

Costs of abandoning the Sure-Thing Principle

Rachael Briggs

Stanford University, Stanford, CA, USA

ABSTRACT

Risk-weighted expected utility theory (REU theory for short) permits preferences which violate the Sure-Thing Principle (STP for short). But preferences that violate the STP can lead to bad decisions in sequential choice problems. In particular, they can lead decision-makers to adopt a strategy that is dominated – i.e. a strategy such that some available alternative leads to a better outcome in every possible state of the world.

ARTICLE HISTORY Received 3 November 2015; Accepted 13 November 2015

KEYWORDS Risk-weighted expected utility; Sure-Thing Principle; expected utility; decision theory; allais paradox


Lara Buchak's risk-weighted expected utility theory (REU theory, for short) is a flexible and well-motivated generalization of standard expected utility theory. But REU theory has a crucial cost: it permits preferences which violate the Sure-Thing Principle (STP). I will argue that violations of the STP come at significant cost, both to decision-makers and to the theorists who model them.

First, consider the decision-makers. A decision-maker whose preferences violate the STP may adopt a *dominated strategy* – i.e. a strategy such that some available alternative strategy leads to a better outcome in every possible state of the world.

In chapter 6, Buchak considers and rebuts a version of this objection which we might call

The Diachronic Challenge

- (1) According to REU theory, it is sometimes permissible to choose a dominated strategy in an extended choice situation.
- (2) Choosing a dominated strategy in an extended choice situation is irrational.

CONTACT Rachael Briggs  rbriggs@stanford.edu

© 2016 Canadian Journal of Philosophy

∴ REU theory fails to provide sufficient conditions for rationality.

Buchak challenges both premises. I'll respond to Buchak's arguments against each premise, and sketch out an alternative possible response to the Diachronic Challenge.

I'll also argue that REU theory's failure to accommodate the STP presents a pragmatic cost to decision theorists who aim to develop simple models of decision-makers. The STP, I claim, is a valuable tool that enables theorists to simplify so-called grand-world decisions into so-called small-world decisions. REU theory is missing this valuable tool.

1. The Sure-Thing principle

STP For all acts f and g , constant acts x and events E , $f_E x > g_E x$ iff $f_E > g_E$ (Buchak 2013, 107). A corollary of the STP is

Savage's STP For all acts f and g , constant acts x and y and events E , $f_E x > g_E x$ iff $f_E y > g_E y$ (Savage 1954, 23). Intuitively, the STP says that a person's preference between two acts should depend only on their outcomes in states where those outcomes differ. So (Savage's STP helps us elaborate) whatever your preference is between f if E and x otherwise and g if E and x otherwise, replacing x with y shouldn't affect that preference.

As an example of an application of the STP, suppose you are planning the route for your next bushwalk. You're choosing between two acts: you can either plan to take your favourite trail (which is reliably pleasant, but offers few surprises), or plan to take a new route (which will bring you stunning mountain views, if you can only manage the difficult hike to the summit). There is a small chance that it will rain, in which case both acts will yield the same outcome: you will cancel the bushwalk and sit indoors reading. The STP says that, when choosing which hike to prepare for, you should ignore the possibilities where it rains (since in that event, both decisions have the same outcome) and consider only what each hike will be like if it does not rain.

I will now give an example illustrating how REU theory violates the STP. Let Rhoda be an REU-maximizer for whom $r(p) = p^2$. At various points, I will find it useful to contrast Rhoda with Eualie, an EU-maximizer for whom $r(p) = p$.

Let us consider how Rhoda and Eualie would respond to the Allais Paradox (Allais 1953). In the Allais paradox, a ticket is drawn at random from a batch of 100 tickets numbered 1–100. Columns in the table correspond to states, which specify which ticket is drawn. Rows correspond to acts. Outcomes are labelled according to their utilities.

	Ticket 1	Tickets 2-11	Tickets 12-100
L_1	1000	1000	1000
L_2	0	2000	1000
L_3	1000	1000	0
L_4	0	2000	0

Rhoda will prefer L_1 over L_2 , and L_4 over L_3 . Eulalie, on the other hand, will prefer L_2 over L_1 and L_4 over L_3 .

