No Expectations

Mark Colyvan

The Pasadena paradox presents a serious challenge for decision theory. The paradox arises from a game that has well-defined probabilities and utilities for each outcome, yet, apparently, does not have a well-defined expectation. In this paper, I argue that this paradox highlights a limitation of standard decision theory. This limitation can be (largely) overcome by embracing dominance reasoning and, in particular, by recognising that dominance reasoning can deliver the correct results in situations where standard decision theory fails. This, in turn, pushes us towards pluralism about decision rules.

Decision theory has always been a rich source of paradoxes. These paradoxes have helped prompt research in new and interesting directions—think of causal decision theory as a response to Newcomb's problem (Joyce 1999). The latest decision theory paradox to hit the streets is Nover and Hájek's (2004) *Pasadena paradox*. While it is tempting to think of the Pasadena paradox as a version of the St Petersburg paradox, that is to misrepresent the new paradox and to understate its significance. In fact, I think the lesson of the Pasadena paradox is quite different from what many take to be the lesson of the St Petersburg paradox. I will argue that despite first appearances, these two paradoxes are not close relatives. The Pasadena paradox draws attention to cases where an act can fail to have an expectation, despite having well-defined probabilities and utilities for each of the relevant states. I will argue that despite the failure of standard decision theory in some interesting cases, many of these cases can be dealt with by other means. This, in turn, suggests a kind of pluralism about decision rules. We should not forget, after all, that maximising expected utility is not the only game in town.

1. The Pasadena game

The Pasadena game consists of a sequence of tosses of a fair coin until the first head appears. At the appearance of the first head, the game is over. So far this is the same as the St Petersburg game (Martin 2001),

Mind, Vol. 115 . 459 . July 2006
doi:10.1093/mind/fzl695

but, unlike its Soviet counterpart, the Pasadena game has the following more complicated payoff schedule:

> If the first head appears on toss n, the payoff is given by $\$(-1)^{n-1}2^n/n$, where a negative amount indicates the punter pays the bookie and a positive amount indicates that the bookie pays the punter.

What is interesting about this game is that the expected utility calculation involves a conditionally convergent series.[1] That is, the series in question,

$$(1)\quad \sum_{j=1}^{\infty} \frac{1}{2^j}(-1)^{j-1}\frac{2^j}{j} = \sum_{j=1}^{\infty} \frac{(-1)^{j-1}}{j} = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \ldots$$

converges (to ln 2), but the related series

$$\sum_{j=1}^{\infty} \left| \frac{(-1)^{j-1}}{j} \right| = 1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \ldots$$

is divergent. Nover and Hájek then invoke the Riemann Rearrangement Theorem (Apostol 1967, p. 412) to show that the resulting expected utility can be made to converge to any finite value or diverge to positive or negative infinity—it all depends on the order of the terms in the series. But the terms in the series do not come in any natural order, and therein lies the problem. The value of the game crucially depends on something that is neither specified by the game nor considered part of the usual specification for a well-posed decision problem.

Nover and Hájek also point out that the situation in Pasadena is decidedly worse than in St Petersburg. For a start, the St Petersburg game is a good game, in the sense that the expected utility is positive.[2] If you are offered the St Petersburg game for free, you would be a fool not to play. With the Pasadena game, since the expectation can be made to fall anywhere in the interval $(-\infty, +\infty)$, we cannot even determine whether it is a good game. The Pasadena game, however, does have a couple of points of contact with its Soviet counterpart: (i) the utilities need to be unbounded; (ii) a partition of the state space into infinitely many states is required. Although both of these assumptions can be

[1] A *conditionally convergent* series $\sum_{j=1}^{\infty} a_j$ is one which is convergent but the related series $\sum_{j=1}^{\infty} |a_j|$ is divergent. A series is *absolutely convergent*, if both series converge.

