Revising Knowledge: A Hierarchical Approach[1]

by Stephen Morris

Department of Economics, University of Pennsylvania
and
Department of Applied Economics, University of Cambridge.

Introduction

A tradition in economics seeks to explain behavior by identifying the underlying preferences which (in a given environment) generate that behavior. This approach was extended in the work of Savage [1954] to decision making under uncertainty. "Belief", and thus "knowledge", are interpreted as properties of individuals' preferences (and thus, in principle, observable choices).

This "revealed preference" approach provides an alternative way of thinking about rules for revising beliefs and knowledge which have been a major concern for philosophers and computer scientists in recent years. Unfortunately, economists typically use the revealed preference type argument to support a model of super-rational actors who use Bayes rule to update beliefs and completely understand the structure of the world. Formally, the standard model of information as a partition of states of the world is equivalent to Kripke's S5 system[2]. This paper describes a project investigating how weaker decision theoretic axioms generate a more plausible description of agents' knowledge and knowledge revision. The results are independent of whether expected utility maximization is assumed to hold.

The paper is organized as follows. The first part of the paper argues why the economists' standard model needs to be weakened. Partition information makes too strong implicit assumptions about decision makers' knowledge, such as that the decision maker can anticipate what he would have known in every possible (but not realized) state of the world. But weakening partition information, while maintaining the standard models of beliefs, is

---

1. Financial support from National Science Foundation Grant #SES-9308515 is gratefully acknowledged.
2. Kripke [1963].

inconsistent. Bayes updating suffers from exactly the same critiques as partition information. Therefore a unified approach to modelling the revision of beliefs and knowledge is required.

The second part of the paper outlines work in this direction. Decision theoretic assumptions about how preferences vary across states can generate a natural model of knowledge and knowledge revision which is weaker than partition information.

The third part of the paper addresses a problem which arises in the work outlined in the second part. There is a circularity in describing decision makers' preferences over state contingent acts when (implicitly) states are distinguished by the preferences of decision makers at those states. It is necessary to carry out a hierarchical treatment of preferences which addresses this circularity, in the same way that hierarchies of belief and knowledge have addressed circularities which arose in those contexts.

Problems with the standard model

There is general agreement among economists on a standard framework for modelling choice under uncertainty with changing information. This general agreement has contributed to the dramatic and valuable growth of information economics in the last twenty years. But this standard framework is flawed in the way the static model of choice under uncertainty is integrated with dynamic descriptions of changing information. The flaw is conceptual: the relationship between preferences, beliefs and knowledge is incoherent. But the flaw also has profound practical implications for the predictions of economic models. In particular, it constrains us to make too strong (and very implausible) assumptions about the nature of knowledge of decision makers.

This section first outlines the standard model, where static choice under uncertainty is generated by expected utility maximization and changing information is generated by partitions over a state space. The assumptions about decision makers' knowledge implicit in this framework are then shown to be implausibly strong. Attempts to splice Bayes updating with weaker (and more plausible) models of knowledge are shown to be incoherent in both theoretical and practical senses.

It is useful to present formally a specific version of the "standard" model, with a finite state space, "subjective" expected utility maximization and partition information. However, the critique of this section would apply to any version of the standard model in common usage.

Suppose there is some finite set, Ω, of possible "states of the world", and a finite set of "consequences", C, which might occur at any of the states of the world. An "act", x, available to the decision maker is assumed to imply a known consequence at each state of the world. Thus the act can be thought of as a vector in C^{Ω}, where the ωth co-ordinate of act x, x_{ω}, is interpreted as the consequence the decision maker would face in state ω if he chose act x. The decision maker is assumed to have a preference relation, $\succeq$, over all possible acts. The decision maker is said to be an expected utility maximizer if there exists a function u: $C \rightarrow \mathbb{R}$ and probability distribution π on Ω such that act x is weakly preferred to act y ($x \succeq y$) if and only if the expected value of the utility of the consequence of act x is no smaller than the expected value of the utility of the consequence of act y. The seminal contribution of Savage [1954] was to show that the decision maker is an expected utility maximizer if and only if the preference relation satisfies completeness, transitivity and certain independence and continuity axioms[3]. The detail and plausibility of those axioms are not our concern here. The key observation for our purposes is that restrictions on preferences are used to guarantee the existence of some beliefs which are somehow reflected in the decision maker's choices.