Rhoda's pattern of preferences violates the STP. Where f is any act that yields 1000 utils if Tickets 1–11 are drawn, and g is any act that yields 0 utils if Ticket 1 is drawn and 2000 utils if Tickets 2–11 are drawn, and E is the proposition that one of Tickets 1–11 is drawn, we have

$$L_1 = f_E 1000$$

$$L_2 = g_E 1000$$

$$L_3 = f_E 0$$

$$L_4 = g_E 0$$

and so by Savage's STP, we have the requirement that $L_1 > L_2$ iff $L_3 > L_4$. Rhoda's preferences violate this biconditional, and hence Savage's STP, and hence STP.

2. Do REU-maximizers choose dominated options?

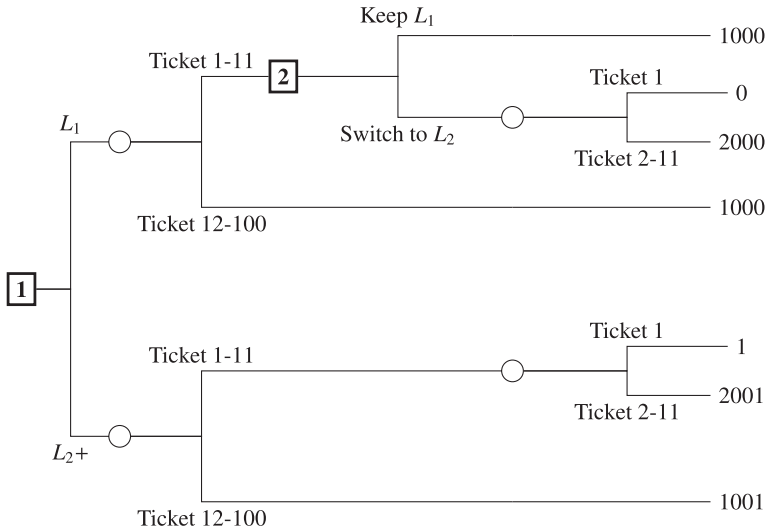
Buchak considers two arguments for Premise 1 of the Diachronic Challenge. (Recall, this is the premise that

1. According to REU theory, it is sometimes permissible to choose a dominated strategy in an extended choice situation.)

The first argument purports to show that agents who violate the STP will pay to go back on earlier decisions. The second purports to show that they will pay to avoid information. (Buchak also considers a third argument meant to establish that the REU-maximizer cannot stick to a plan, but this argument does not, strictly speaking, involve financial exploitation. I will omit discussion of this third argument; everything I have to say about it is covered in my discussion of Problem 1.)

Throughout my comments, I will focus on the special case of Rhoda. But what I say can be generalized to any REU-maximizer who violates the STP, as Buchak shows in the book.

Problem 1 Suppose Rhoda is offered the following sequence of choices. First, she may choose either L_1 or L_2 + (a version of L_2 that is sweetened by adding one util to each outcome). Next, a ticket is drawn. If Rhoda has chosen L_1 and the ticket is numbered 1–11, then she is given the opportunity to switch to an unsweetened L_2 . In extended form, the decision looks like this:



At node [1], Rhoda will choose L_1 over L_2+ , since L_1 has a higher REU. But if she reaches node [2] and learns that one of tickets 1–11 has been drawn, she will switch to L_2 , since L_2 will now have a higher REU than L_1 in the light of her new information. So Rhoda’s strategy will be to choose L_1 at [1] and switch to L_2 at [2], if she gets there. But this strategy is dominated. The strategy of choosing L_2+ at node [1] results in a better outcome no matter which ticket is drawn.

