Game Playing Under Ignorance

Brian Weatherson, Cornell University

Summary

In earlier work (Weatherson 2002) I argued that using 'vague probabilities' did not ground any argument for significantly adjusting Bayesian decision theory. In this note I show that my earlier arguments don't carry across smoothly to game theory. Allowing agents to have vague probabilities over possible outcomes dramatically increases the range of possible Nash equilibria in certain games, and hence arguably (but only arguably) increases the range of possible rational action.

Introduction

Many theorists in recent years have proposed that we make a small amendment to traditional Bayesian psychology. Rather than saying the partial beliefs of a rational agent are represented by a unique probability function, they represent those beliefs by a set of probability functions. Bas van Fraassen (1990) calls this set the agent's *representor*. The representor, *S*, is related to the agent's comparative probability judgments in a natural way. She regards *p* as at least as probable as *q* iff for all $Pr \in S$, $Pr(p) \geq Pr(q)$. If *S* is not a singleton, then it is possible the agent regards the *more probable than* relation as being a partial preorder, rather than a total preorder. (A partial preorder is a relation that is reflexive and transitive, and a total preorder is a relation that is reflexive and transitive and *complete* in that for all *a*, *b*, either *aRb* or *bRa*.)

Given the agent's utility function *U*, from any probability function *Pr* in *S* we can generate an expected utility function EU_{Pr}. It is a sufficient condition for the agent preferring A to B that for all $Pr \in S$, $EU_{Pr}(A) > EU_{Pr}(B)$. Whether this is a necessary condition is something that we shall return to below.

There are many reasons for preferring this way of representing uncertain agents to the traditional, single probability function approach. For one thing, it reduces the idealising demands on agents. (Richard Jeffrey calls this approach "Bayesianism with a Human Face".) For another, the arguments that *more probable than* should be a total preorder are quite weak, especially compared to the arguments for the other constraints Bayesians support. More speculatively, it is arguable that the considerations about risk and uncertainty that arise in Keynes's work on probability and economics (Keynes 1921, 1937) can be used to ground a *requirement* that an agent's subjective probabilities not form a total preorder. But for present purposes we need not address those arguments. All that matters here is that it is permissible for agents to be represented by a non-singleton set of probability functions.

An agent who is represented this way is sometimes referred to as having 'vague' or 'imprecise' probabilities (Hájek 2000). This way of talking can be quite convenient, but it is possibly misleading. An agent's comparative probability judgments can be as determinate and precise as one likes, but if those judgments generate a partial preorder, it will be necessary to represent her using a non-singleton representor. Having flagged that concern, I'll use the language of vague probabilities (and their opposite, precise probabilities) in what follows.

Decision Making with Vague Probabilities

Isaac Levi (1980, 1986) has argued that the use of vague probabilities should cause us to alter our decision theory. In particular, he defends the following two claims.

- When EU_{Pr}(A) is greater than EU_{Pr}(B) according to some, but not all, *Pr* in *S*, a rational agent will use a maximin-type rule to choose between A and B.
- Using the best such rule over a number of choices can lead to violations of various axioms Bayesians adopt, such as Independence of Irrelevant Alternatives, even though any agent with precise probabilities should conform to those axioms.

In earlier work (Weatherson 2002), I argued against such claims. I argued that the above sufficient conditions for rational preference were the *only* rationally mandatory conditions. On that basis I defended the following two theses. (I didn't frame them this way in the earlier paper, but the intent was clear enough.)

> **Rational Irrelevance of Vagueness**. Let C be a class of actions rationally performed by an agent with vague probabilities. Then there is some (possible) rational agent R with precise probabilities who performs all the actions in C.
>
> **Practical Irrelevance of Vagueness**. In evaluating any decision situation, we can proceed as if each agent is represented by a single probability function.

The arguments for this thesis were not particularly striking. They mostly consisted in pointing out that the prima facie compelling arguments for Bayesian constraints on decision-making did not make essential use of the axiom that *more probable than* is a total preorder. So I was reasonable confident the arguments could hold up against future challenges. But at the time I did not consider game theory.

Matching Pennies With Opt-Out

In orthodox game theory, agents are allowed to adopt one of the following two kinds of strategy.

- A pure strategy, where the agent picks one of the available options.
- A mixed strategy, where the agent assigns a probability to several available options, and uses a randomising device to select which option she will take.

If we are using vague probabilities, then a third kind of strategy is, in principle, available.

- A vague strategy, where the agent assigns a vague probability to several available options, and uses a randomising device to select which option she will take.

We can use vague strategies in practice by letting our final decision be guided by events about which we do not have a precise probability. (That is, if S is our representor, then there are Pr_1, Pr_2 both in S such that Pr_1(that event occurs) $\neq$ Pr_2(that event occurs).) For example, imagine Sam has an unpunctual but unpredictable colleague. Call him Lazy. It is hard to tell whether Lazy will turn up to today's faculty meeting on time. Sam uses Lazy to adopt the following strategy in a play-by-mail game of Diplomacy she is playing. If Lazy is on time, move her army into Trieste, and if Lazy is late, move her army into Venice. If she does not assign a precise probability to Lazy being on time, she not only does not know what move she will make, she does not even know what the probability is that she will, for example, move to Venice.