Criticisms of the standard model in recent years have focussed on the expected utility hypothesis. A series of well-documented experimental violations of that hypothesis and theoretical work showing the tractability of weaker formulations have together made a strong case for rejecting the expected utility hypothesis[4]. The work described in this paper, however, is concerned with some orthoganal problems which exist independent of whether the expected utility model is accepted or rejected. These problems arise with the integration of whichever static model of choice under uncertainty is chosen, and the model of changing information. For ease of exposition, however, I will continue to focus on the expected utility case[5].

Savage's theory, and most expected utility and non-expected utility theories before and since, are essentially *static*. They describe preferences at a point in time over acts or state-contingent consequences. Having devoted considerable intellectual resources in developing static

3. Savage [1954] proved his theorem for an infinite state space. Gul [1992] gives an analogous result for the finite state space, subjective case outlined here.
4. Machina [1989].
5. Morris [1992b] considered the problem of defining knowledge in a situation where static preferences do *not* reflect expected utility maximization.

results, the next step to incorporate changing information is typically rather cursory. Represent information by a *partition* of the state space, P. Thus P is a collection of disjoint subsets of Ω, whose union is Ω. The interpretation is that if the true state is ω, the decision maker knows only that the true state is in $P(\omega)$, the (unique) element of P which contains ω. Changing information can then be represented by a sequence of partitions, $\{P_t\}_{t=1,..T}$, where P_t represents the decision maker's information at time t. It is then natural to impose the restriction that information improves through time, so that if t is greater than s, $P_t(\omega)$ is a subset of $P_s(\omega)$, for every state ω.

Now a dynamic expected utility representation of preferences is generated by the assumption that, at each state ω and date t, the decision maker will choose an act which maximizes expected utility *conditional* on the event $P_t(\omega)$.

It was noted above that the combination of the partition representation of information and static expected utility maximization has proved analytically powerful. But what does it *mean*? In the static framework, assumptions on preferences implied the existence of beliefs and a utility function which could be used to represent preferences. Our dynamic representation is exogenously specifying which states the decision maker considers possible. At the very least, then, it is making implicit assumptions on the decision maker's preferences through time. I will return to this issue at the end of this section. But first I consider an independent problem with the partition representation of information.

A partition seems at first to be an innocuous representation of knowledge. An intuitive justification might go as follows. Suppose the description of each state includes every relevant aspect of the world. Define a relation, ~, between states, with the interpretation that $\omega \sim \omega'$ means "the decision maker cannot be distinguish state ω from ω'". The interpretation seems to require that the relation be reflexive and symmetric, so that ~ is an equivalence relation. But now the equivalence classes of the relation, ~, generate a partition of Ω, which can be thought of as representing the decision maker's information.

This justification is fine if we think of the decision maker understanding (in some informal, unmodelled, way) the structure of Ω. He can conceive of every possible state. But we will see that this implicit assumption is not necessary for formal analysis and is too strong in a specific sense that will become clear. I do so by outlining work which derives the partition model of information in another way, from explicit assumptions about knowledge. The most elegant and natural way to do this is by *constructing* the state space from truth valuations of propositions, including propositions about decision makers' knowledge. This approach has been used in the economics literature by Shin [1992] and Samet [1990]. For many purposes, however, it is sufficient to identify propositions with the event where they are true, an approach which makes clearer the relation to standard economic models. Such an approach is outlined, for example, in Aumann [1989], Binmore and Brandenburger [1990], and Geanakoplos [1992].