Furthermore, the problem is essentially linked to the failure of the STP. At node [1], Rhoda strongly prefers choosing L_1 and sticking to L_1 (call this option ‘stick’) over choosing L_1 and sticking to L_2 (call this option ‘switch’). But although

$$stick > switch,$$

where E is the proposition that a ticket from 1–11 is drawn, we also have both

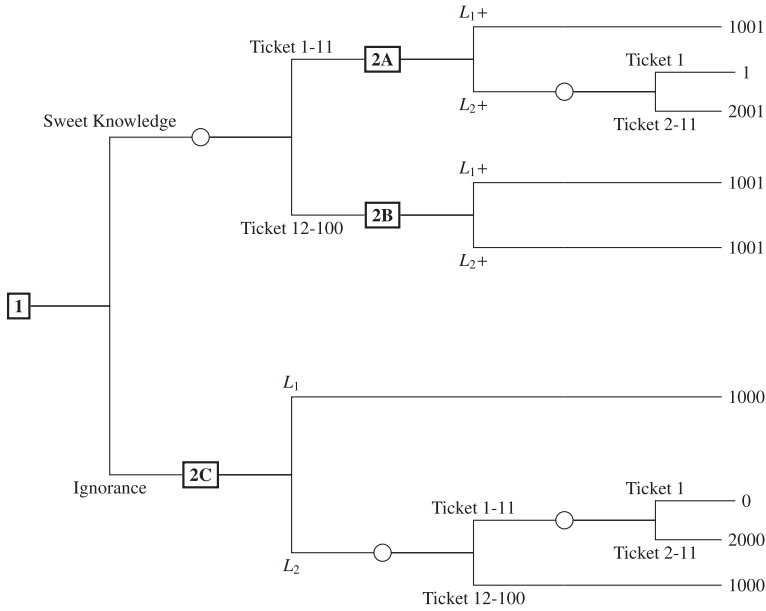
$$switch_E > stick_E$$

and

$$switch_{\neg E} \succsim stick_{\neg E}$$

Rhoda strictly prefers sticking to switching, even though she weakly prefers switching to sticking conditional on both E and $\neg E$.

Problem 2 Suppose Rhoda will be offered a choice between L_1 and L_2 . Beforehand, she has the opportunity to decide whether to make this choice in complete ignorance, or whether to make it with knowledge of whether the ticket drawn was numbered between 1–11 or between 12–100. Additionally, the knowledge comes with a sweetener: if she chooses knowledge, an extra util is added to each outcome.



Rhoda can then reason as follows. If she gets to state [2A], she will choose L_2+ (since it will have a higher REU than L_1+); if she ends up in [2B], then she will end up with the same outcome no matter what she chooses. If she instead ends up at [2C], she will choose L_1 . At state [1], she assigns higher REU to L_1 than to L_2+ . So at state [1], she should choose ignorance (which will ensure that she ends up choosing L_1) over sweet knowledge (which may result in her choosing L_2+). But the strategy of choosing ignorance at [1], followed by L_1 at [2C], is dominated by the strategy of choosing sweet knowledge at [1], followed by choosing L_1+ at [2A] and at [2B].

Again, the problem is closely linked to the violation of the STP. We have all of the following:

$$\begin{aligned}
 L_1 &> L_2+ \\
 L_2+E &> L_1E \\
 L_2+-E &> L_1-E
 \end{aligned}$$

Rhoda strictly prefers L_1 to L_2+ , even though she strictly prefers L_2+ to L_1 conditional on both E and $\neg E$. From her L_1 perspective, she should avoid anything that will lead her to choose L_2+ over L_1 – which is precisely what knowledge whether E will do.

2.1. Preference and choice

I have argued that Rhoda will choose the dominated option in Problem 1, and that she will choose the dominated option in Problem 2. Buchak claims that both arguments are unsound.

She points out that both arguments rely on contentious assumptions about the relationship between preference and choice. In fact, Buchak notes, the assumptions needed to prove that Rhoda will choose a dominated strategy in Problem 1 contradict the assumptions needed to prove that Rhoda will choose a dominated strategy in Problem 2, so it cannot be that both arguments are sound. To unpack these assumptions, it will be helpful to discuss Rhoda's preferences in each of the problems in more detail.

First, consider Problem 1. The following table gives Rhoda's preference ranking over all the strategies at each node, as determined by REU theory. (The left-hand column lists the node at which the preferences are held; the right-hand column lists the strategies from most-preferred at the top, to least-preferred at the bottom. All preferences are strict.)