[2] Of course the exact expectation of the St Petersburg game is controversial (is it *really* infinite?), but it seems clear that the expectation is defined (though see Broome 1995 for disagreement on this) and that it is positive. That is all I am claiming here.

questioned, I will not head down that path. Nover and Hájek make a convincing case for accepting infinite utilities and infinite decision tables. Less convincing is their case against the view that decision theory 'goes silent'. In keeping with the geographic nomenclature, let us call this response the *St Lucia solution*. I articulate and defend it in the next section.

2. The St Lucia solution

In a nutshell, my proposal is that the Pasadena game does not have an expectation (or alternatively, there is no fact of the matter about what the expectation of the game is). The resulting decision problem (play or not play the Pasadena game) is thus ill posed. We must insist that all well-posed decision problems have acts with well-defined (i.e. unique though not necessarily finite) expected utilities.[3] In particular, the expectations should not depend on the order of the states in the decision table. That is, we should insist on the relevant expectation series being absolutely convergent. The lesson of the Pasadena paradox might be taken to be that decision problems where we only have conditional convergence are, despite appearances, not fully specified, and thus ill posed. And bear in mind that we have no hesitation rejecting as ill posed a decision problem where some of the probabilities or utilities are not given. My suggestion is that the problem of whether or not to play the Pasadena game is essentially the same.

Now this response to the Pasadena paradox might be thought to be *ad hoc* but I do not think it is. After all, insisting that in well-posed decision problems, all acts have well-defined expectations is not *ad hoc*. Indeed, we already know that there are cases where there is no well-defined expectation and we have no hesitation in rejecting the corresponding decision problems as ill posed. Consider, for example, a game in which the expected utility of playing is given by:

$$\sum_{n=1}^{\infty}(-1)^{n-1}$$

This series is divergent and since it does not diverge to positive or negative infinity, the game does not have a well-defined expectation. Or consider a game where the utilities or probabilities are vague over an interval. Again it seems uncontroversial that there is no well-defined

[3] For the moment we will only consider decision theory as maximisation of expected utility. In the final section I will relax this constraint and advance a slightly more liberal conception of decision theory.

expectation for such a game. That there are acts where there is no well-defined expectation is not news. What *is* news is that there are cases in which the expectation series is conditionally convergent (as in (1)) and that these too fail to have well-defined expectations. What is puzzling about these latter cases (of which the Pasadena game is one example) is that it seems as though there should be an expectation.

It is worth noting that Nover and Hájek (2004) do accept that the expectation of the Pasadena game is undefined (p. 305). Despite this, they seem to be of the view that the decision problem of whether or not to play the game is well defined (p. 310). They do not argue for the latter.[4] They simply rely on our intuitions to sweep us along and accept that the decision problem of whether or not to play the Pasadena game is well posed.[5] Indeed, as they point out, the case for the Pasadena game having an expectation is rather weak, and once we learn a little about infinite series, the case is untenable. And, as I have pointed out above, given that we are already aware of cases where decision theory goes silent (despite having well-defined utilities and probabilities for each act–state pair) we need some justification for accepting as well posed the decision problem of whether or not to play the Pasadena game. This justification is not forthcoming. But I do not want to resort to simply pushing the burden of proof around. Nover and Hájek draw attention to some of the shortcomings of the St Lucia solution and I will respond to these in the next section.

3. An uncomfortable silence?

Nover and Hájek point out that if we accept that the Pasadena game has no expectation, we cannot compare the Pasadena game to other games. For a start, we cannot say that it is worse than the St Petersburg game. We cannot even say that it is worse than the neighbouring Altadena game, whose payoffs are one dollar higher than the Pasadena game.

But intuitions are not entirely reliable here. After all, intuitions are bad enough when it comes to infinities and infinitesimals, but when we

[4] Nor do they explain what they mean by '[t]he game is apparently well defined' (Nover and Hájek 2004, p. 310). Usually we think of *decision problems* as being well defined (or not), not the games that are involved in those problems. I take it that they mean that the corresponding decision problem of whether or not to play the Pasadena game is well defined.