Again let Ω be a finite set of possible worlds. A decision maker's knowledge can be represented by a mapping $K: 2^{\Omega} \rightarrow 2^{\Omega}$, with the interpretation that the decision maker knows event A at state ω if and only if $\omega \in KA$. Discussion of knowledge in other contexts suggest that the following properties of such *knowledge operators* will be important. Write -A for the complement in Ω of event A. Then I will be concerned with the following seven properties of knowledge, assumed to hold for all events $A,B \subset \Omega$:-

Properties **[K1]** $K(\Omega) = \Omega$
[K2] $K(\emptyset) = \emptyset$
[K3] $K(A) \cap K(B) \subset K(A \cap B)$
[K4] $A \subset B$ implies $K(A) \subset K(B)$
[K5] $K(A) \subset A$
[K6] $K(A) \subset K(K(A))$
[K7] $-K(A) \subset K(-K(A))$

[K1] requires that events which are necessarily true (true at all possible worlds) are always known to be true. [K2] requires that events which are necessarily false are never known to be true. [K3] requires that if A is known and B is known, then $A \cap B$ is known. [K4] requires that if A implies B and A is known, then B is known. [K5] requires that if something is known to be true, it is indeed true. [K6] requires that if something is known, then it is known that it is known. [K7] requires that if something is not known, it is known that it is not known.[6]

Axioms [K1] through [K6] are standard in the philosophical literature on knowledge. Axiom [K7], however, is making a very strong requirement on knowledge. It has been shown that all seven axioms together imply a system of knowledge which is equivalent to partition information. In particular, it is natural to say that the decision maker cannot distinguish state ω from state ω' if ω' is an element of every event known at state ω. This indistinguishability

6. Note that there is considerable redundancy in the set of seven axioms. In particular, [K5] implies [K2] and [K5] and [K7] together imply [K6].

relation can then be shown to be reflexive and symmetric if and only if the seven axioms are satisfied.

There is no reason to suppose that axiom [K7] is natural in economic models of choice under uncertainty[7]. Therefore a number of authors have been led to analyze the implications of weaker models of knowledge (while maintaining the static expected utility assumption). For example, Aumann [1976] showed, in a partition information setting, that it cannot be common knowledge that decision makers with common prior beliefs have different posterior beliefs about some event. Samet [1990] showed that this result *still* holds when the objectionable axiom [K7] is dropped. A series of papers have shown that differences in information alone cannot lead decision makers with common prior beliefs to want to trade assets (see, for example, Milgrom and Stokey [1982]). This is often interpreted as implying *no speculation*. Geanakoplos [1989] showed necessary and sufficient conditions on knowledge for this kind of result in different settings. In no setting was the objectionable axiom [K7] required. Rubinstein and Wolinsky [1990], Brandenburger, Dekel and Geanakoplos [1992] and Morris [1991] give further results in this spirit.

This work certainty illustrates the feasibility of looking at weaker knowledge systems in an economic context. The weakness of this approach lies in the way beliefs are introduced. It is assumed that there exists a common prior distribution over all states of the world, and that given the set of states which the decision maker believes possible (with his possibly non-partitional information), his posterior beliefs are derived by standard Bayes updating. But a decision theoretic justification of Bayes updating requires some kind of dynamic consistency assumption, which *also* implies partition information (this issue was explored in Morris [1992a]). This suggests that a *unified* treatment of beliefs and knowledge is required. How would this be done?

---

7. Philosophers make a distinction between belief with probability one (which has no objective implications) and knowledge, or "true belief", which in addition entails truth. Property [K7] may then be accepted as a property of belief (with probability one) but not as a property of knowledge. As will become clear below, I am not making that important distinction in this paper. An alternative - informal - critique of [K7] is that it seems to imply (in a way that, for example, [K6] does not) that individuals understand the global structure of the state space. Modica and Rustichini (1993) have recently attempted to formally model such "unawareness" of certain contingencies.