Node 1	[1] L_1 , [1] L_2+ [1] L_1 , [2] Switch
Node 2	[1] L_2+ [1] L_1 , [2] Switch [1] L_1 [2] Stay

Next, consider Problem 2. Rhoda does not care (at any node) what she chooses at [2B] – if two strategies differ only with respect to what she chooses at [2B], then she is always indifferent between them. I will therefore consider her preferences among coarse-grained strategies, which do not specify what she chooses at [2B]. Her preference ranking is as follows at each node.

Node 1	[1] Sweet Knowledge [1] Ignorance [1] Sweet Knowledge [1] Ignorance	[2A] L_1+ [2C] L_1 [2A] L_2+ [2C] L_2
Node 2A	[1] Sweet Knowledge [1] Ignorance [1] Sweet Knowledge [1] Ignorance	[2A] L_2+ [2C] L_2 , [2A] L_1+ [2C] L_1
Node 2C	[1] Sweet Knowledge [1] Ignorance [1] Sweet Knowledge [1] Ignorance	[2A] L_1+ [2C] L_1 [2A] L_2+ [2C] L_2

There are three possible ways to derive choices from preferences in extended-form decision problems. All three coincide in the above problems for Eulalie, but they come apart for Rhoda.

Naive Choice At each node, choose the action that belongs to the best strategy (according to your current preferences) available at that node.

Sophisticated Choice Assume that if you reach a final choice node n , you will choose the action with the best outcome (according to your preferences at n). Assume that no other outcome is possible once you reach node n . Work backward through the tree until you reach the first choice node (Hammond 1988).

Resolute Choice Choose the best strategy at the first node (according to your preferences at the first node) and adhere to it at all other nodes, regardless of your later preferences (McClennen 1990). In Problem 1, naive Rhoda will adopt the dominated strategy of choosing L_1 at [1] and switching at [2]. But neither the sophisticated nor the resolute chooser will take up a dominated plan. Sophisticated Rhoda will pick L_2+ at [1] to avoid switching at [2]. And resolute Rhoda will pick L_1 and stick to it.

In Problem 2, it is sophisticated Rhoda who will adopt the dominated strategy of choosing Ignorance at [1] (to stave off the possibility of choosing L_2+ at [2A]), followed by L_1 at [2C]. Naive and resolute Rhoda will pick Sweet Knowledge at [1]. (Naive Rhoda will follow up with L_2+ , should she reach [2A], while resolute Rhoda will follow up with L_1+ .)

Buchak does not think that sophisticated Rhoda will choose a dominated option in Problem 2. She offers the following rebuttal. (I have altered Buchak's notation to match my own.)

If [objectors like Machina (1989) and McClennen (1990)] are correct that sophisticated choosers can end up with a dominated option, then sophisticated choosers are in trouble. However, the sophisticated chooser can respond to this argument by pointing out that although she ends up with L_+ , L_1+ is not actually an available option for the agent at [1], since at [1] the agent knows she will not choose L_1+ at [2C]. The view of agency which makes sophisticated choice attractive in the first place is one on which once our future preferences are fixed, certain logically possible future options are not within our grasp. Granted, the fact that the option is unavailable to the agent at [1] because of her own future choices might reveal that she has diachronically inconsistent preferences, a possibility which I will examine in the following section; but her plan does yield, for each given time-slice, the best consequence available to that time-slice [189].

The objection is that by the sophisticated chooser's lights, Rhoda's options are *acts*. But there is no time at which L_1 is the result of any act that is available to sophisticated Rhoda. It is not available at [1], because if she gets to node [2A], she will choose L_2+ and not L_1+ . It is not available at node [2A], because, being the kind of person who will choose sweet ignorance at [1], she will never reach [2A]. And it is not available at [2C], because L_1+ is no longer available at [2C].

Nonetheless, there is a sequence of actions, each of which Rhoda can perform, that will guarantee that she ends up with L_1+ . (All she has to do is choose sweet knowledge at [1], and choose L_1+ at [2A], should she get there.) In other words, there is a strategy that is available to Rhoda that results in her ending up with L_1+ . Buchak holds that on the view that motivates sophisticated choice,

we are not entitled to evaluate entire strategies, since there is no time at which strategies are the object of choice. But the objector should press the point that a strategy can be available, even if there is no one time at which it is available.

