

Making Decisions in Large Worlds

Ken Binmore
Economics Department
University College London
Gower Street
London WC1E 6BT, UK

Rational Decisions in Large Worlds¹

by Ken Binmore

The look before you leap principle is preposterous if carried to extremes. . .

Leonard Savage's *Foundations of Statistics*

1 Bayesianism

This paper argues that we need to look beyond Bayesian decision theory for an answer to the general problem of making rational decisions under uncertainty.

The view that Bayesian decision theory is only genuinely valid in a *small world* was asserted very firmly by Leonard Savage [18] when laying down the principles of the theory in his path-breaking *Foundations of Statistics*. He makes the distinction between small and large worlds in a folksy way by quoting the proverbs "Look before you leap" and "Cross that bridge when you come to it". You are in a small world if it is feasible always to look before you leap. You are in a large world if there are some bridges that you cannot cross before you come to them.

As Savage comments, when proverbs conflict, it is proverbially true that there is some truth in both—that they apply in different contexts. He then argues that some decision situations are best modeled in terms of a small world, but others are not. He explicitly rejects the idea that all worlds can be treated as small as both "ridiculous" and "preposterous".

The first half of his book is then devoted to a very successful development of the set of ideas now known as Bayesian decision theory for use in small worlds. The second half of the book is an attempt to develop a quite different set of ideas for use in large worlds, but this part of the book is usually said to be a failure by those who are aware of its existence.²

Frank Knight [15] draws a similar distinction between making decision under risk or uncertainty.³ The pioneering work of Gilboa and Schmeidler [7] on making

¹This paper is the text of a lecture given at the 2006 ADRES conference in Marseilles. It develops ideas from Binmore [3]. There will eventually be a book with the same title.

²The same is true if one regards Savage's discussion of "microcosms" as an attempt to reduce any given large world to a small world.

³A decision is risky if probabilities can be reliably assigned to relevant events. Otherwise we are faced with a problem in (Knightian) uncertainty. Behavioral economists speak of *ambiguity* rather than uncertainty, but I see this as a retrograde step, since it suggests that the difficulty with uncertain situations is only that the decision-maker may be in two minds about what probability to assign to an uncertain event (Camerer *et al* [5, 10]). But are we so sure that it always makes sense to assign subjective probabilities to all uncertain events?

case-based decisions in what Savage would call a large world has been particularly influential in advancing Knight's research program. Nevertheless, Bayesianism remains the prevailing orthodoxy amongst the vast majority of economists.

Bayesianism is understood in this paper to be the philosophical principle that Bayesian methods are always appropriate in all decision problems, regardless of whether the relevant set of states in the relevant world is large or small. For example, the world in which financial economics is set is obviously large in Savage's sense, but the suggestion that there might be something questionable about the standard use of Bayesian updating in financial models is commonly greeted with incredulity or laughter.

Someone who acts as if Bayesianism were correct will be said to be a Bayesianite. It is important to distinguish a Bayesian like myself—someone convinced by Savage's arguments that Bayesian decision theory makes sense in small worlds—from a Bayesianite. In particular, a Bayesian need not join the more extreme Bayesianites in proceeding as though:

- All worlds are small.
- Rationality endows agents with prior probabilities.
- Rational learning consists simply in using Bayes' rule to convert a set of prior probabilities into posterior probabilities after registering some new data.

Bayesianites are often understandably reluctant to make an explicit commitment to these principles when they are stated so baldly, because it then becomes evident that they are implicitly claiming that David Hume [11] was wrong to argue that the principle of scientific induction cannot be justified by rational argument. However, I do not think I am merely attacking a straw man. What matters for this purpose is not so much what people say about their philosophical attitudes, but what models they choose to construct. As it says in the book of Matthew: By their fruits shall ye know them.

For example, recent winners of the Nobel prize for economics include John Harsanyi and Robert Aumann, both of whom must surely be counted as Bayesianites. One cannot but admire Aumann's [2] vision in modeling the entire universe and everything in it as a small world to which Bayesian decision theory applies, but nor can one escape the conclusion that such a vision excludes the possibility that any lesser world can be large.

2 Using Savage's theory

One may distinguish between objective, subjective, and logical probabilities. Objective probabilities are determined by long-run frequencies. The probabilities assigned to the numbers on a roulette wheel are a standard example. Subjective probabilities are usually understood as being embodied in a person's choice behavior. For example, the odds at which Pandora will bet on different horses at the race track

can be thought of as revealing information about her subjective probabilities that various horses will win.⁴ An adequate theory of logical probability would solve the age-old problem of scientific induction. Is the universe infinite? Does the Higgs boson exist? When is the next earthquake? Does my boyfriend really love me? Just put the evidence in a computer programmed with the theory, and out will come the appropriate probability.

Bayesianites believe that the subjective probabilities of Bayesian decision theory can be reinterpreted as logical probabilities without any hassle. Its adherents therefore hold that Bayes' rule is the solution to the problem of scientific induction. No support for such a view is to be found in Savage's theory—nor in the earlier theories of Ramsey Ramsey³¹, de Finetti [6], or Von Neumann and Morgenstern [21]. Savage's theory is entirely and exclusively a consistency theory. It says nothing about how decision-makers come to have the beliefs ascribed to them; it asserts only that, if the decisions taken are consistent (in a sense made precise by a list of axioms), then they act as though maximizing expected utility relative to a subjective probability distribution.