I will say that a decision maker *knows* an event if he assigns it probability one. Economists have typically used this language, at least for the finite state space setting that I am considering here. The reader is warned that this terminology is inconsistent with the clear and important distinction which philosophers make between belief with probability one (a subjective notion) and knowledge, which also entails truth. I nonetheless maintain this terminology because it will turn out in what follows that - under the assumptions I make - belief with probability one does indeed imply knowledge.

With knowledge *defined* as a property of beliefs, it ought to be possible to derive any interesting properties of knowledge from underlying properties on beliefs. But Savage's [1954] expected utility representation theorem showed that it is possible to think of beliefs as a parameter of preferences. Therefore it is possible to derive rules for updating beliefs through time, and use those rules to derive rules about updating knowledge. This is the approach in Morris and Shin [1993]. However, it is also possible to think of knowledge directly as a property of preferences, skipping the role of beliefs. This is the approach pursued in the next section. But the axioms on knowledge involve statements about what is known at different states of the world. Thus, if knowledge is to be interpreted as a property of preferences, we will have to be concerned with different sets of preferences at different states of the world.

An axiomatic derivation of knowledge from preferences

The standard model must be extended to deal with the problems outlined in the previous section. A first attempt to do so proceeds as follows. Take the approach of the standard model, with a finite state space, Ω, capturing all relevant states of the world, a finite set of consequences, C, and state-contingent "acts" in C^{Ω}. But now instead of studying a single preference relation over acts, define a *different* preference relation over acts at each state of the world ω, $\succsim_{\omega}$. Thus $x \succsim_{\omega} y$ is interpreted to mean that, if the true state is ω, the decision maker prefers act x to act y. This section builds on results in Morris [1992b] to show how restrictions on preferences across states can be used to derive standard but weak axioms of knowledge. These results are independent of whether expected utility maximization is assumed.

The key to this approach is to identify *knowledge* as a property of preferences. In Morris [1992b], a number of different ways to do so are identified, but here I focus on just one. For event $A \subset \Omega$, let x_A denote the tuple $\{x_{\omega}\}_{\omega \in A}$, i.e. the (ordered) collection of consequences implied on A by act x. Now say that the decision maker knows an event A at state ω, if his preferences over acts never depend on the consequences outside event A. Thus:-

$$K(A) = \{ \omega \in \Omega \mid (x_A, y_{-A}) \succeq_\omega (x_A, z_{-A}) \ \forall\, x,y,z \in C^\Omega \}$$

This is closely related to Savage's notion of null events: a decision maker knows event A if the complement of A is Savage-null. Thus knowledge is defined as a property of preferences. The basic axioms of knowledge discussed in the previous section can be related to properties of preferences. In particular, knowledge satisfies [K1] if and only if preferences are reflexive. Knowledge satisfies [K2] if and only if preferences are non-trivial, i.e $x \not\succeq_\omega y$ for some x, y in C^Ω. Knowledge satisfies [K3] if preferences are transitive. Finally, knowledge always satisfies [K4].

But thus far there have been no axioms relating together the choices made at different states of the world. That is, I have said that $\succeq_\omega$ must satisfy some property, for each ω, but not considered properties which relate the collection $\{\succeq_\omega\}_{\omega\in\Omega}$ to each other. This is exactly what I must do to derive the plausible properties of knowledge operators, that something known is true (K5), and that if something is known, it is known that it is known (K6).

I will say that decision maker preferences are *coherent* if, in any decision problem, the decision maker makes optimizing decisions at each state (given his preference ordering at that state), the resulting profile of choices cannot make him worse off (given his preferences at *any* state) than any constant decision rule. This notion is generalization of the idea that information is *valuable* in the sense of Blackwell [1951]. Information cannot make you worse off. This is certainly true of partition information. I am requiring it to be true for our more general information structures. This property was key in Geanakoplos' [1989] results about no speculation in economic environments, discussed in the previous section.

