be anomalous (though not regarded as impossible) in the new one.

The discovery process described by Kuhn involves adherence to relatively conservative 'normal science' most of the time, with occasional radical changes. This is also true, in a more nuanced way, of Lakatos' idea of competition between scientific research programs.

10.2 Constrained-optimal discovery

An alternative approach to modelling endogenous discovery begins with a notion of optimality, derived from the external perspective. Suppose that bounded rationality is represented by a cost-of-calculation function, increasing in the lattice order, so that consideration of more refined trees (larger sets of propositions) is more costly. Then, given a known choice rule, a constrained-optimal choice of subjective tree may be defined from the external viewpoint. The constrained-optimal path may be considered, either as a possible descriptive model or as a normative benchmark.

Informally, consideration of a larger set of propositions (including possible decisions) may be seen as having a value akin to information value or option value. Thus, on a constrained optimal path, new propositions will be discovered when the associated information value exceeds the marginal calculation cost of the associated refinement. At least as a first approximation, this seems reasonable, reflecting the idea that necessity is the mother of invention. As existing representations of the state space become less satisfactory, there will be pressure to explore new ideas.

11 Choice

The development of a fully-fledged theory of choice is beyond the scope of this paper. However, we may sketch an outline of the approach indicated by the developments above.

11.1 Consequences

First, to each node in τ_0 we may attach a consequence, either expressed as an element of some outcome space or directly in utility terms. For present purposes it is most usefull to assume monetary payoffs. Now for each node in a subjective tree τ , we similarly attach a payoff selected from the set r^{-1} . Note that there is no loss of generality in selecting a unique payoff for each subjective node - if multiple payoffs are considered possible we can refine the subjective tree accordingly.

A natural benchmark is obtained from the position of the unboundedly rational outside observer, assumed to share the same preferences as the subjective individual. In particular, we assume the same ranking of lotteries over completed histories. For simplicity, we will posit preferences given by expected utility in each period, with stationary discounting over time.

11.2 Expected utility

As has already been shown, Bayesian updating of beliefs on a restricted state space is consistent, under the condition that unconsidered propositions are stochastically independent of considered propositions. Suppose that, in addition, outcomes are expressed in terms of monetary payoffs (or, more generally, some interval on the real line) and that preferences display CARA. Then, from the external perspective, uncertainty about the outcomes arising from the realisation of unconsidered propositions does not affect optimal choices over acts with payoffs measurable with respect to considered events. Thus, maximization of expected utility on the restricted domain is optimal.

Now consider discovery. This gives the individual access to previously unconsidered propositions and events, including, in general, decision propositions corresponding to choices with outcomes measurable with respect to the new event space. Suppose that the independence assumption remains valid; that is, remaining unconsidered events are stochastically independent of the newly discovered events as well as those that are already considered. Then, expected utility is optimal for the newly discovered choices. Moreover, the independence of the newly discovered and previously independent events means that, under CARA, the individuals optimisation problem may be subdivided into separate optimisation problems, 'new' and 'old'.

Extending this idea over time, the optimal strategy for a boundedly rational individual in these circumstances is to apply expected utility in each 'small world' as the new choice problem is revealed. This seems a more plausible account of reasonable EU-maximising behavior than does the standard model of a single contingent strategy, adopted at the beginning of the individual's lifetime and implemented consistently thereafter in response to observed signals and the associated updating of the state space.

11.3 Minmax EU and other multiple-priors preferences

Much of the discussion of EU may be extended to the case of multiple priors. The requirements for consistent updating have already been discussed, and the CARA assumption plays a similar role. The main interest, therefore is in the choice between minmax EU, maxmax EU and intermediate possibilities such as α -CEU (see for example, **Jaffray and Philippe 1997**). As a rule

for choice in the presence of unconsidered propositions, minmax EU may be seen as embodying a maximally pessimistic assumption, namely that, whatever decisions the individual makes, 'surprises' will be maximally unpleasant, while the opposite is true for maxmax.

Such assumptions are embodied in many pieces of popular wisdom such as 'Murphy's Law', and, in the 'precautionary principle', widely advocated as a basis for choices in the presence of environmental uncertainty. The rationale for such a principle is evident from considering the engineering context from which Murphy's law emerges. Any complex engineering system has many possible states, of which only a small minority are likely to yield optimal performance as intended by the designers. Furthermore, these states are likely to be relatively well-understood in terms of propositions considered in the associated design theory. The outcome of attempts to operate the system outside the limited parameter (sub)space associated with these propositions will depend on a range of unconsidered propositions, and at least some of these propositions are likely to turn out badly.