Objections to Savage's axiom system can be made, although it is no objection when discussing rational behavior to argue, along with Allais [1] and numerous others, that real people often contravene the axioms. People also often get their sums wrong, but this is no good reason for advocating a change in the axiomatic foundations of arithmetic! However, it is not the soundness of Savage's consistency axioms that are at issue here, since concerns I am expressing would be equally valid if Savage's theory were replaced by a rival theory in which his consistency axioms were replaced by those advocated by one of his critics. What is being denied is that *any* passive descriptive theory can be reinterpreted as an active prescriptive theory at negligible cost.

A reasonable decision-maker will presumably wish to avoid inconsistencies. A Bayesianite therefore assumes that it is enough to assign prior beliefs to a decision-maker, and then forget the problem of where beliefs come from. Consistency then forces any new data that may appear to be incorporated into the system via Bayesian updating. That is, a posterior distribution is obtained from the prior distribution using Bayes' rule.

The naiveté of this approach doesn't consist in using Bayes' rule, whose validity as a piece of algebra isn't in question. It lies in supposing that the problem of where the priors came from can be quietly shelved.

Savage did argue that his descriptive theory of rational decision-making could be of practical assistance in helping decision-makers form their beliefs, but he didn't argue that the decision-maker's problem was simply that of selecting a prior from a limited stock of standard distributions with little or nothing in the way of soul-searching. His position was rather that one comes to a decision problem with a whole set of subjective beliefs derived from one's previous experience that may or

⁴Dutch book arguments based on this scenario can be used instead of Savage's theory to provide a foundation for subjective probabilities.

may not be consistent.

In a famous encounter with Allais, Savage himself was trapped into expressing inconsistent beliefs about a set of simple decision problems. The response he made is very instructive. He used his theory to adjust his beliefs until they became consistent. Luce and Raiffa [16, p.302] explain the process by means of which such a consistent set of final beliefs is obtained as follows:

Once confronted with inconsistencies, one should, so the argument goes, modify one's initial decisions so as to be consistent. Let us assume that this jockeying making snap judgments, checking on their consistency, modifying them, again checking on consistency, etc leads ultimately to a bona fide, a priori distribution.

For Savage therefore, forming beliefs was more than a question of attending to gut-feelings. It was a matter for calculation—just as the question of whether you or I prefer $\$17 \times 29$ to $\$19 \times 23$ is a matter for calculation.

But why should we wish to adjust our gut-feelings using Savage's methodology? In particular, why should a rational decision-maker wish to be consistent? After all, scientists aren't consistent, on the grounds that it isn't clever to be consistently wrong. When surprised by data that shows current theories to be in error, they seek new theories that are inconsistent with the old theories. Consistency, from this point of view, is only a virtue if the possibility of being surprised can somehow be eliminated. This is the reason for distinguishing between large and small worlds. Only in the latter is consistency an unqualified virtue.

One might envisage the process by means of which a decision-maker achieves a consistent set of subjective beliefs in a small world as follows. Pandora knows that subjective judgments need to be made, but prefers to make such judgments when her information is maximal rather than minimal. She therefore asks herself, for every conceivable possible course of future events: what would my beliefs be after experiencing these events? Such an approach automatically discounts the impact that new knowledge will have on the basic model that Pandora uses in determining her beliefs—that is, it eliminates the possibility that she will feel the need to alter her basic model after being surprised by a chain of events whose implications she had not previously considered. Next comes the question: is this system of contingent beliefs consistent? If not, then the Pandora may examine the relative confidence that she has in the "snap judgments" she has made, and then adjust the corresponding beliefs until they are consistent. With Savage's definition of consistency, this is equivalent to asserting that the adjusted system of contingent beliefs can be deduced, using Bayes' rule, from a single prior probability distribution.

At the end of the story, the situation is as envisaged by Bayesianites: the final "massaged" posteriors can indeed be formally deduced from a final "massaged" prior using Bayes' rule. This conclusion is guaranteed by the use of a complex adjustment process that operates until consistency is achieved. As far as the massaged beliefs are concerned, Bayes' rule has the status of a tautology—like $2 + 2 = 4$. Together with the massaged prior, it serves essentially as an indexing system that keeps track of the library of massaged posteriors.

Notice that it isn't true in this story that Pandora is *learning* when the massaged prior is updated to yield a massaged prior—that rational learning consists of no more than applying Bayes rule. On the contrary, Bayesian updating only takes place after all learning is over. The actual learning takes place while Pandora is taking into account the effect that possible future surprises may have on the basic model that she uses to construct her beliefs, and continues as she refines her beliefs during the massaging process. Bayesianites therefore have the cart before the horse. Insofar as learning consists of deducing one set of beliefs from another, it is the massaged prior that is deduced from the unmassaged posteriors.

A caveat is necessary before proceeding. When the word "learning" is used in the preceding paragraph, it is intended in the sense of "adding to one's understanding" rather than simply "observing what happens". Obviously, if Pandora has perfect recall, she will have more facts at her disposal at later times than at earlier times, and it is certainly true that there is a colloquial sense in which she can be said to "learn" these facts as time goes by. However, it seems to me that this colloquial usage takes for granted that Pandora is also sorting and classifying the facts she learns into some sort of orderly system with a view to possibly making use of her knowledge in the future. Otherwise it wouldn't seem absurd to say that a camera is "learning" the images it records.

In any case, it is not the simple recording of facts that is intended when I discuss "Bayesian learning". Any proposal for a rational learning scheme will presumably include recording the facts (if the cost of so doing is negligible). What distinguishes "Bayesian learning" from its alternatives must therefore be something else.

In spite of this caveat about what I intend when speaking of learning, the suggestion that Bayesian updating in a small world involves no genuine learning at all commonly provokes expressions of incredulity. Is it being said that we can only learn when deliberating about the future, and never directly from experience? The brief answer is no, but I have learned directly from experience that a longer answer is necessary.