Let us state the coherence condition formally. A *decision problem* is a finite subset of acts, $D \subset \mathbb{R}^\Omega$. For each subset, I can define for each state the set of best choices, $B_\omega[D] = \{x \in D \mid x \succeq_\omega y, \text{ for all } y \in D\}$. By finiteness of D, and completeness and transitivity of each $\succeq_\omega$, these best choice sets are non-empty for every decision problem. Now what will happen if the decision maker makes an optimizing choice at each state of the world? A *decision rule* is a function $f: \Omega \to D$ and an optimizing decision rule satisfies $f(\omega) \in B_\omega[D]$ for all $\omega \in \Omega$. Such an optimizing decision rule generates a new act y, not necessarily in D, with $y_\omega = f_\omega(\omega)$, for all $\omega \in \Omega$, (where $f_{\omega'}(\omega)$ is the ω'th co-ordinate of $f(\omega)$). Thus the class of acts generated by optimizing decision rules is $B^*[D] = \{x \in \mathbb{R}^\Omega \mid x_\omega = f_\omega(\omega), \text{ for all } \omega \in \Omega, \text{ for some } f:\Omega \to B_\omega[D]\}$. A minimal coherence condition on preferences is that acts in $B^*[D]$ (which may or may not be contained in D) must be at least as good as acts in the original decision problem, D. Preferences satisfy *coherence* if, for each finite $D \subset \mathbb{R}^\Omega$, there exists $x \in B^*[D]$ such that $x \succeq_\omega y$, for all $\omega \in \Omega$, $y \in D$.

It is shown in Morris [1992b] that if preference satisfy certain monotonicity and continuity properties (strictly weaker than those implying static expected utility maximization), then there exist preferences satisfying the usual conditions and coherence if and only if knowledge satisfies axioms [K1] through [K6]. This does not require the final and dubious axiom [K7] about knowledge, that if something is not known, then it is known that it is not known.

These results are very powerful. Unfortunately, this approach too has a problem. Decision makers are assumed to have preferences, at each state of the world, over acts which depend on the true state of the world. This in itself is not objectionable. But, by not assuming partition information, I have implicitly assumed that in some sense decision makers do not completely understand the structure of the state space. There is a circularity here: decision makers' knowledge is described in terms of their preferences over acts which depend on the states of the world which (implicitly) contain a description of their knowledge. The natural way to deal with this circularity is to introduce a hierarchical way of thinking about preferences. This is discussed in the next section.

A hierarchical approach to preferences and knowledge

It will be useful to review some issues concerning hierarchies of knowledge and belief, before moving on to our plan for constructing a hierarchical description of preferences. Since our whole approach is that knowledge is a special case of beliefs and beliefs are a property of preferences, it is hardly surprisingly that there are close connections.

In contexts involving many decision makers, it is often not enough to know what each decision maker believes about the physical world. Each decision maker also has beliefs about what other decision makers believe. A number of authors (Mertens and Zamir [1985], Böge and

Eisele [1979] and Brandenburger and Dekel [1993]) have constructed hierarchies of such beliefs about beliefs in the following manner. Suppose there is some set of possible physical states of the world. The first order beliefs space consists of the cross product of the original set of states and the set of possible beliefs over that space for each decision maker. The second order beliefs space consists of the cross product of the first order beliefs space and the set of possible beliefs over that space for each decision maker. Thus all higher order beliefs spaces can be constructed iteratively. But a key *coherence* restriction on beliefs is imposed at each iteration: it is required that each decision maker's kth and (k+1)th order beliefs generate the same beliefs over events of order k or lower. These constructions are used to argue that the circularity that arises in assuming common knowledge of the structure of the environment can be got around in some sense by a correct hierarchical treatment of the problem.