Natural environments may similarly be seen as complex systems. The precautionary principle is associated with a general view of these systems as being finely-adjusted, with a propensity for breakdown in response to shocks. Hence, it seems reasonable to assume the worst as regards the consequences of poorly-understood innovations, such as the introduction of new pesticides. By contrast, optimistic views of the environment as highly resilient lead to the opposite view that restrictions based on environmental concerns should be imposed only on the basis of proven damages.

11.4 Rank-dependent, Choquet and NEO-EU

Rank-dependent choice representations are, in a formal sense, closely related to multiple-priors, at least within the framework of a fully-specified state space. For example, any RDEU model with concave (convex) transformation of the cumulative distribution function can be represented as a multiple prior model with minmax (maxmax) EU. However, the main case of interest, that of overweighting extreme outcomes, has a different motivation and distinct properties.

Suppose, instead of independence, that the outcomes associated with the

realization of unconsidered events are likely to be positively correlated with those associated with the realization of considered events (given the assignment of consequences to nodes in the subjective tree). Then an optimal decision strategy, expressed in terms of the distribution of considered outcomes, will place more weight on the extreme events than on intermediate events.

The NEO-EU model (Chateauneuf, Eichberger and Grant, 2006) is a polar case of this, with a particularly appealing interpretation in the current context. The weight on the linear component may interpreted as confidence in the implicit beliefs used to derive probabilities for considered events. Conversely, the combined weight on the extremes may be seen as an estimate of the likelihood that the implicit beliefs will be violated in the given choice context, with their second parameter determining the likelihood ratio of favorable to unfavorable surprises.

On the one hand the occurrence of 'unlikely' events will reduce confidence in the implicit beliefs, leading to an increase in the combined weight placed on the extremes, as in Eichberger, Grant and Kelsey (2006). On the other hand, the discovery of new propositions and the associated revision of beliefs may lead to a more plausible model of the world, and therefore a lower combined weight placed on the extremes. Thus, in the model presented here, unlike a model of learning with no new discovery, the combined weight on the extremes need not increase over time.

11.5 Behavioral implications

The model presented here is very general, as is the alternative of a state-space model with unbounded rationality. Hence, any behavioral test requires more precise specification of the maintained hypothesis, and associated auxiliary hypothesis and the alternatives being tested. A natural choice of maintained hypothesis is that, under unbounded rationality, preferences would obey subjective expected utility for some appropriately updated Bayesian prior. We may then consider tests of this hypothesis, with auxiliary hypotheses of varying strength, against alternatives requiring rejection of the joint maintained and auxiliary hypotheses.

If states and probabilities are assumed to be known, then the maintained

hypothesis becomes the standard von Neumann-Morgenstern model of EU under risk, which may be tested for possible violations such as Allais common consequence and common ratio effects. However, this does not seem to be a satisfactory test, since the combination of the maintained and auxiliary hypotheses is too strong. Most analysts would prefer to explain Allais effects with reference to non-EU preferences, without discarding the whole idea of a state space. However, as shown above, rank-dependent and NEO-EU models may be interpreted as heuristic responses to incompleteness of the state space.

If the state space is assumed known, but probabilities are assumed to be subjective, the maintained hypothesis becomes the Savage SEU model, which may be rejected in favour either of probabilistically sophisticated alternatives (by showing that the axioms of Machina and Schmeidler are satisfied and then testing EU for the implied probabilities) or of models in which probabilities are unknown or ambiguous using tests such as those proposed by Ellsberg. Again, most analysts would not feel it necessary to abandon the idea of a known state space in such cases, although the derivation of the multiple priors model presented above suggests that this is one possible explanation.

Perhaps the most satisfactory test arises from considering the treatment of unspecified state spaces by Dekel, Lipman and Rustichini (2001). They provide conditions under which the existence of a subjective state space with EU may be inferred from choices over menus, even though neither the state space nor the probabilities are specified in the choice problem. Conversely, given the maintained hypothesis of EU, an experimental rejection of the conditions derived by Dekel, Lipman and Rustichini might be interpreted as evidence that there exists no completely specified state space.