In the first place, the manner in which you and I (and off-duty Bayesianites) learn things about the real world is not necessarily relevant to the way a Bayesian learns. Still less is it relevant to the way in which a Bayesianite learns when on duty. Experimental data offers very little evidence in favor of the proposition that we are natural Bayesians of any kind. Indeed, what evidence there is seems to suggest that, without training, even clever people are quite remarkably inept in dealing with simple statistical problems involving conditional probabilities. In my own game theory experiments, no subject has ever given a Bayesian answer to the question "Why did you do what you did?" when surveyed after the experiment even though, in most cases, the populations from which the subjects were drawn consisted entirely of students who had received training in Bayesian statistics. I therefore think introspection is unlikely to be a reliable guide when considering what learning for a Bayesian may or may not be.

The fact that real people actually learn from experience is therefore not relevant to whether Bayesian updating in a small world should count as genuine learning.

The worlds about which real people learn are almost always large and, even when they are confronted with a small world, they almost never use Bayesian updating.

Bayesian statisticians are an exception to this generalization. They use Bayesian updating all the time, but, just like real people, they are almost never working in a small world. That is to say, they have not asked themselves why a knee-jerk adherence to consistency requirements is appropriate, but simply update from a prior distribution chosen on the basis of what their past experience has shown to work out well in analogous situations. I do not argue that such a procedure is necessarily nonsensical. On the contrary, it often leads to descriptions of the data that provide much insight. Nor do I argue that a Bayesian statistician who updates from a prior distribution chosen on a priori grounds isn't learning. All I have to say to such a Bayesian statistician is that I see no grounds for him to claim that he is learning optimally, or that his methodology is necessarily superior to those of classical statistics.

The problem of how "best" to learn in large world is unsolved. Gilboa and Schmeidler [7], for example, make no such claims for their approach. Probably it is one of those problems that has no definitive solution. But, unless the problem of scientific induction is solved, any learning procedures that we employ in the context of large world will necessarily remain arbitrary to some extent.

However, we are still not through with the question of whether Bayesian updating in a small world can properly count as learning. So far, it has been argued that the fact that real people clearly learn from experience is irrelevant to this question. The same is true of Bayesian statisticians operating in a large world, (or a potentially small world that they have failed to close). This leaves us free to focus on what is genuinely at issue. For this purpose, I want to draw an analogy between how a Bayesian using the massaging methodology I have attributed to Savage learns, and how a child learns arithmetic. It is true that the Bayesian is envisaged as teaching himself, but I do not think this invalidates the comparison.

When Alice learns arithmetic at school, her teacher does not know what computations life will call upon her to make. Amongst other things, he therefore teaches her an algorithm for adding numbers. This algorithm requires that Alice memorize some addition tables. In particular, she must memorize the answer to $2 + 3 = ?$. If the teacher is good at his job, he will explain why $2 + 3 = 5$. If Alice is an apt pupil, she will understand his explanation. One may then reasonably say that Alice has learned that, should she ever need to compute $2 + 3$, then the answer will be 5.

Now imagine Alice in her maturity trying to complete an income tax form. In filling the form, she finds herself faced with the problem of computing $2 + 3$, and so she writes down the answer 5. Did she just learn that the answer to this problem is 5? Obviously not. She learned this in school. All that one can reasonably say that she "learned" in filling the form is that filling the form requires computing $2 + 3$. But such simple registering of undigested facts is excluded by the caveat that identifies learning with "augmenting one's understanding". Of course, there may be children who are such poor students that they grow to maturity without learning their addition tables. Such a person might perhaps use her fingers to reckon with

and thereby discover or rediscover that $2 + 3 = 5$ while filling the tax form. She would then undoubtedly have learned something. But she would not be operating in a small world, within which all potential surprises have been predicted and evaluated in advance of their occurrence.

How is it that Bayesianites succeed in convincing themselves that rational learning consists of no more than the trivial algebraic manipulations required for the use of Bayes' rule? My guess is that their blindness is only a symptom of a more serious disease that manifests itself as a worship of mathematical formalism. A definition-axiom-theorem-proof format is designed to close the mind to irrelevant distractions. But the aspects of the learning process that are neglected by Savage's formalism are not irrelevant. How decision-makers form and refine their subjective judgments really does matter. But the fact that Savage's theory leaves these aspects of the learning process utterly unmodeled creates a trap into which Bayesianites readily fall. The trap is to proceed as though anything that is not expressed in the formalism to which one is accustomed does not exist at all.

In game theory, however, the question of where beliefs come from cannot sensibly be ignored. Bayesianite decision theory provides an adequate account of why we should study equilibria, but fails to make any headway at all with the problem of equilibrium selection. Game theorists therefore cannot afford to fall victim to Bayesianite newspeak⁵ if they hope to break out of the bridgehead they currently occupy.

3 Bayesianism in Game Theory

This section looks very briefly at two approaches to the problem of founding game theory on Bayesian principles. The second approach, due to Robert Aumann citeAumann87b, hangs together very much better than the first. But this is because Aumann's approach doesn't attempt to do more than justify game theorists' obsession with the notion of an equilibrium. However, the first approach aims to say things about which equilibrium should be selected.

Harsanyi and Selten's [19] theory is perhaps the best known of the avowedly Bayesian approaches to the problem of equilibrium selection. However, it is too baroque a theory to lend itself to easy discussion in a paper like this.

