# Williams on Decision Making in cases of Indeterminate Survival

Brian Weatherson

Robbie Williams (2013) has raised the following fascinating question: how should our decision making be affected if we learn that our survival is indeterminate? Williams focusses on a case where it is indeterminate that the agent survives. I want to look at a related case, where it is indeterminate how the agent survives, to bring out some striking consequences of Williams's view. I take these 'striking consequences' to be problematic, but I'm not going to argue for that evaluation here, just draw out the consequences.

Here's the case I'm going to focus on, which is a variant of a case Williams adapts from van Inwagen (1990).

> **The Splitter**
>
> Alpha is going to go into a box, and two people are going to come out of the box, Beta and Gamma. It is determinate that Alpha will survive going into the box, but indeterminate whether she survives as Beta or Gamma. Beta and Gamma are rather different people, but they each share key characteristics with Alpha. We will leave the story open enough that you can fill in key characteristics in such a way that it turns out to be indeterminate whether Alpha is Beta or Gamma. For example, if you think it is indeterminate whether personal identity goes by physical or psychological continuity, you can imagine that Beta is a physical continuant of Alpha who is psychologically rather different, while Gamma is a psychological continuant who is physically rather different. But if your views on personal identity are different, change the example to suit. All that matters is that there are multiple features that enter in some way into identity judgments, and Beta shares many of them with Alpha, while Gamma shares the others with Alpha. The difference between Beta and Gamma should thereby be pronounced above that it is clear that they are distinct persons, not a bi-located person.

Assume that Alpha knows all of the above, is rational and purely self-interested. And assume, for simplicity, that over small enough quantities money has a constant marginal utility for Alpha. Finally, assume as a baseline that Beta and Gamma will each get half of Alpha's current money, and this baseline is independent of anything Alpha can do. The general problem will be that Alpha has to choose between various outcomes that modify this baseline by treating Beta and Gamma differently, and we want to know how she should think about the fact that it is indeterminate whether she will be Beta or Gamma.

Williams's answer starts with the idea of indeterminate credences, developed in work by Levi (1974), van Fraassen (1990), Walley (1991), and since widely adopted by many formal epistemologists. So rather than a rational agent having a credal state that is represented by a probability function, it is represented by a *representor*, which is a non-empty set of probability functions.

But Williams adds two distinctive features to this picture. First, he says that there should be one function in the set for each of the various possibilities that are not determinately not the case. He doesn't say that the set should be 'closed' in any way. So in our example, Alpha should have one function in her representor that says the probability of being Beta is 1, another function that says the probability of being Gamma is 1, and that's it, at least unless there is any non-quantifiable uncertainty about another relevant proposition.

Williams mentions in a footnote that some authors think the representor should be a convex set of probability functions, and that Richard Jeffrey (1983) rightly objected to that constraint. Jeffrey argued that there were times that the representor should be the set of functions such that $p$ and $q$ were probabilistically independent, and this is not a convex set.[1] But there are a lot of weaker constraints than convexity that the representors that Williams uses violate. Here are four more plausible constraints, all of which are violated by the representor Williams attributes to Alpha.

- Contiguity: The set of functions in the representor should be contiguous.
- Conditional contiguity: The set should be contiguous, and this should be preserved under conditionalisation on any proposition to which the functions in the representor assign probabilities.
- Intervality: For any $x, y, z \in [0,1]$ such that $x < y < z$ and proposition $p$, if there are functions $\Pr_1, \Pr_3$ in the representor such that $\Pr_1(p) = x$ and $\Pr_3(p) = z$, then there is a $\Pr_2$ in the representor such that $\Pr_2(p) = y$.
- Conditional intervality: Intervality should be preserved under conditionalisation on any proposition to which the functions in the representor assign probabilities.

Conditional intervality is a rather strong constraint, though I think it is plausible. But continguity is a very weak constraint, and it is much more plausible.

[1]For reasons related to the discussion in Moss (2011) I think it's less clear than Jeffrey thought that this is a reasonable representor, but I think something like it can work.

Nevertheless, Williams's model of a rational agent violates it. The consequences of Williams's view that I'm going to draw out can largely be traced to this feature.