An analogous exercise for knowledge was performed by Fagin, Halpern and Vardi [1991]. Assumptions of *correctness*, *introspection*, and *extendibility* play a role analogous to coherence in the beliefs case. Samet's [1990] construction of a state space from fundamental propositions and higher order propositions addresses different questions, but has a hierarchical flavor (when we interpret higher order propositions about knowledge to be higher in the hierarchy). Samet imposes the standard knowledge assumptions that everything known is true (K5) and if something is known, it is known that it is known (K6), which then play a role analogous to the coherence conditions in the beliefs hierarchy. Shin's [1992] construction goes one step deeper, constructing a hierarchy of propositions based on levels of knowledge, where the coherence requirement is a *provability* condition imposed throughout the hierarchy: a proposition is known if it can be proved from other known propositions. This turns out to be equivalent to Samet's assumptions.

I conjecture that a powerful, reasonable and usable model of knowledge, beliefs and preference will come from constructing a hierarchy of statements about preference. The role analogous to all the coherence conditions in the earlier hierarchies will be played by a condition requiring that information is valuable, as discussed in the previous section. In the remainder of this paper, I give a sketch of how such a hierarchy could be constructed. It involves turning statements about decision makers' preferences into propositions, following Ramsey [1926] in his work anticipating Savage's expected utility results. The following construction closely follows Shin's [1992] construction for knowledge statements.

Consider some set of propositions Φ. Assume that Φ is closed under the logical operations of negation, $\sim p$, and conjunction, $p \wedge q$. Thus if there is some proposition p in Φ, there is another proposition "not p" (written $\sim p$) also in Φ. If there are propositions p and q in Φ, there is another proposition "p and q" (written $p \wedge q$) in Φ. A truth assignment is a function specifying, for each proposition, whether it is true of false. Thus let ω be a function $\omega: \Phi \rightarrow \{0, 1\}$, with the interpretation that $\omega(p) = 1$ means proposition p is true and $\omega(q) = 0$ means proposition q is false. Write $t(\Phi)$ for the set of such truth assignments, i.e. $t(\Phi) = \{ \omega \mid \omega: \Phi \rightarrow \{0, 1\} \}$. I will be concerned with truth assignments which are logically consistent, so that "p" and "not p" cannot simultaneously be true,

(N) $$\omega(\sim p) = 1 - \omega(p)$$

and "p and q" is true if and only if both are true:-

(C) $$\omega(p \wedge q) = \omega(p)\, \omega(q)$$

I will be interested in logically consistent truth assignments, that is those truth assignments where properties (N) and (C) hold for every proposition in Φ. Write $t^*(\Phi)$ for the set of such consistent truth assignments, i.e. $t^*(\Phi) = \{ \omega \in t(\Phi) \mid \text{(N) and (C) hold for all } p,q \in \Phi \}$. The set of consistent truth assignments, $t^*(\Phi)$, represents the set of possible "states of the world" generated by the set of propositions, Φ.

I want to go on to incorporate statements about our decision maker's preferences over consequences contingent on the state of the world. Suppose C is some set of "consequences". A given "act" x of the decision maker would imply a certain consequence depending on which propositions are true. Thus act x is a function, $x: t^*(\Phi) \rightarrow C$. We will be interested in the set of acts. Thus let $A(\Phi, C) = \{ x \mid x: t^*(\Phi) \rightarrow C \}$. Decision maker i's preferences can then be represented by a binary relation, $\succeq_i$, on $A(\Phi, C)$. Recall that mathematically a binary relation on set $A(\Phi, C)$ can be thought of as a subset, X_i, of the product set, $A(\Phi,C) \times A(\Phi,C)$. Thus $x \succeq_i y$ is simply shorthand for the statement $(x,y) \in X_i$.