The notion of a tracing procedure lies at the heart of their model. Their procedure seeks to trace the manner in which Bayesian players will reason their way to an equilibrium. Other authors offer alternative accounts of how such reasoning might proceed. Skyrms [20] gives a particularly clean description of how he sees the deliberative process operating inside the head of a Bayesian player.

Skyrms [20] follows Harsanyi and Selten and others in supposing that, while deliberating, the players assign interim subjective probabilities to the actions available

⁵Newspeak is an invented language in Orwell's *1984* in which it is impossible to make politically incorrect statements.

to their opponents. If these subjective probabilities are common knowledge, along with the fact that everyone is a maximizer of expected utility, then an inconsistency will arise—unless the players' beliefs happen to be in equilibrium. When such an inconsistency arises, the players are assumed to update their subjective probabilities using Bayes' rule.

Various candidates for the likelihood function can be considered (of which Skyrms offers a small sample). However, the modeling judgment made at this level is irrelevant to the point I want to make. My criticism of this and similar models will be clear. By hypothesis, the players haven't looked ahead to preview all possible lines of reasoning they might find themselves following in the future. They are therefore operating in a world that is definitely large. In such world, no special justification for the use of Bayesian updating exists. One might seek to rescue the special status of Bayesian updating by departing from Skyrms' story and postulating that the players have indeed previewed all the possible lines of reasoning open to them. But, after the previewing is over, there would then be no scope for Bayesian updating, because no there would be no new information to incorporate into the system when the player began to reason for real. In summary, one might say that the conditions that justify the use of Bayes' rule in this context are satisfied if and only if there is nothing for Bayes' rule to update.

One cannot make the same criticism of a well-known paper by Kalai and Lehrer [13]. They envisage a game being played repeatedly in real time. The circumstances under which the repetition takes place need not concern us. For our purposes, it is enough that the players use Bayes' rule to update their beliefs after each repetition, and that Kalai and Lehrer give conditions under which there is convergence on a Nash equilibrium. What does such a conclusion mean?

It is certainly a very reassuring consistency result for those like myself who regard Nash equilibrium as the basic tool of game theory. But is the result also a contribution to equilibrium selection theory? It is true, as Kalai and Lehrer remark, that the limit equilibrium is a function of the players' prior beliefs, but it seems to me that much care is necessary in interpreting this piece of mathematics. If we take seriously the notion that a players' prior beliefs are simply a summary of a set of massaged posterior beliefs, we have to abandon the idea that the players in Kalai and Lehrer's model are learning which equilibrium to play as the future unfolds. The players' already knew what equilibrium would be played under all possible future contingencies. Their initial snap judgments necessarily incorporate preconceptions on this subject that the model leaves unexplained. Any learning took place during the unmodeled introspection period before the play of the game when the players previewed all possible courses the game might take and predicted how the game would end up being played after each of these possible sets of common experience.

It should be emphasized that the last thing I wish to do is to criticize anyone for seeking to model the process by means of which equilibrium is achieved. Far from decrying such work, I believe that the reason game theorists have made so little progress with the equilibrium selection problem is because of their reluctance to confront such questions. I do not even object to Bayesian updating being used as

a learning rule in this context, provided that nobody is claiming any special status for it, beyond the fact that it possesses some pleasant mathematical properties. However, other learning rules also have virtues, and the decision to use Bayes' rule in the context of a small world is no less *ad hoc* than the decision to use one of the rival rules. My own preferred research strategy on this subject is not to make any a priori choice at all of a learning rule, but to let one emerge endogenously as a consequence of the operation of evolutionary pressures. However, this is an approach fraught with many difficulties.

Aumann's [2] attempt to provide Bayesian foundations for game theory is very different in character from the work discussed so far in this section. Nobody learns anything or even decides anything in his very static model. Things are "just the way they are", and we are offered the role of a passive observer who sits on the sidelines soliloquizing on the nature of things. Such a model is not well-adapted to the equilibrium selection problem. Its purpose is to clarify what kinds of equilibria should lie in the set from which a selection needs to be made.

Aumann postulates a universe of discourse whose states are all-inclusive. A description of such a state includes not only what players know and believe about the world and the knowledge and beliefs of other players, but also what all the players will do in that state. In such a framework, it becomes almost tautological that players whom fate has decreed will be Bayesian-rational in every state will necessarily operate some kind of equilibrium. Aumann then notes that, if what the players know always includes what strategy they find themselves using, then they will necessarily be frozen into what he calls a "subjective correlated equilibrium".

This brief summary is adequate only to make it clear that Aumann's world is not small. On the contrary, his world is as large as a world could possibly be, since its states encompass everything that might conceivably matter. However, Aumann evades the traps that await more naive Bayesianites by refusing to classify his theory either as descriptive or as prescriptive. He describes his model as "analytic" to indicate that all the action takes place in the head of an otherwise passive observer.

Aumann's model certainly cannot be prescriptive because there is no point in offering advice to players who "just do what they happen to do" and "believe what they happen to believe". Nor can the model be descriptive of a world in which people actively make conscious choices after transferring their experience into subjective judgments about the way things are. But it seems to me that the latter is precisely the kind of world with which game theory needs to grapple. I want to argue now that such a world is necessarily large in Savage's sense. The case for this is even stronger than the standard claim that the world within which physics is discussed is large. Or, to say the same thing more flamboyantly, inner space is necessarily even more mysterious than outer space. The reason is that, if the thinking processes of a player are to be modeled, then we are no longer free to envisage that all possible mental processes have been completed. A player cannot exhaustively evaluate all contingencies in a universe that includes his own internal deliberations and those of other players like himself. The issue is more fundamental than whether Bayesianism is applicable or not, since one cannot even rely on the epistemology that Bayesianites

take for granted.