[5] They do provide an indirect argument, by putting pressure on what I am calling the St Lucia solution. This, they claim, is not a happy solution and we will get to their concerns in this regard shortly.

get to undefined quantities, intuitions abandon us altogether.[6] So in reply to this pure intuition argument, an initially tempting response is to bite the bullet. Since both the Altadena and the Pasadena games do not have expectations, they are incomparable. If this is unintuitive, then so be it. It might also seem unintuitive that 2/0 is not comparable to 1/0, but so much the worse for our intuitions.

But perhaps we can do better than this. For example, the intuition that the Pasadena game is worse than the Altadena game can be backed up by dominance reasoning.[7] What is interesting about dominance reasoning is that it depends only on utilities, not on expectations. So in cases such as the Pasadena game versus the Altadena game, dominance reasoning gives clear advice: choose the Altadena game.[8] Standard (expected utility) decision theory, on the other hand, offers us no advice.[9] But this just serves to highlight the problematic relationship between dominance reasoning and the principle of maximising expected utility.[10] Perhaps the right attitude here is a kind of pluralism

[6] Consider an "intuitive" argument as to why $2/0 > 1/0$: because for all positive x, $2/x > 1/x$, then consider the limit as $x \to 0$ and conclude that $2/0 > 1/0$. But of course 2/0 and 1/0 are undefined and are thus not comparable via the greater-than relation with either each other or to any real number. Interestingly, the comparison of the values of the Pasadena and Altadena games is similar in some respects. Think of these two games as limits of a sequence of truncated, finite games. For each finite game, the expectation of the Altadena game is higher than the Pasadena game. That is, each pair of finite games is well behaved and comparable but in the limit things go awry. But in expected utility decision theory it is the limits we are interested in, and here intuitions about finite cases are of no use. I will return to this line of thought in a moment. For, as we shall see, there is a way of making good the intuition that the Altadena game is better than the Pasadena game. But to do this we need to step beyond expected utility decision theory.

[7] Dominance reasoning suggests that one ought to choose act A_1 over act A_2 if in every state the utilities associated with A_1 are never less than the corresponding utilities for A_2, and in at least one state the utility of A_1 is higher than the corresponding utility for A_2. It should also be mentioned that this rule is only applicable when the states are independent of the acts.

[8] Note that issues concerning the order of presentation do not arise here. For we are only considering utilities (not expected utilities) and these are attached to specific act–state pairs. All we need is that each utility in the Altadena game is not lower than the corresponding utility (i.e. the utility for the same state) in the Pasadena game and that in at least one state the Altadena game's utility is higher. This is guaranteed by construction.

[9] A similar dominance argument can be employed in other cases where the principle of maximising utility is silent. For example, dominance can be used to justify our intuitive preference for the St Petersburg game over the Pasadena game. Moreover, dominance reasoning can be invoked to justify preference for a variant of the St Petersburg game in which the first payoff is increased by a dollar over the regular St Petersburg game. We might even be able to use dominance reasoning to justify the choice of some vaguely specified games over others. (For example, consider a toss of a fair coin with payoffs $0 for heads and something between $4 and $6 for tails. Dominance reasoning tells us that this game is better than the same game with payoffs $0 for heads and something between $2 and $3 for tails.)

[10] Newcomb's problem (Joyce 1999; Nozick 1969), for instance, can be seen as a conflict between these two principles.

about decision rules. At least in the case of the Pasadena game versus the Altadena game we do not have conflicting advice. We simply have dominance reasoning telling us to choose the Altadena game and the principle of utility maximisation remaining silent. This is only problematic if you think that whenever dominance reasoning applies, it should offer the same advice as the principle of utility maximisation. But this might be asking for too much. Unequivocal and sound advice from one decision rule is surely enough.