In this paper, we have discussed the case of a single agent, focusing on learning and discovery over time. However, the lattice structure adopted here is similar to that of the multi-agent model developed by Heifetz, Meier and Schipper (2006), who examine 'interactive unawareness' (see also Heifetz 2006). Grant, Kline and Quiggin (2006) show that the ideas of bounded rationality can be used to model ambiguity in contracting, giving rise to a range of behavioral implications potentially observable in market settings.

12 Concluding comments

Observation of, and introspection about, decision-making reveals many phenomena that are excluded from standard decision-theoretic models. Plans are revised, or abandoned, in the light of unforeseen contingencies. Beliefs held with certainty are falsified by events. New ideas arise spontaneously or in response to the discovery that previous understandings of the world were incomplete or unsatisfactory. In this paper, we have presented a framework within which such phenomena can be modelled.

References

- Allais, Maurice. "Le comportement de l'homme rationel devant le risque: critique des axioms et postulates de l'école Americaine." Econometrica 21 (1953): 503 – 46.
- [2] Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite and David Schmeidler. "Fact-Free Learning." American Economic Review 95, no. 5 (2005): 1355 – 1368.
- [3] Chew, Soo Hong. "A Generalisation of the Quasilinear Mean With Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox." Econometrica 51, no. 4 (1983): 1065 92.
- [4] Belnap, Nuel Perloff, Michael and Xu, Ming. Facing the Future: Agents and Choices in Our Indeterminist World. Oxford University Press, Oxford, 2001
- [5] Chateauneuf, Alain, Jürgen Eichberger and Simon Grant. "Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities." mimeo (2006).
- [6] Cubitt, Robin, and Robert Sugden. "On Money Pumps." Games and Economic Behavior 37 (2001): 121 – 60.

- [7] Dekel, Eddie. "An Axiomatic Characterization of Preferences Under Uncertainty: Weakening the Independence Axiom." Journal of Economic Theory 40 (1986): 304 – 18.
- [8] Dekel, Eddie, B. Lipman, and A. Rusticchini. "A Unique Subjective State Space for Unforeseen Contingencies." Econometrica 69 (2001): 891 - 34.
- [9] Eichberger, Jürgen, Simon Grant and David Kelsey. "Updating Choquet Beliefs." mimeo (2006).
- [10] Ellsberg, Daniel. "Risk, Ambiguity and the Savage Axioms." Quarterly Journal of Economics 75, no. 4 (1961): 643 – 69.
- [11] Epstein, L. and Massimo, Marinacci (2006), 'Coarse contingencies', mimeo.
- [12] Ghirardato, Paolo. "Coping with ignorance: unforeseen contingencies and non-additive uncertainty." Economic Theory 17, 247 – 276 (2001).
- [13] Gilboa, Itzhak, and David Schmeidler. "Maxmin Expected Utility With Non-Unique Prior." Journal of Mathematical Economics 18, no. 2 (1989): 141 – 53.
- [14] Gilboa, Itzhak, and David Schmeidler. "Updating Ambiguous Beliefs." Journal of Economic Theory 59 (1994): 33 – 49.
- [15] Gilboa, Itzhak, and David Schmeidler. "Case-Based Decision Theory." Quarterly Journal of Economics 110 (1995): 605 – 39.
- [16] Grant, Simon, and John Quiggin. "Unforeseen Contingencies: a Propositional Approach." (2006), in Munier, B. et al (eds), Proceedings of FUR XI Conference, Paris 2004, Springer, forthcoming.
- [17] Halpern, Joseph. "Alternative Semantics for Unawareness." Games and Economic Behavior 37 (2001): 321 – 39.
- [18] Halpern, Joseph. Reasoning About Uncertainty. Cambridge, Massachusetts: The MIT Press, 2003.