4 Epistemology

Bayesian decision theory assumes that a decision-maker's knowledge at any time partitions the universe of discourse into a collection of disjoint possibility sets. The partitioning property of these possibility sets is then inherited by the information sets that Von Neumann introduced into game theory (Binmore [4, p.454]). An assumption of "perfect recall" is then usually made to ensure that a player's knowledge partition is always a refinement of his previous partitions. Such knowledge partitions can be formulated in terms of the idea of a knowledge operator.

We identify an event E with a subset of a given universe of discourse denoted by Ω . The event in which Pandora knows that E has occurred is denoted by $\mathcal{K}E$, where \mathcal{K} is her knowledge operator. The event in which Pandora thinks it possible that E has occurred is denoted by $\mathcal{P}E$, where \mathcal{P} is her possibility operator.

If we make the identification $\mathcal{P} = \sim\mathcal{K}\sim$, then we establish a duality between \mathcal{K} and \mathcal{P} . Either of the following lists will then serve as a rendering of the requirements of what philosophers call the modal logic S-5:

$$\begin{array}{ll}
 (\text{K0}) & \mathcal{K}\Omega = \Omega \\
 (\text{K1}) & \mathcal{K}(E \cap F) = \mathcal{K}E \cap \mathcal{K}F \\
 (\text{K2}) & \mathcal{K}E \subseteq E \\
 (\text{K3}) & \mathcal{K}E \subseteq \mathcal{K}^2E \\
 (\text{K4}) & \mathcal{P}E \subseteq \mathcal{K}\mathcal{P}E \\
 (\text{P0}) & \mathcal{P}\emptyset = \emptyset \\
 (\text{P1}) & \mathcal{P}(E \cup F) = \mathcal{P}E \cup \mathcal{P}F \\
 (\text{P2}) & \mathcal{P}E \supseteq E \\
 (\text{P3}) & \mathcal{P}E \supseteq \mathcal{P}^2E \\
 (\text{P4}) & \mathcal{K}E \supseteq \mathcal{P}\mathcal{K}E
 \end{array}$$

Axioms (K2) or (P2) are infallibility or "consistency" axioms—they say that what Pandora knows is actually true. The seemingly innocent (K0) and (P0) will be called "completeness" axioms.

These ideas are linked with knowledge partitions by defining the possibility set $P(\omega)$ to be the set of states that Pandora thinks is possible when the true state of the world is ω . In Bayesian decision theory, the minimal requirements for such possibility sets are usually taken to be:

$$\begin{array}{ll}
 (\text{Q0}) & \text{The collection } \{P(\omega) : \omega \in \Omega\} \text{ partitions } \Omega \\
 (\text{Q1}) & \omega \in P(\omega)
 \end{array}$$

The second of these assumptions is the infallibility requirement.

To establish an equivalence between the two approaches, it is only necessary to define \mathcal{P} and P in terms of each other using the formula:

$$\omega \in \mathcal{P}E \Leftrightarrow P(\omega) \cap E \neq \emptyset. \quad (1)$$

With this definition, (P0)–(P3) can be deduced from (Q1)–(Q2) and vice-versa. However, the role of the completeness axiom (P0) is peripheral. If we dispense with

(P0) and redefine both (P1)–(P3) and (1) so that they apply only to non-empty events, then (new P1)–(new P3) are equivalent to (Q1)–(Q2).

It is significant that (P0) can be eliminated, because the result of the next section can be regarded as saying that (P0) and (P2) cannot both hold in a large enough world when the possibility operator is algorithmic.

How does does Pandora decide what she knows or thinks is possible? The formalism we have been discussing packs this question away inside a black box, which we now unpack by postulating that Pandora is equipped with a ‘Leibniz engine’ or possibility machine J that makes judgments on what events are possible.

The reference to a Leibniz engine is intended to signal that justification is to be understood as algorithmic. The possibility machine J is taken to be a Turing machine that sometimes answers NO when asked questions that begin:

Is it possible that ...?

Issues of timing are obviously relevant here. How long does one wait for an answer before acting? Such timing problems are idealized away by assuming that Pandora is able to wait any finite number of periods for an answer.

As in the Turing halting problem, we suppose that $[N]$ is some question about the Turing machine N . We then take $\{M\}$ to be the question:

Is it possible that M answers NO to $[M]$?

Let T be the Turing machine that outputs $[x]$ on receiving the input $\{x\}$, and employ the Turing machine $I = JT$ that first runs an input through T , and then runs the output of T through the justification machine J . Then the Turing machine I responds to $[M]$ as J responds to $\{M\}$.

An event E is now defined to be the set of states in which I responds to $[I]$ with NO. We then have the following equivalences:

$$\begin{aligned} \omega \in E &\Leftrightarrow I \text{ responds to } [I] \text{ with NO} \\ &\Leftrightarrow J \text{ reports it to be impossible that } I \text{ responds to } [I] \text{ with NO} \\ &\Leftrightarrow \omega \notin \sim PE \end{aligned}$$

It follows from (P2) that

$$\sim PE = E \subseteq PE.$$

This identity only holds when $PE = \Omega$. Since $E = \sim PE$, it follows that $E = \emptyset$, and so $P\emptyset = \Omega$. That is to say, we are led to the following apparent contradiction:

Proposition. If the states in Ω are sufficiently widely construed and knowledge is algorithmic, then infallibility implies that the decision-maker always thinks it possible that nothing will happen.

I think the apparent contradiction built into the preceding proposition signals a failure of the model to capture the extent to which familiar assumptions from small-world decision theories need to be modified when moving to a large world.