The upshot is that if Rhoda is either naive or sophisticated, she will choose a dominated strategy in one of the above examples. If Rhoda is resolute, she will never choose a dominated strategy – not in the above problems, and not in any other extended decision problem.¹ Unless the REU theorist is prepared to accept that resolute choice is the right approach to extended decision problems, she is compelled to accept Premise 1.

3. Is choosing a dominated strategy a sign of irrationality?

The Diachronic Challenge against REU theory has not yet succeeded. Buchak also presents arguments against Premise 2, which states:

2. Choosing a dominated strategy in an extended choice situation is irrational.

The first part of Buchak's argument is negative: she points out that several apparently tempting diagnoses of the irrationality fail. For instance, it is *not* true that Rhoda is incapable of picking a strategy and sticking to it. If Rhoda is sophisticated or resolute, she will stick to the strategy she chooses at the first node of the decision tree.

Nor, Buchak argues, can Rhoda be accused of having inconsistent preferences across time. Her preferences over outcomes remain constant throughout Problem 1 and Problem 2. True, her preferences over lotteries change, but so do Eulalie's preferences over lotteries: Eulalie strictly prefers L_2 over L_1 , but will reverse this preference if she learns that E .

The objector might make the following complaint: REU theory allows preferences to change in a predictable direction, while the more restrictive EU theory does not. If an EU-maximizer is prepared to revise her assessment of a lottery *upward* on learning E , she must also be prepared to revise it *downward* on learning $\neg E$. Not so for the REU-maximizer Rhoda, who revises her opinion of L_2 upward whether she learns E , or $\neg E$.

In addition to her negative account of what is *not* wrong with choosing a dominated option in REU theory, Buchak has a positive explanation of why REU-maximizers are willing to pay to avoid information. The positive account is aimed not just at *explaining* the choices of sophisticated REU-maximizers, but at *vindicating* them – showing that they are not inconsistent. The positive diagnosis is as follows.

Because REU theory allows for foreseeable preference reversals, Buchak argues, it counts new information about E as bad, from Rhoda's initial perspective. This information will foreseeably reverse Rhoda's preferences, causing her to make a decision whose potential costs (when appropriately weighted according to their probabilities, by her own lights) outweigh its potential benefits. So,

Rhoda has a pragmatic reason to avoid this information, even if she has to pay to avoid it (as in Problem 2).

The trouble (from Rhoda's initial perspective) is that the information may be misleading – and the prospect of being misled is both bad enough and likely enough to make the information bad overall. There are some states of the world where Rhoda learns E , and accepts L_2 on the basis of that information, even though L_2 in fact leads to a very bad outcome in those states. (The relevant state in the example is that Ticket 1 is drawn.) Furthermore, because Rhoda is risk-averse, the possibility of being misled has a stronger negative impact on the overall value of the information than it would for Eulalie.

From Rhoda's later perspective, however, the information is good on balance. While it still has a chance of being misleading, the outcome in which she is misled is not bad enough and likely enough to make the information bad overall. Instead, the information is either good (if she learns E) or neutral (if she learns $\neg E$). From her later perspective, she should act on this overall good information, even if it means going back on an earlier, less-informed decision (as in Problem 1).

But is the information good overall, or bad overall? There is no answer to this question, says Buchak. The information is bad from the earlier perspective and good from the later perspective – there is no neutral perspective from which to reconcile them. While it is true that Rhoda's later views about the value of information are guaranteed to diverge from her earlier views, there is no requirement that her earlier and later views agree. Thus, it is perfectly consistent for Rhoda's views about the value of information to change in a predictable way.

So the disagreement between Buchak and her objector can be recast as a disagreement about whether there must always be a neutral perspective from which to evaluate new information: the objector thinks there must be, and Buchak thinks there needn't be. Both views are internally consistent. To mount a good argument against the REU theorist, the objector needs to either find a compelling internal argument that REU theory is inconsistent, or find a compelling external reason that the REU theorist's assumptions about rationality are implausible.