In general, a vague strategy for a player faced with options $O_1, \ldots, O_n$ is a set S of probability functions defined over the O_i such that

- For all $Pr \in S$, and $i \neq j$, $Pr(O_i \wedge O_j) = 0$
- For all $Pr \in S$, $Pr(O_1 \vee \ldots \vee O_n) = 1$
- For any $Pr_1, Pr_2 \in S$, and $x \in [Pr_1(O_i), Pr_2(O_i)]$, there exists a $Pr \in S$ such that $Pr(O_1) = x$

The first two conditions says that the agent determinately assigns probability 1 to making exactly one choice. The third condition says that the vague probability the agent has for choosing O_i can be represented by an interval.

When another player plays the mixed strategy S, I assume that a pure or mixed response strategy R_1 is better than a rival (pure or mixed) response R_2 iff for all $Pr \in S$, $EU_{Pr}(R_1) > EU_{Pr}(R_2)$, equal to R_2 iff for all $Pr \in S$, $EU_{Pr}(R_1) = EU_{Pr}(R_2)$, and worse than R_2 iff for all $Pr \in S$, $EU_{Pr}(R_1) < EU_{Pr}(R_2)$, otherwise R_1 and R_2 are incomparable. A pure or mixed strategy R is *basic optimal* iff it is worse than no pure or mixed strategy. A strategy is *optimal* iff it is basic optimal or a set of basic optimal strategies. (For these purposes I equate a pure strategy *Do O* with the mixed strategy $Pr(O) = 1$.) These definitions look quite plausible, especially if we are presupposing some version of **Irrelevance of Vagueness**, but note that I've bracketed the quite hard question of how we, in general, compare two vague strategies. That question is beyond the scope of this short note.

Given all that, consider the following game, called **Matching Pennies with Opt-Out**.

	A	**B**	**C**
A	(5,0)	(0,5)	(3,3)
B	(0,5)	(5,0)	(3,3)
C	(3,3)	(3,3)	(4,4)

If we restrict attention to pure and mixed strategies, the only Nash equilibrium is for each player to play C. But if we allow vague strategies, the following Nash equilibrium arises. Each player plays the vague strategy determined by $\{Pr: Pr(C) = 0\}$. This is a Nash equilibrium because if the other player plays it, literally any two pure or mixed responses will be incomparable, so any such responses will be optimal, so the set {*Pr*: *Pr*(C) = 0} will be one of very many optimal vague strategies to be played in response.

Assuming my definitions of optimality are correct, this leads to a problem for **Practical Irrelevance of Vagueness** given two extra premises.

> **Nash Permissibility**. For any game G in which each player makes exactly one move, and any Nash equilibrium N of G, there is a possible world where the players in G know each other to be fully rational and play their part of N.
>
> **Nash Requirement**. For any game G in which each player makes exactly one move, and where the players know the other players to be fully rational, the moves made will form a Nash equilibrium.

Given those premises, we can use **Matching Pennies with Opt-Out** to refute **Practical Irrelevance of Vagueness**. **Nash Requirement** says that any playing of **Matching Pennies with Opt-Out** under circumstances of common knowledge of rationality will, if we work within the framework that all players have precise probabilities, lead to each player playing C. So **Practical Irrelevance of Vagueness** says that we can assume the same thing if we drop the requirement of precise probabilities. But **Nash Permissibility** says that there is a possible world in which each plays a vague strategy that leads to either playing A or B. So **Practical Irrelevance of Vagueness** is false.

It does not, however, follow that **Rational Irrelevance of Vagueness** is false. All that principle requires is that for any action an agent performs, there is a possible agent with precise probabilities in that position who performs the same action. Now recall that if the other player plays the vague strategy {*Pr*: *Pr*(C) = 0}, *any* pure or mixed strategy is optimal. So whatever the player does, a rational duplicate of her with precise probabilities, i.e. following a pure or mixed strategy, does the same thing. The important point is that it's the *other* player being vague that opens up the permissibility of doing A or B, not the vagueness of the player who plays A or B. That suffices to defend **Rational Irrelevance of Vagueness**. This makes it, I think, doubly surprising that **Practical Irrelevance of Vagueness** fails here, as it most surely does. (Thanks to Andy Egan and Alan Hájek for comments on an earlier draft of this paper.)

Ithaca-Melbourne
21 December 2004

References

Hájek, Alan (2000) "Objecting Vaguely to Pascal's Wager" *Philosophical Studies* 98(1): 1-16.

Jeffrey, Richard (1983a) "Bayesianism with a Human Face", in J. Earman (ed.) *Testing Scientific Theories*, Minneapolis: University of Minnesota Press.

Keynes, John Maynard (1921) *Treatise on Probability*. London: Blackwell.

Keynes, John Maynard (1937a) "The General Theory of Employment", *Quarterly Journal of Economics* 51: 209-23

Levi, Isaac (1980) *The Enterprise of Knowledge*. Cambridge, MA.: MIT Press.

Levi, Isaac (1986) *Hard Choices*. Cambridge: Cambridge University Press.

van Fraassen, Bas (1990) "Figures in a Probability Landscape" in J. Dunn and A. Gupta (eds) *Truth or Consequences*. Amsterdam: Kluwer, pp 345-356.

Weatherson, Brian (2002) "Keynes, Uncertainty and Interest Rates" *Cambridge Journal of Economics*

Weatherson, Brian (forthcoming) "Should we Respond to Evil with Indifference" *Philosophy and Phenomenological Research*.