For example, we think of an event E as having occurred if the true state ω of the world has whatever property defines E . But how do we determine whether ω has this property?

If we are committed to an algorithmic approach, we need an algorithmic procedure for the defining property P of each event E . This procedure can then be used to interrogate ω with a view to getting a YES or NO answer to the question:

Does ω have property P ?

We can then say that E has occurred when we get the answer YES, and that $\sim E$ has occurred when we get the answer NO.

But in a sufficiently large world, there will necessarily be properties for which the relevant algorithm sometimes will not halt. Our methodology will then classify ω as belonging neither to E nor to $\sim E$. If we insist on retaining a formalism that only make proper sense in a small world, we are then forced to place ω in \emptyset —but then the symbol \emptyset can no longer denote the set with no elements in the usual way. To write $P\emptyset = \Omega$ would then need to be interpreted as meaning that Pandora necessarily thinks it possible that the true state of the world will remain unclassified according to any of the properties recognized by her algorithmic classification system.

5 A Minimal Extension

In this section, I describe the minimal extension of Bayesian decision theory that seems compatible with the requirement that a decision-maker be algorithmic.

Following Von Neumann and Morgenstern [21], we first simplify the problem to be considered by allowing only two consequences, winning and losing. Which of these will occur depends on some process about which Pandora is only partially informed. A Turing machine $M = M(\omega)$ will be used to model the manner in which the Pandora evaluates her partial information. The input to M is therefore the data $D = D(\omega)$ about the unknown process available to Pandora.

Since our aim is to find a minimal extension to Bayesian decision theory, let us restrict the machine M to answering either YES or NO when asked probabilistic questions like:

Is $\text{prob}(E) \leq x$?

As in the previous section, there will be occasions when M will calculate for ever when offered such a question.

This approach leads immediately to the notion of “upper and lower probabilities” with which many decision theorists have toyed. We take the upper probability $\bar{\pi}$ of an event E to be infimum of all x to which M answers the question “Is $\text{prob}(E) \leq x$?” in the affirmative. We take $\underline{\pi}$ to be the supremum of all x to which M answers the question “Is $\text{prob}(E) \geq x$?” in the affirmative. With some mild assumptions about the consistency of replies made by the machine M , we are then led to the idea that the probabilistic properties of any event E can be summarized by associating

it with an interval $[\underline{\pi}, \bar{\pi}]$, in which $\underline{\pi}$ is the lower probability of E and $\bar{\pi}$ is its upper probability.⁶

Measurable events E can be defined to be those for which $\underline{\pi} = \bar{\pi}$. The measure or probability $\pi = \text{prob}(E)$ of such a measurable event E is simply $\pi = \underline{\pi} = \bar{\pi}$. The events E for which $\underline{\pi} < \bar{\pi}$ are nonmeasurable. The inner and outer measure of a nonmeasurable set can be defined in the usual way. We should presumably insist that the inner measure is no larger than $\underline{\pi}$, and the outer measure is no smaller than $\bar{\pi}$. For a *minimal* extension, I believe we should make $\underline{\pi}$ equal to the inner measure of E and $\bar{\pi}$ equal to the outer measure of E .

Suppose that Pandora has preferences over lotteries in which the prizes are themselves gambles in which she wins or loses according to whether a nonmeasurable event obtains or not. If these preferences satisfy the Von Neumann and Morgenstern rationality postulates, then we can introduce the Von Neumann and Morgenstern utility $u[\underline{\pi}, \bar{\pi}]$ of a nonmeasurable event. One will presumably wish to insist that

$$u[\underline{\pi}, \underline{\pi}] \leq u[\underline{\pi}, \bar{\pi}] \leq u[\bar{\pi}, \bar{\pi}]$$

and to normalize so that $u[\pi, \pi] = \pi$, but it is not clear what further rationality requirements are appropriate. However, if Pandora evaluates $\underline{\pi}$ and $\bar{\pi}$ “separately”, then an argument of Keeney and Raiffa [14] shows that $u[\underline{\pi}, \bar{\pi}]$ must take one of the two forms:

$$u(\underline{\pi}) + u(\bar{\pi}) \quad \text{or} \quad u(\underline{\pi}) \times u(\bar{\pi}). \quad (2)$$

John Milnor [17, 16] gives a list of axioms for “decision making under complete ignorance” that lead to a decision criterion originally proposed by Leo Hurwicz [12]. In the current context, the Hurwicz criterion takes the form

$$u[\underline{\pi}, \bar{\pi}] = h\underline{\pi} + (1 - h)\bar{\pi}, \quad (3)$$

and hence provides a simple instance of the first of the two functional forms (2).

6 Ellsberg Paradox

The Ellsberg paradox will serve as an example of how the Hurwicz criterion works. Recall that an urn contains 300 billiard balls of which 100 are known to be red.

⁶One way that Bayesianites proceed is to insist that their critics compare bets on events for which the appropriate probabilities are uncontroversial with bets on events to which the critics are reluctant to assign subjective probabilities. In the case of Pandora using the machine M , they would therefore seek two events between which the decision-maker is indifferent: a measurable event for which M yields a probability p and a nonmeasurable event for which M yields $[\underline{\pi}, \bar{\pi}]$. However, although Pandora may express such an indifference, it doesn't follow that she is saying that she regards the two events as interchangeable. She will be expressing a preference not a belief. It is true that, with Savage's consistency axioms, these ideas merge (when the underlying processes are understood to be independent). But it isn't reasonable to insist on such strong consistency requirements in a large world.

The remaining 200 balls are black and white, but nothing is known about how many are black and how many are white.