I think there is another line of argument available to the objector here. There is a simpler story available about why choosers who settle on dominated strategies are practically irrational. Being disposed to settle on a dominated strategy is not *indicative* of some rational flaw such as inconsistency. Rather, it already *constitutes* practical irrationality.

This argument has both an internal and an external version. According to the internal version, choosing dominated strategies is also irrational by the REU-maximizer's own lights. Notice that REU lets us evaluate both acts and strategies. On any continuous r function with a minimum of 0 and a maximum of 1, L_1+ has a strictly higher REU than L_v and L_2+ has a strictly higher REU than L_2 . So dominated strategies are worse than the strategies that dominate them

by the REU-maximizer's own lights. By their own lights, both the sophisticated and the naive REU-maximizer will sometimes choose a worse strategy when a better one is available.

According to the external version of the argument, the aim of decision theory is to enable us to choose options with the best consequences. To choose a dominated option (whether 'option' is understood as 'act' or 'strategy') frustrates this aim, in a way that is foreseeable *a priori*. The following is an independently compelling claim about rationality: if it is knowable *a priori* that strategy *a* yields a better result than strategy *b*, then it is pragmatically irrational to choose strategy *b* when strategy *a* is available.

3.1. *Where things stand so far*

So far, I have raised arguments against Buchak's response to the Diachronic Argument. I have argued in favour of Premise 1, which says that every REU-maximizer who violates the STP will sometimes choose a dominated strategy. Adopting a resolute account of rational choice solves the problem, but naive and sophisticated choice do not. Therefore, I claim REU theorists can deny Premise 1 only at the cost of embracing resolute choice.

And I have argued in favour of Premise 2, which says that if an REU-maximizer is disposed to choose a dominated strategy, then that REU-maximizer is irrational. I claim that being disposed to choose a dominated strategy is constitutive of irrationality, whether or not it is also the sign of some deeper inconsistency.

I would like to briefly suggest a third way out: why not require that the REU-maximizer's *r* functions change in response to new information? In both Problem 1 and Problem 2, I assumed that $r(p) = p^2$ both before and after Rhoda learned *E*. And in both problems, the exploitation of Rhoda turned on foreseeable preference reversals: there were some node (*n*) partition $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$ and gambles *f* and *g* such that

- At node *n*, Rhoda would learn which member of \mathcal{E} was true.
- Rhoda preferred *f* to *g*.
- For every member of \mathcal{E} , if Rhoda were to learn *E*, Rhoda would prefer *g* to *f*.

To get around the problem, an REU theorist would need to claim that Rhoda's preferences were not just *permitted* to change, but *required* to change in the light of new information. A fully fleshed out version of this argument would include rules about how to change *r* in the light of new evidence, but I propose it as a promising avenue of future research.

4. A lesson about simplifying lotteries

I think there is another lesson to be drawn from Rhoda's preference reversals – one that brings out a hidden cost of REU theory. Buchak has already pointed out

that for Rhoda, information lacks a stable value. I suggest an underlying reason for this: for Rhoda, sub-acts lack stable utility values.

In claiming that sub-acts lack stable values, I don't aim to say only that the utilities of sub-acts can change in the light of new information. Rhoda and Eulalie agree that $L_2 +_E$ loses value in the light of the information that Ticket 1 was drawn, because this information rules out some of the states in E – it tells us something about which states obtain if E is true. But for Rhoda, even propositions entailed by E – that is, even propositions that give no information about which states obtain if E is true – will affect the value of sub-acts on E .

To see why sub-acts lack stable values for Rhoda, consider $L_2 +_{E'}$ which yields 1 util if Ticket 1 is drawn, and 2001 utils if one of tickets 2–11 is drawn. What, we might ask, is the utility of $L_2 +_{E'}$?