- [19] Halpern, Joseph, and Leandro Chaves Rêgo. "Interactive Unawareness Revisited." unpublished working paper (2005)
- [20] Halpern, Joseph, and Leandro Chaves Rêgo (2006a), 'Reasoning About Knowledge of Unawareness', Proceedings of the Conference on Knowledge Representation and Reasoning
- [21] Halpern, Joseph, and Leandro Chaves Rêgo (2006b), 'Extensive Games with Possibly Unaware Players', Conference on Autonomous Agents and Multiagent Systems.
- [22] Hart, Sergiu, Aviad Heifetz and Dov Samet. "Knowing Whether', 'Knowing That', and the cardinality of state spaces." Journal of Economic Theory 70 (1996): 249 – 256.
- [23] Heifetz, Aviad, Martin Meier, and Burkhard Schipper. "Interactive Unawareness." Journal of Economic Theory forthcoming (2005).
- [24] Heifetz, Aviad. "Game states", mimeo (2006).
- [25] Hendry, David (1987). Econometric Methodology: A personal perspective, in T. F. Bewley (ed.) Advances in Econometrics Fifth World Congress, Volume II, Cambridge University Press, Cambridge.
- [26] Jaffray, Jean-Yves and Fabrice Philippe. "On the Existence of Subjective Upper and Lower Probabilities." Mathematics of Operations Research 22, no. 1 (1997): 165-185.
- [27] Jeffrey, Richard C. Formal Logic: Its Scope and Limits, McGraw-Hill, New York 1990.
- [28] Kahneman, Daniel, and Amos Tversky. "Prospect Theory: an Analysis of Decision Under Risk." Econometrica 47, no. 2 (1979): 263 – 91.
- [29] Kaneko, M. and J. Jude Kline. "Inductive Game Theory: A Basic Scenario", Tsukuba University, Department of Social Systems and Management Discussion Paper Series No. 1143 (2006).

- [30] Karni, E. (2005), 'Subjective Expected Utility Theory without States of the World', Johns Hopkins University School of Economics Working Paper,WP523.
- [31] Karni, Edi and Zvi Safra (1990), 'Behaviorally consistent optimal stopping rules', Journal of Economic Theory, 51(2), 391-402.
- [32] Kreps, David. "A representation theorem for 'Preference for flexibility'." Econometrica 47, 565 – 576 (1979).
- [33] Kreps, David. "Static choice and unforeseen contingencies," in Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn, ed. by Partha Dasgupta, Douglas Gale, Oliver Hart and Eric Maskin. Cambridge, MA: MIT Press, 259 – 281 (1992).
- [34] Kuhn, Thomas (1962) The Structure of Scientific Revolutions, University of Chicago Press, Chicago.
- [35] Lakatos, Imre. "Falsification and the methodology of scientific research programmes". in Lakatos, I. M. and A. (ed.), Criticism and the Growth of Knowledge, Cambridge University Press, Cambridge, 1970.
- [36] Machina, Mark. "Expected Utility' Analysis Without the Independence Axiom." Econometrica 50, no. 2 (1982): 277 – 323.
- [37] Modica, Salvatore, and A. Rusticchini. "Awareness and Partitional Information Structures." Theory and Decision 37 (1994): 107 – 24.
- [38] Modica, Salvatore, and A. Rusticchini. "Unawareness and Partitional Information Structures." Games and Economic Behavior 27, no. 2 (1999): 265 – 98.
- [39] Mukerji, Sujoy. "Understanding the nonadditive probability decision model." Economic Theory 9, 23 – 46 (1997).
- [40] Nehring, Klaus. "Preference for flexibility in a Savage framework." Econometrica 67, 101 – 120 (1999).
- [41] Popper, Karl. Conjectures and Refutations: the Growth of Scientific Knowledge. New York: Harper and Row, 1963.

- [42] Popper, Karl. " On the Sources of Knowledge and of Ignorance." In Conjectures and Refutations: the Growth of Scientific Knowledge, New York: Harper and Row, 1963.
- [43] Savage, L.J. Foundations of Statistics. N.Y.: Wiley, 1954.
- [44] Schmeidler, David. "Subjective probability and expected utility without additivity." Econometrica 57 (1989): 571–87.
- [45] Segal, Uzi. "Anticipated Utility: a Measure Representation Approach." Annals of Operations Research 19 (1989): 359 – 74.
- [46] Tversky, Amos, and Daniel Kahneman. "Cumulative Prospect Theory: an Analysis of Attitudes Towards Uncertainty and Value." Journal of Risk and Uncertainty 5, no. 3 (1992): 297 – 323.
- [47] Yaari, Menahem. "The Dual Theory of Choice Under Risk." Econometrica 55, no. 1 (1987): 95 – 115.



Figure 1: Decision Tree for Electron Beam Experiment without Unconsidered Possibility of Gamma Ray Emission.



Figure 2: Tree for Electron Beam Experiment that Includes Possibility of Gamma Ray Emission.



Figure 3: A Coarsening of the Tree in Figure 2



Figure 4: Direct (and Indirect) Mapping from tree τ'' to coarser tree τ (via tree τ').



Figure 5: Commuting Information Correspondences for tree $\tau \sqsubseteq \tau_0$



Figure 6: Commuting Information Correspondences for $\tau \sqsubseteq \tau'$.