Pandora is offered a choice between the gamble **J** in which she wins \$1m if and only if a red ball is drawn from the urn, and a gamble **K** in which she wins \$1m if and only if a black ball is drawn from the urn. She is also offered a choice between the gamble **L** in which she wins \$1m if and only if a red ball is not drawn from the urn, and a gamble **M** in which she wins \$1m if and only if a black ball is not drawn from the urn.

In real life, people mostly strictly prefer **J** to **K**, and **L** to **M**. However, if a Bayesian's subjective probabilities for red, black, and white are $r = \frac{1}{3}$, b , and w , these choices imply that $b < \frac{1}{3}$ and $w < \frac{1}{3}$. But then we have the inconsistency $r+b+w < 1$. It is commonly argued that real people therefore depart from Bayesian decision theory by revealing an aversion to uncertainty (or ambiguity).

However, it isn't clear that real people see the situation as the kind of small world that is in the mind of the experimenter. Everybody knows that psychological experimenters are sneaky, and perhaps they have used their knowledge of how people choose in such experiments to fix the proportion of black and white balls to the disadvantage of the subjects.⁷ If Pandora reasons in this way, she will be faced with a large world, because it contains the mind of the experimenter.

Applying the Hurwicz criterion to Pandora's comparison of **J** and **K**, we find

$$\frac{1}{3} = u(\mathbf{J}) > u(\mathbf{K}) = h \cdot 0 + (1-h) \frac{2}{3} \Leftrightarrow h > \frac{1}{2}.$$

Similarly, $u(\mathbf{L}) > u(\mathbf{M}) \Leftrightarrow h > \frac{1}{2}$. Thus, rather than being regarded as an irrational phenomenon in a small world, uncertainty aversion can be treated as a rational analogue of risk aversion that becomes relevant in a large world. For this reason, one can think of the Hurwicz coefficient h ($0 \leq h \leq 1$) as a measure of Pandora's aversion to uncertainty.

7 Battle of the Sexes

Von Neumann restricted his attention to two-person, zero-sum games because he was interested in identifying a unique rational solution of a game. It was presumably for this reason that he was dismissive when Nash showed him his proof that Nash equilibria always exist in finite games. My guess is that Von Neumann thought he could have proved the result equally well himself, but that there was no point in doing so without proposing a solution to the equilibrium selection problem.

The Battle of the Sexes poses the equilibrium selection problem in an aggravated form. The game is played by a honeymoon couple, Adam and Eve, who didn't agree at breakfast whether to go to a boxing match or the ballet in the evening. During the day they are separated in the crowds, and so must make their decisions of where

⁷Recall Newcomb's so-called paradox, in which a box contains a big prize if and only if the experimenter predicts that the subject will not choose it.

to go in the evening independently.

	<i>box</i>	<i>ball</i>
<i>box</i>	1	0
<i>ball</i>	0	2
	0	1

Battle of the Sexes

The payoffs in this game are to be strictly interpreted as Von Neumann and Morgenstern utilities. For example, we could divide each payoff in the table by 2 and interpret the resulting numbers as probabilities with which the relevant player will win some fixed prize.

The Battle of the Sexes has two Nash equilibria in pure strategies: (box, box) and $(ball, ball)$. But in the absence of any pre-play opportunity to break symmetries, neither is a candidate for the unique rational solution of the game, because any argument in favor of one is an equally valid argument in favor of the other. There remains the mixed equilibrium in which Adam plays *ball* with probability $p = \frac{1}{3}$ and Eve plays *box* with probability $q = \frac{2}{3}$. At this equilibrium, each player receives an expected payoff of $\frac{2}{3}$.

However, there is an objection to identifying the (skew) symmetric mixed equilibrium as the rational solution of the game. The players' security levels are both $\frac{2}{3}$ (which Adam guarantees by playing $p = \frac{2}{3}$ and Eve guarantees by playing $q = \frac{1}{3}$). So why don't the players switch from playing their supposedly rational strategies, from which they expect to get a payoff of only $\frac{2}{3}$, to their security strategies, which *guarantee* a payoff of at least $\frac{2}{3}$? Harsanyi [9] thought this argument sufficiently convincing that he proposed abandoning the requirement that the rational solution of a game be a Nash equilibrium in such situations.⁸

One way out of such dilemmas is to propose expanding the set of mixed strategies to "multiplex" strategies in much the same way that the set of pure strategies is expanded to obtain mixed strategies. To expand the set of pure strategies, one requires a randomizing device. To expand the set of mixed strategies, one requires a computer \mathcal{C} programmed to print a sequence of 1s and 0s. The computer \mathcal{C} would be regarded as an adequate substitute for a random device that generates 1 with probability π if the average value of the first n terms converged to π , but no pattern could otherwise be distinguished in its output.⁹ However, the average value of the

⁸One can respond to Harsanyi that Eve profits by sticking to her Nash equilibrium strategy if Adam switches to his security strategy—but the opposite is true in Australian Battle of the Sexes, which is obtained by reversing the signs of all the payoffs in regular Battle of the Sexes.

⁹Von Mises and Kolmogorov offer possible formal characterizations of the latter requirement.

sequence printed by \mathcal{C} doesn't converge. Its limit superior $\bar{\pi}$ is strictly larger than its limit inferior $\underline{\pi}$.

The standard methodology proposed by Bayesianites for dealing with a “weighted coin” like \mathcal{C} is to begin with a prior distribution with support $[0, 1]$ over all possible weighting probabilities of the coin, and to use Bayes' rule to update this distribution as data from \mathcal{C} is received. But this methodology won't lead to the standard result here, because the posterior distribution won't gradually concentrate around some particular probability π . The convergence properties of the posterior distribution will only allow us to exclude points outside the interval $[\underline{\pi}, \bar{\pi}]$.