If Rhoda is at [2A] in Problem 2, where she is certain of E (the proposition that a ticket numbered 1–11 is drawn), then we can calculate the utility of $L_2 +_{E'}$ as the utility of the outcome o such that she would be indifferent between $L_2 +_{E'}$ and an assured prize of o :

$$u(o) = REU(L_2 +_{E'}) = 1 + \left(\frac{10}{11}\right)^2 \times (2001 - 1) = \frac{500,121}{121} = 991.1$$

If Rhoda is at node [1] in Problem 2, then we can calculate the utility of $L_2 +_{E'}$ as the utility of the outcome o such that she would be indifferent between $L_2 +_{E'}$ and $o_E L_2 +_{E'}$ (the act that yields o in the event that E and the same outcome as $L_2 +_{E'}$ in the event that $\neg E$). Since $L_2 +_{E'}$ yields 1001 utils in all $\neg E$ states, I will also write this act as $o_E 1001$.

$$\begin{aligned} REU(L_2 +_{E'}) &= REU(o_E L_2 +_{E'}) \\ 990.1 &= REU(o_E 1001) \end{aligned}$$

Since $REU(o_E 1001) < 1000$, we know that $u(o) < 1001$. Thus:

$$\begin{aligned} 990.1 &= u(o) + \left(\frac{11}{100}\right)^2 1001 \\ u(o) &= 979 \end{aligned}$$

So the same sub-act $L_2 +_{E'}$ is worth 991.1 utils at node [1] and 979 utils at node [2]. And this despite the fact that Rhoda's credences conditional on E are the same at the two nodes.

This instability in the values of sub-acts explains why information has no fixed value. Learning which member of a partition $\{E_1, E_2, \dots, E_n\}$ is true is good insofar as it might lead the agent to choose a better sub-act on each of the E_i s in the partition, and bad insofar as it might lead her to choose a worse sub-act. If there are no stable facts about the values of these sub-acts, then there are no stable facts about the value of information.

Another consequence of this instability is that REU theorists lose a key strategy for simplifying ‘grand world’ decisions by recasting them as ‘small world’ decisions. This consequence, too, is best illustrated by an example.

Suppose I am deciding whether to buy a frogurt, which looks delicious, but which I fear may be cursed. I’m also not sure whether the frogurt comes with a free topping (that’s good), or whether the free topping contains potassium benzoate, also known as E212 (that’s bad, since E212 may mildly irritate my skin, eyes, and mucous membranes). The following matrix depicts the possible states of the world, with their probabilities noted (columns), the acts available to me (rows) and the utilities of the outcomes in each act-state pair (cells).

	cursed topping E212 ($p = 0.12$)	cursed topping no E212 ($p = 0.12$)	cursed no topping ($p = 0.36$)	not cursed topping E212 ($p = 0.08$)	not cursed topping no E212 ($p = 0.08$)	not cursed no topping ($p = 0.24$)
buy	-50	-30	-45	10	30	15
don't buy	0	0	0	0	0	0

This is a complicated decision problem, with six different states to keep track of. (In fact, a truly accurate portrayal of the decision problem would be vastly more complicated: my mucous membranes might or might react to the E212; the free topping might be delicious chocolate fudge or boring old caramel sauce; my astrologer might help me evade the curse, or fate may catch up with me no matter what I do, and then there are many possible ways to be cursed...) I’d like to reduce the complicated decision problem to a simpler, more coarse-grained decision problem with the following form:

	cursed $p = ???$	not cursed $p = ???$
buy	???	???
don't buy	???	???

But how can I fill in the blank spots in the smaller table? Given a probability function p over fine-grained states, and a utility function u over fine-grained outcomes, how do I generate a probability function p^* over coarse-grained states, and a utility function u^* over coarse-grained outcomes.

The probability function is easy enough. Each s^* in the coarse-grained space of states corresponds to an event $E(s^*)$ in the fine-grained space of states. For instance, the state cursed in the coarse-grained space corresponds to the event $\{\langle \text{cursed, topping, E212} \rangle, \langle \text{cursed, topping, no E212} \rangle, \langle \text{cursed, no topping} \rangle\}$ in the fine-grained space. In general, the probability of a coarse-grained state can be obtained by adding up the probabilities of the fine-grained states that make it up:

$$p^*(s^*) = p(E(s^*))$$

For EU-maximizers, deriving u^* is also simple: The utility of a coarse-grained outcome of act A in state s^* is the expected utility of the sub-act $A_{E(s^*)}$. Furthermore, this expected utility depends only on how things stand within $E(s^*)$: it is a function of the probabilities of states in $E(s^*)$, and the utilities of outcomes that result from A in these states. In the frogurt example, to calculate the utility of buying a cursed frogurt, we take an average of the values of the more fine-grained outcomes that might result from buying a cursed frogurt, weighted by the conditional probabilities of the states that give rise to the outcomes, given that the frogurt is cursed.