There are two other features of Williams's view that are distinctive. He thinks agents should choose strategies, not acts. And he thinks agents should make choices between incomparable options randomly, not capriciously. I'll spend some time explaining each of these distinctions.

One vivid way to bring out the difference between strategy choice and act choice uses a (mild variant of) an example from Seidenfeld (1994). I'll describe the case in words, then with a graph.

> **Dominance Game**
>
> Before Alpha goes into The Splitter (as described above), a coin will be tossed. As things stand, Beta and Gamma are deemed to have bet \$10 on the coin toss. If the coin lands heads, \$10 will be transferred from Gamma to Beta, and if the coin lands tails, \$10 will be transferred from Beta to Gamma. Right now, Alpha can pay \$2 to be blindfolded and led into the box, and the coin will be tossed without Alpha knowing about it, or having the chance to do anything else before the box operates. If she passes up this chance, she'll see the coin be tossed. Then she'll have the chance to pay \$8 to cancel the bet. After that, she'll go into the box, and unless she paid the \$8, the money will be transferred between Beta and Gamma's accounts before the comes out.

To graph the game, it helps to have some names for various propositions. Let $p$ be that the coin lands heads, so $\neg p$ is that it lands tails, $q$ be that Alpha is Beta, so $\neg q$ is that Alpha is gamma, and let $r$ be the material biconditional $p \leftrightarrow q$. Let $2m$ be Alpha's cash stock before going into the machine, so the baseline is that Beta and Gamma both get $m$. Let $B$ be the act of paying \$2 for the blindfold, and going straight into the machine, and $C$ be the act of paying \$8 to cancel the bet, while the opposite act $A$ is accepting the bet. The payouts in the following table are how much Alpha ends up with. (Remember that if she pays some amount before going into the machine, only half of that comes out of her final cash stock, since the other 'split' in effect pays the other half.)
Following Seidenfeld, I've listed the payouts in terms of what Alpha receives if $r$ or $\neg r$ happens. This is slightly odd; if Alpha chooses $\neg B$, then she doesn't care about $r$ as such. But it makes it easier to compare the left and right sides of the chart.

In particular, it lets us see that $B$ is dominated by other possible strategies. Note that Alpha has five possible strategies at her disposal.

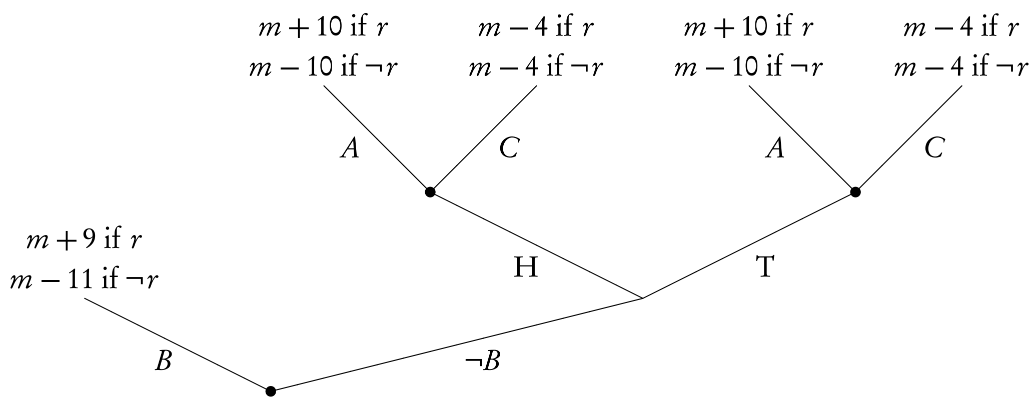


**Dominance Game**

1. Choose $B$, which we'll denote as $B$
2. Choose $\neg B$; then $A$ if the coin lands heads and $A$ if it lands tails, which we'll denote as $AA$.
3. Choose $\neg B$; then $A$ if the coin lands heads and $C$ if it lands tails, which we'll denote as $AC$.
4. Choose $\neg B$; then $C$ if the coin lands heads and $A$ if it lands tails, which we'll denote as $CA$.
5. Choose $\neg B$; then $C$ if the coin lands heads and $C$ if it lands tails, which we'll denote as $CC$.