Suppose that Adam and Eve both choose multiplex strategies in the Battle of the Sexes generated by independent computers, both with $\bar{\pi} = 1$ and $\underline{\pi} = 0$. If they evaluate their situation using the Hurwicz criterion (with the same value of the coefficient h), then we have a Nash equilibrium in multiplex strategies in which each player's payoff is $h0 + (1 - h)2$.

This equilibrium payoff exceeds $\frac{2}{3}$ if and only if $h < \frac{2}{3}$. That is to say, we have found a skew-symmetric Nash equilibrium in multiplex strategies that generates higher payoffs for the players than their security levels—provided that they aren't too uncertainty averse. If $h < \frac{5}{8}$, then their payoff at the equilibrium in multiplex strategies exceeds the maximum of $\frac{3}{4}$ possible with any skew-symmetric ($p = 1 - q$) pair of mixed strategies, whether or not in equilibrium.

8 Consistent Updating?

The discussion of upper and lower probabilities given above is proposed as the minimal extension of Bayesian decision theory consistent with an algorithmic decision-maker. It amounts to accepting that not all events can be measurable, and assessing nonmeasurable sets in terms of their inner and outer measures. It follows that Bayes' rule is not lost in the extension, because it remains the appropriate method of updating the probabilities of measurable sets, and hence the appropriate method of updating the inner and outer measures of nonmeasurable sets.

For example, in computing Eve's utility when she and Adam independently choose multiplex strategies in the Battle of the Sexes, her upper probability is obtained by multiplying the upper probabilities associated with the players' two independent multiplex strategies (which is the result of applying Bayes' rule in the independent case). Similarly, her lower probability is obtained by multiplying the lower probabilities associated with the players' two independent multiplex strategies.

One might think that Eve has the alternative of first using the Hurwicz criterion to calculate the utility for each of her pure strategies given Adam's choice of a multiplex strategy, and then substituting these values in the Hurwicz criterion again to calculate the utility of her own choice of a multiplex strategy. But such an alternative calculation can't be expected to yield the same answer, because we know that insisting on this kind of consistency is another way of characterizing Bayesian decision theory (Hammond [8]). Even the minimal extension of Bayesian

decision theory discussed in the latter part of this paper therefore fails the most elementary of consistency tests in its updating procedure.

How worried should we be that consistency breaks down? I think we should simply adopt the same attitude as mathematicians to nonmeasurable sets. The Axiom of Choice allows them to split the surface of a sphere into five sets that can be reassembled, using only translations and rotations, into two spheres that are each congruent to the first. But they prefer to live with this version of the Banach-Tarski Paradox than to give up the Axiom of Choice. I don't think we even have the alternative of giving up the Axiom of Choice, unless we simultaneously plan to abandon any hope of modeling rational choice in worlds that are as complicated as we are ourselves.

9 Conclusion

My intention in this paper was to draw attention to the limitations of Bayesian decision theory by demonstrating its failure to cope even with the small move in the direction of realism that results from assuming that the decision-maker is not able to decide mathematically undecidable propositions.

References

- [1] M. Allais. Le comportement de l'homme rationnel devant le risque: critique des postulants et axiomes de l'école Américaine. *Econometrica*, 21:503–546, 1953.
- [2] R. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55:1–18, 1987.
- [3] K. Binmore. Debayesing game theory. In B. Skyrms, editor, *Studies in Logic and the Foundations of Game Theory: Proceedings of the Ninth International Congress of Logic, Methodology and the Philosophy of Science*. Kluwer, Dordrecht, 1992.
- [4] K. Binmore. *Fun and Games*. D. C. Heath, Lexington, MA, 1992.
- [5] C. Camerer and M. Weber. Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5:325–370, 1992.
- [6] B. de Finetti. La prevision: Ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937.
- [7] I. Gilboa and D. Schmeidler. *A Theory of Case-Based Decisions*. Cambridge University Press, Cambridge, 2001.

- [8] P. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.
- [9] J. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge, 1977.
- [10] M. Hsu, M. Bhatt, R. Adolphs, D. Tranel, and C. Camerer. Neural systems responding to degrees of uncertainty in human decision-making. *Science*, 310:1680–1683, 2005.
- [11] D. Hume. *A Treatise of Human Nature (Second Edition)*. Clarendon Press, Oxford, 1978. (Edited by L. A. Selby-Bigge. Revised by P. Nidditch. First published 1739).
- [12] L. Hurwicz. Optimality criteria for decision making under ignorance. Cowles Commission Discussion Paper, Statistics 370, 1951.
- [13] E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61:1019–1045, 1993.
- [14] R. and H. Raiffa Keeney. Additive value functions. In J. Ponsard et al, editor, *Theorie de la Decision et Applications*. Fondation national pour L'Enseignement de la Gestion des Enterprises, Paris, 1975.
- [15] F. Knight. *Risk, Uncertainty, and Profit*. Houghton-Mifflin, Boston, 1921.
- [16] R. Luce and H. Raiffa. *Games and Decisions*. Wiley, New York, 1957.
- [17] J. Milnor. Games against Nature. In *Decision Processes*. Wiley, New York, 1954. (Edited by R. Thrall, C. Coombs, and R. Davies).
- [18] L. Savage. *The Foundations of Statistics*. Wiley, New York, 1951.
- [19] R. Selten and J. Harsanyi. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA, 1988.
- [20] B. Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge, MA, 1990.
- [21] J. Von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.