$$u^*(\text{buy} \wedge \text{cursed}) = u_{\text{cursed}}(\text{buy})(-50 \times 0.4) + (-30 \times 0.4) + (-45 \times 0.6) = -43$$

In general, where A is an act, s^* is a coarse-grained state, E is the set of fine-grained states corresponding to s^* , p and u are probabilities over the fine-grained space, and p^* and u^* are probabilities and utilities over the coarse-grained space:

$$u^*(A \wedge s^*) = \sum_{s \in E} p(s|E)u(s)$$

Applying this idea to the frogurt example gives us the following table:

	cursed $p = 0.6$	not cursed $p = 0.4$
buy	-43	17
don't buy	0	0

Our REU-maximizer can find ways of filling in the cells in of the small-world problem, assigning $u^*(\text{buy} \wedge \text{cursed})$ and $u^*(\text{buy} \wedge \neg \text{cursed})$ so that

$$REU(\text{buy}) = (0.6)^2[u^*(\text{buy} \wedge \text{cursed}) - u^*(\text{buy} \wedge \neg \text{cursed})]$$

However, these values will not be a function of the probabilities and utilities of states in which the frogurt is cursed. They will also depend on both the probabilities and utilities of the states in which the frogurt is not cursed.

Thus, REU theory makes it more difficult to simplify grand world problems into small world problems. In chapter 6, Buchak gives us an extensive discussion of one way in which this complexity plays out: bets turn out not to have stable values independent of background patterns of risk. I have pointed out another way in which the complexity plays out: sub-acts turn out not to have stable values independent of the larger acts in which they are embedded.

5. Conclusion

Where does this leave the STP? I have argued that all the key premises in the Diachronic Argument are appealing: both naive and sophisticated REU-maximizers will sometimes choose dominated strategies, and that being disposed to choose dominated strategies is constitutive of irrationality. REU

theorists can get out of the argument by endorsing a resolute view of choice, but this amounts to a substantive commitments, with some drawbacks (in particular, at some nodes, resolute choosers will choose acts they disprefer over acts they prefer, simply because of their past preferences). Another avenue for REU theorists to explore is to prescribe changing one's r values in the light of new information.

For REU theorists, sub-acts, like bets, lack context-independent values. Just as the values of bets depend on background distributions of risk, the values of sub-acts depend on the larger acts in which sub-acts are embedded. This means that lotteries cannot be treated like constant acts, for the purpose of simplifying decision problems.

Neither of these points is a knock-down objection to REU theory. In response to the first point, REU theorists can adopt a resolute account of choice, or prescribe changing r values. And in response to the second point, REU theory is a global theory, so in some sense it is no surprise that it should stymie our attempts to break complex decisions down into simple parts. But each point highlights a cost. The real question, then, is whether the costs are worth paying.

Notes

1. More precisely, as long as Strategy 1 weakly dominates Strategy 2, and the set of states in which Strategy 1 yields a strictly better outcome than Strategy 2 is greater than 0, the REU of Strategy 1 will be strictly greater than the REU of strategy 2 at the first node of the decision tree.

Funding

This work was supported by the Stanford University.

Notes on contributors

Rachael Briggs is a professor at Stanford University. Her research interests include decision theory, formal epistemology, and metaphysics.

References

- Allais, Maurice. 1953. "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine." *Econometrica* 21 (4): 503–546.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford: OUP.
- Hammond, Peter J. 1988. "Consequentialist Foundations for Expected Utility." *Theory and Decision* 25 (1): 25–78.
- Machina, Mark. 1989. "Dynamic Consistency and Non-expected Utility Models of Choice Under Uncertainty." *Journal of Economic Literature* 27 (4): 1622–1668.
- McClennen, Edward F. 1990. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Wiley Publications in Statistics.