Now an iterative approach to incorporating propositions about preference is natural. Begin with a set of propositions Φ^0, closed under logical operations. Think of each proposition as a description of some feature of the world which is external to the decision makers (not under anyone's control). In particular, Φ^0 contains no reference to either the knowledge or the preferences of any decision maker. I will to expand this set of propositions to include, in addition, propositions of the form "decision maker i prefers act x to act y" (written $x \succsim_i y$), where x and y are acts in $A(\Phi^0, C)$. Formally, let Φ^1 be the intersection of all sets which satisfy the following three conditions:-

(i) $\Phi^0 \subset \Phi^1$

(ii) for all $p,q \in \Phi^1$, $\sim p \in \Phi^1$ and $p \wedge q \in \Phi^1$

(iii) for all $x,y \in A[\Phi^0, C]$, $i \in I$, $x \succsim_i y \in \Phi^1$

Now first order propositions about knowledge can be defined in terms of the preference propositions. As noted in the previous section, there are different natural ways of defining knowledge in terms of preferences, but one natural one is the following. *Define* the proposition "decision maker i knows p" (or $k_i p$) as being the conjunction of every proposition of the form "decision maker i (weakly) prefers act x to act y" where x and y do not differ at those states where p is true. Formally,

$$k_i p = \bigwedge_{\{x,y \in A(\Phi^0, C) \mid x(\omega) = y(\omega),\ \forall\ \omega \text{ such that } \omega(p)=1\}} x \succsim_i y$$

Thus restrictions on propositions about preference implicitly place restrictions on statements about knowledge. Notice also that by construction of Φ^1, every proposition of the form $k_i p$, where $p \in \Phi^0$, is contained in Φ^1.

Now I have expanded the set of propositions to include first order statements about preference, and thus first order statements about knowledge. But the set of relevant propositions is not exhausted when I expand the set of propositions to include first order statements about knowledge. Second order statements, like "decision maker i knows that the decision maker j knows that p", are relevant too. Thus higher order propositions about preferences must be included. These will consists of propositions of the form "act x is preferred by i to act y" where acts x and y themselves depend on propositions about preference. Formally, I can repeat the expansion of the set of propositions I carried out, iteratively. Thus, for each integer $k \geq 1$, let Φ^k be the intersection of all sets which satisfy the following three conditions:-

(i) $\Phi^{k-1} \subset \Phi^k$

(ii) for all $p,q \in \Phi^k$, $\sim p \in \Phi^k$ and $p \wedge q \in \Phi^k$

(iii) for all $x,y \in A[\Phi^{k-1}, C]$, $i \in I$, $x \succeq_i y \in \Phi^k$

Now I have a well-defined hierarchical description of propositional spaces. But it should be clear that I am including far too many propositions about preferences at each level. That is, the state space generated by logically consistent truth valuations of the level k set of propositions, $t^*(\Phi^k)$, will include many states which are nonetheless inconsistent with our notion of preference. It is at this point that standard restrictions on preferences (reflexivity, transitivity, non-triviality), together with a valuable information type condition as the relevant coherence requirement, are imposed throughout the hierarchy.

Once a complete description of the construction of the hierarchy is completed, I can examine the properties of the hierarchy. The belief hierarchies discussed earlier had an attractive *limit-closure* property. The infinite hierarchy of beliefs exhausts each decision maker's beliefs in the sense that any one decision maker's infinite sequence of beliefs (at different levels) uniquely determines his beliefs over the whole infinite sequence of beliefs of other decision makers. The knowledge hierarchy of Fagin, Halpern and Vardi fails to satisfy the analogous limit-closure condition. Fagin, Geanakoplos, Halpern and Vardi [1992] identify general conditions under which such hierarchies satisfy such a condition. An important question would be the interpretation of those conditions in the context of the preferences hierarchy.