And the immediate thing to note is that $AA$ dominates $B$. Whatever happens, whether the coin lands heads or tails, whether Alpha is Beta or Gamma, the result is \$1 better if $AA$ is chosen than if $B$ is chosen. So, if we are in the business of choosing strategies, $B$ should be out from the start. That's what Williams says should happen, and it's what I think should happen too.

But it's not what Seidenfeld thinks should happen. Indeed, the point of his paper is to argue that there are reasonable choice procedures that lead to choosing $B$. It's true that if we started out with the normal form representation of the game, with five possible strategies and four possible choices, we would see that $AA$ dominates $B$, so would not choose $B$. But Seidenfeld thinks that in the extensive form representation, as above, $B$ is a reasonable choice.

To see why, we need to say a bit about security based decision procedures. Say that two choices are *incommensurable* for an agent iff for some probability function in her representor, the expected value of the first is higher than the expected value of the second, and for some other function in the representor, the expected value of the second is higher than the expected value of the first. A security based decision procedure gives the agent a method for choosing between incommensurable options.

There are two important kinds of security based decision procedures, but they give the same result in this case. An absolute value security procedure says that in a choice between incommensurable options, the agent should choose the option that maximises the minimum possible payout. An expected value security procedure says that in such a choice, the agent should choose the option that maximises the minimum expected payout, as we look at the expected payouts of the options according to each function in the representor. Isaac Levi (1986) defends the absolute value security procedure, but it's possible to motivate both of them, I think. As it turns out, both of them end up recommending that the agent choose $B$. (This is a really nice feature of Seidenfeld's example, and the main reason I've adopted it here.) It doesn't matter whether the representor is the one Williams recommends, with just the two functions in it, or the convex closure of that function, so I'll just assume we agree with Williams so far.

To see why either security procedure recommends choosing $B$, we need to use backwards induction. Work backwards from the end of the graph. After Alpha sees the coin flip, she'll have to choose between $A$ and $C$. These choices are incommensurable. She knows which of Beta and Gamma will win the bet, but it's indeterminate which will, and hence it is indeterminate whether she is the winner. So the choices of accepting or declining the bet are incommensurable. At this stage the security procedures kick in. She'll note that the lowest possible payout for $C$ is $m-4$, while the lowest possible payout for $A$ is $m-10$. Similarly, the lowest expected payout for $C$ is $m - 4$, since that result is guaranteed, while according to the probability function that says Alpha has probability one of being the person she now knows to be the loser, the expected return of $A$ is $m - 10$. So at this stage, according to either security measure, it is better to choose $C$.

Now think about whether Alpha should choose $B$ or $\neg B$. We've already worked out that if she chooses $\neg B$, she will follow up by choosing $C$, and end with $m - 4$. If she chooses $B$ she has a 1 in 2 chance of ending with $m + 9$, and a 1 in 2 chance of ending with $m - 11$. And that's true whether she is Beta or Gamma. So the expected value of choosing $B$, according to every probability function in her representor, is $m - 1$. So $B$ is not incommensurable with $\neg B$; it is better than it according to every probability function. So Alpha should do $B$.

There are a couple of things we could conclude here. We could say that security based decision procedures are plausible, so this is a case where choosing a dominated strategy, namely $B$, is plausible. Or we could say that choosing a dominated strategy is never plausible, and this is a reason to reject security based decision procedures. I think the second option is better, but I'm not going to offer many more reasons to those who accept these procedures to change their mind.

More interestingly for present purposes, Williams's decision procedure rules out choosing $AA$. To see this, we need a simple table. This table lists the expected value of each of the five strategies, according to each of the two probability functions in her representor. We'll call $\Pr_\beta$ the function that says the probability that Alpha is Beta is 1, and $\Pr_\gamma$ the function that says that the probability that Alpha is Gamma is 1.

| | $\Pr_\beta$ | $\Pr_\gamma$ |
|---|---|---|
| $B$ | $m-1$ | $m-1$ |
| $AA$ | $m$ | $m$ |
| $AC$ | $m+3$ | $m-7$ |
| $CA$ | $m-7$ | $m+3$ |
| $B$ | $m-4$ | $m-4$ |

The key cells to look at are shaded. Each of them has an expected return greater than $m$, because the probability function thinks that Alpha plans to cancel the bet (at a small cost) iff she is going to lose it. So those are the only options that Williams thinks are acceptable for the agent to take.