There are a couple of interesting consequences of my suggested pluralist attitude towards decision theory. First, I note that it is common to see the principle of maximising expected utility as the more general decision rule. Dominance reasoning is thought to apply in only a limited number of cases. But if what I am suggesting here is correct, neither rule can properly be thought of as more general. Maximising expected utility will apply when there are well-defined expectations (and the Pasadena paradox shows us that this is not simply a matter of having well-defined probabilities and well-defined utilities for each act–state pair). Dominance reasoning will apply whenever the problem satisfies the appropriate dominance conditions and the states are independent of the acts. What we find is that there is a class of cases where we can rightfully use dominance reasoning but not the principle of maximising expected utility. These cases include familiar cases where the probabilities are unknown as well as a newly identified class of cases where the expectation of one or more acts is undefined because it involves a conditionally convergent series. Moreover, the two rules of decision theory should not be thought of as being in conflict here (as they seem to be in Newcomb's problem). It is just that they each have different (though not mutually exclusive) domains of applicability.

A related issue is what counts as a well-posed decision problem. Previously, we thought we simply needed to ensure that the states formed a partition of the entire state space and that we had well-defined (and consistent) probability and utility assignments for each act–state pair. But the Pasadena paradox shows us that that is not always enough. If we are to use the principle of maximising expected utility we must also insist that the expectations of each act are well defined. On the other hand, if we wish to use the dominance principle, we need to insist that the usual dominance conditions are satisfied. So what is a well-posed decision problem? Well, on my view it is a disjunctive matter: a decision problem is well posed if *either* it has well-defined expectations for each act (and, as usual, the states form a partition of the states space), *or* it satisfies the dominance conditions (outlined in footnote 7).

So, to return to Nover and Hájek. They are right that the expectation of the Pasadena game *is* vexing. It is vexing because there is *no expectation*. And this is quite different from the St Petersburg game where the expectation is vexing because it is infinite. Nover and Hájek also suggest that accepting what I have called the St Lucia solution means that decision theory goes silent. They go on to argue that this is an uncomfortable silence. I have suggested a way to avoid some of the most uncomfortable silences: we adopt a pluralist attitude toward decision rules. This allows dominance reasoning to fill the conversational void left when the principle of maximising expected utility has nothing to say. So although on my view the Pasadena game has no expectation and the corresponding decision problem (play or not play) is not well posed, this does not imply that there are no well-posed decision problems featuring the Pasadena game. The problem of whether to play the Pasadena game or the Altadena game, for instance, can be settled by dominance reasoning. Any residual uncomfortable silences I suggest we learn to live with. Without wishing to sound Wittgensteinian here, sometimes silence is the right response—uncomfortable or otherwise.[11]

Department of Philosophy MARK COLYVAN
University of Sydney
Sydney, NSW, 2006
Australia
mcolyvan@mail.usyd.edu.au

References

Apostol, T.M. 1967: *Calculus*, Vol. 1. New York: John Wiley and Sons.
Broome, J. 1995: 'The Two-envelope Paradox'. *Analysis*, 55, pp. 6–11.
Joyce, J. 1999: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
Martin, R. 2001: 'The St Petersburg Paradox', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Fall 2001 Edition)* [online

[11] Thanks to Alan Hájek, Storrs McCall, Harris Nover, and Juliana Weingaertner for conversations or email exchanges on the topic of the Pasadena paradox, and to Alan Hájek for many very helpful comments on an earlier draft. I am also grateful to the Center for Philosophy of Science at the University of Pittsburgh where I held a visiting Fellowship in 2004. Some of the work on this paper was completed while enjoying the stimulating research environment afforded by this centre.

encyclopedia], Fall 2001 <http://plato.stanford.edu/archives/fall2001/entries/paradox-stpetersburg/>.

Nover, H. and A. Hájek 2004: 'Vexing Expectations'. *Mind*, 113, pp. 237–249.

Nozick, R. 1969: 'Newcomb's Problem and Two Principles of Choice', in Rescher 1969.

Rescher, N. (ed.) 1969: *Essays in Honour of Carl G. Hempel*. Dordrecht: Reidel.