There is a larger challenge once the hierarchy is constructed and its properties explored. The hierarchy avoids the circularity in the kind of models in the previous section relating beliefs, knowledge and preference. Thus it provides an additional justification for such an approach. The ultimate purpose, then, is to develop the methods of modelling revision of beliefs and knowledge into a usable form for economists explaining the kind of real learning that takes place in economic environment, as opposed to the trivial learning involved in revising partitions and updating by Bayes rule.

REFERENCES

Aumann, R. [1976]. "Agreeing to Disagree," *The Annals of Statistics*, 4:6, 1236-1239.

Aumann, R. [1989]. "Notes on Interactive Epistemology", Yale University.

Binmore, K. and A. Brandenburger [1990]. "Common Knowledge and Game Theory," in Binmore [1990], *Essays on the Foundations of Game Theory*. Oxford: Basil Blackwell.

Blackwell, D. [1951]. "The Comparison of Experiments" in *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, 93-102. University of California Press.

Böge, W. and T. Eisele [1979]. "On Solutions of Bayesian Games," *International Journal of Game Theory*, **8**, 193-215.

Brandenburger, A. and E. Dekel [1993]. "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, **59:1**, 189-198.

Brandenburger, A., E. Dekel and J. Geanakoplos [1992]. "Correlated Equilibrium with Generalized Information Structures," *Games and Economic Behavior*, **4**, 182-201.

Fagin, R., J. Halpern and M. Vardi [1991]. "A Model-Theoretic Analysis of Knowledge," *Journal of the Association for Computing Machinery*, **38**:2, 382-428.

Fagin, R., J. Geanakoplos, J. Halpern and M. Vardi [1992]. "The Expressive Power of the Hierarchical Approach to Modeling Knowledge and Common Knowledge," IBM Almaden Research Center.

Geanakoplos, J. [1989]. "Game Theory without Partitions, and Applications to Speculation and Consensus," Cowles Foundation Discussion Paper #914, Yale University.

Geanakoplos, J. [1992]. "Common Knowledge," in *4th Conference on Theoretical Aspects of Reasoning about Knowledge*, edited by Y. Moses. San Mateo, California: Morgan Kaufman.

Gul, F. [1992]. "Savage's Theorem with a Finite Number of States," *Journal of Economic Theory*, **57:1**, 99-110.

Kripke, S. [1963]. "Semantic analysis of modal logic," *Z. Math. Logik Grundlag. der Math.*, **9**, 67-96.

Machina, M. [1989]. "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, 27, 1622-1668.

Mertens, J.-F. and S. Zamir [1985]. "Formulation of Bayesian analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14:1, 1-29.

Milgrom, P. and N. Stokey. [1982]. "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26:1, 17-27.

Modica, S. and A. Rustichini [1993]. "Unawareness: a Formal Theory of Unforeseen Contingencies."

Morris, S. [1991]. "The Role of Beliefs in Economic Theory." Ph.D. dissertation, Yale University.

Morris, S. [1992a]. "Dynamic Consistency and the Value of Information," CARESS Working Paper #92-17, University of Pennsylvania.

Morris, S. [1992b]. "Revising Knowledge: A Decision Theoretic Approach," CARESS Working Paper #92-27, University of Pennsylvania.

Morris, S. and H. Shin [1993]. "Noisy Bayes Updating and the Value of Information," CARESS Working Paper #93-02, University of Pennsylvania.

Ramsey, F. [1926]. "Truth and Probability," in *The Foundations of Mathematics and other Logical Essays*. London: Kegan Paul [1930].

Rubinstein, A. and A. Wolinsky [1990]. "On the logic of "Agreeing to Disagree" type results," *Journal of Economic Theory*, 52:1, 190-207.

Samet, D. [1990]. "Ignoring Ignorance and Agreeing to Disagree," *Journal of Economic Theory*, 52:1, 190-207.

Savage, L. [1954]. *The Foundations of Statistics*. John Wiley and Sons.

Shin, H. [1992]. "Logical Structure of Common Knowledge," forthcoming in *Journal of Economic Theory*.