This seems, I think, rather odd. It seems to me that $AA$ is a perfectly acceptable strategy for the agent to adopt. It is, I think, surprising that a symmetric situation requires an asymmetric response.

We can make this intuition more pronounced by changing the game a little. Say that **Information Game** is just like **Dominance Game**, except that if Alpha chooses to be blindfolded at the first stage, she is paid \$4, rather than being charged \$2. The expected value of each strategy is given by the following table.

| | $\Pr_\beta$ | $\Pr_\gamma$ |
|---|---|---|
| $B$ | $m+2$ | $m+2$ |
| $AA$ | $m$ | $m$ |
| $AC$ | $m+3$ | $m-7$ |
| $CA$ | $m-7$ | $m+3$ |
| $B$ | $m-4$ | $m-4$ |

Note that the best strategies have not changed. Still the only strategies that come out best, and hence the only strategies that Williams thinks the agent should randomly choose between, are *AC* and *CA*.

It is, in general, true that agents should prefer more information to less. Part of what's odd about the security based decision procedures in **Dominance Game** is that the agent is paying to avoid information. So it is not in principle absurd that the agent should forego $2 in order to see the result of the coin toss.

What's harder to see is why the agent should pay for just this information. Given the setup of the case, knowing the result of the coin flip doesn't give the agent any reason to prefer one option to another. It gives each of her representors such a reason, but not the agent herself. (The arguments here owe something to the arguments in Walley (1991) against identifying imprecise credal states with sets of precise states.) And so it doesn't seem plausible that the agent must be prepared to pay for this information.

There is an easy fix to this. If Williams imposed any of the convexity or contiguity constraints described above on Alpha's representor, there would be functions in the set according to which *B* had a higher expected return than *AC* or *CA*. It has a higher return if the probability that Alpha is Beta is 1/2, for example.

The arguments of Williamson (2007) and Cappelen (2012) have convinced many people that philosophy does not, and should not, rely on intuitions. I largely agree, but I also note that appeal to intuitions plays a much more important role in game theory, especially as practiced by economists, than in philosophy. For evidence, see the appeals to intuition in Cho and Kreps (1987), or any of the thousands of papers following up from Cho and Kreps' work. So I'll close with a simple appeal to intuition. Intuitively, it is acceptable to play *AA* in **Dominance Game**, and *B* in **Information Game**. According to Williams, it is not. That is a reason to reject Williams's view, and prefer a view that imposes convexity, or at least contiguity, constraints on representors.

## References

Cappelen, Herman. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press.

Cho, In-Koo and Kreps, David M. 1987. "Signalling Games and Stable Equilibria." *The Quarterly Journal of Economics* 102:179–221, doi:10.2307/1885060.

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

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

Levi, Isaac. 1974. "On Indeterminate Probabilities." *Journal of Philosophy* 71:391–418.

—. 1986. *Hard Choices*. Cambridge: Cambridge University Press.

Moss, Sarah. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120:1053–1069, doi:10.1093/mind/fzs007.

Seidenfeld, Teddy. 1994. "When Normal and Extensive Form Decisions Differ." In Brian Skyrms Dag Prawitz and Dag Westerståhl (eds.), *Logic, Methodology and Philosophy of Science*, 451–463. Amsterdam: Elsevier.

van Inwagen, Peter. 1990. *Material Beings*. Ithaca: Cornell University Press.

Walley, Peter. 1991. *Statisical Reasoning with Imprecise Probabilities*. London: Chapman & Hall.

Williams, J. Robert G. 2013. "Decision Making under Indeterminacy." Downloaded from `http://www.personal.leeds.ac.uk/~phljrgw/wip/dmuinew3.pdf`.

Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Blackwell Pub. Ltd.