Interpolating decisions
Jonathan Cohen and Elliott Sober[1]

Abstract: Decision theory requires agents to assign probabilities to states of the world and utilities to the possible outcomes of different actions. When agents commit to having the probabilities and/or utilities in a decision problem defined by objective features of the world, they may find themselves unable to decide which actions maximize expected utility. Decision theory has long recognized that work-around strategies are available in special cases; this is where dominance reasoning, minimax, and maximin play a role. Here we describe a different work around, wherein a rational decision about one decision problem can be reached by "interpolating" information from another problem that the agent believes has already been rationally solved.

1. Introduction

Pascal's wager, or something very like it, can be found in Islamic philosophy centuries before Pascal. In his book *The Alchemy of Happiness*, Abu Hamid Al-Ghazali [1110/1909: 71] begins his version of the argument by saying that it should not be wasted on anyone who is *certain* of God's nonexistence, for "the case of such a man is hopeless; all one can do is to leave him alone." Instead Al-Ghazali aims his argument at someone who thinks that God and an afterlife are at least possible. He starts with several real-life examples involving risk. These include a one-rupee charm that is purported to cure a terrible illness and a warning that your food may have been contaminated with snake venom. In each of these examples, the cost is small (paying one-rupee or forgoing a bit of food), but the benefit that might be gained is momentous – saving your life! Al-Ghazali then connects these mundane examples with the question of whether to believe in God. He writes: "People take perilous voyages in ships for the sake of merely probable profit. Will you not suffer a little pain of abstinence now for the sake of eternal joy hereafter?" He concludes that even if you are "doubtful about a future existence, reason suggests that you should act as if there were one, considering the tremendous issues at stake." Al-Gazali's argument deploys a strategy that we think is familiar; to answer the agent's question of whether to believe in God, Al-Gazali interpolates information from other decision problems in which the agent thinks it's clear which action is best among a set of alternatives. Here we describe in more detail how interpolation works.

2. Interpolation

Standard decision theory is based on the assumption that what it's rational for an agent to do in a given decision problem is determined by those of the agent's probabilities and utilities that pertain to the actions and states of the world considered . This might be called a "locality"

---

assumption.[2] These probabilities and utilities dictate what the expected utilities are for different possible actions, and then the rule of maximizing expected utilities kicks in and dictates what the agent ought rationally to do.[3] Theorists have described a handful of strategies for extending this standard procedure to settings in which agents don't know point values for probabilities and utilities; these workaround strategies include appeals to dominance, maximin, and minimax rules. Here we describe a different technique for extending the applicability of standard decision theory in which agents use information about their previous decisions, and the belief that those previous decisions were rational, to solve a current decision problem.[4] We expect that this pattern of reasoning will be familiar to readers, once it is described.

Here's an example of interpolation. You discover that you have a suboptimal mix of good and bad cholesterol, and, on the advice of your physician, you consider taking drug A. You consider the relevant costs, benefits, and probabilities involved, and decide to take the drug. A year later, your physician tells you that there is a new drug on the market, drug B, that is relevant to a different health problem that you have. Flu season is about to begin, and drug B is a flu vaccine. Your doctor informs you that drug B has fewer health costs and more health benefits

---

[2] In his discussion of decision-making with vague probabilities, Weatherson [ms.] endorses a decision rule in which decision-making concerns sets of actions, not singletons. Our motivation for violating locality is different. The issue of locality also comes up in the theory of estimation developed in statistical decision theory in response to Stein's paradox; see [Vassend, Sober, and Fitelson 2017].

[3] As we understand the framework, a commitment to evaluating actions in terms of expected utilities is compatible with not knowing the values of the quantities one would need in order to compute expected utilities for candidate actions, and, therefore, without knowing what this commitment entails about concrete situations; we return to this point in §3 below. We note also that the rule of maximizing expected utility has been challenged on various grounds; in this paper, we use that rule without attempting to justify it. However, in §7, we describe how interpolation can be used by agents who reject the rule of maximizing expected utility because they think that variance in utility is relevant. Finally, our subject is also separate from Rubinstein's [1988] use of similarity considerations to make progress on solving the Allais Paradox, which paradox is a threat to the rule of maximizing expected utility.

[4] Interpolation is similar in some ways to a different proposal for assessing the rationality of new actions by comparing them to old ones, due to Gilboa and Schmeidler [1995, 2001]. On their proposal, one derives point values for the utilities and probabilities of world states relevant to assessing a novel action A from those of old actions, together with a stipulated (and undefined) "similarity" measure between A and those other actions, and then feeds the resulting determinate point values into the standard framework. However, comparisons of actions enter the discussion in a different place for us. Our contention is that comparative information can be informative in assessing A's rationality even in the absence of point values for the utilities of outcomes and the probabilities of world states relevant to assessing A (*a fortiori* in cases where they are not derived from comparisons, per Gilboa and Schmeidler).

A referee suggested that we use the term 'analogy' for what we are calling 'interpolation' and drew our attention to Shafer's [1986] proposal that "an analysis of a decision problem by subjective expected utility is merely an argument, an argument that compares that decision problem to the decision problem of a gambler in a pure game of chance. This argument by analogy may or may not be cogent. In some cases other arguments are more cogent." Our interpolation differs from Shafer's analogy; what matters for us is a system of inequalities that connects an old decision problem to a new one. It doesn't matter to us how analogous the two problems are. For us, the two problems may concern entirely different possible actions, states of the world, and expected utilities, and there is no need for either of them to be about a gambler thinking about a game of chance.

than drug A, and drug B has a higher probability than drug A had of being safe; she also says that there is no interaction between the two medications. You reason as follows:

> Taking drug B is rational, if I was rational a year ago to take drug A.
> I was rational a year ago to take drug A.
> _____
> I now am rational to take drug B.

In this argument, you use the assumption that you were rational a year ago to take drug A to solve your present decision problem about whether to take drug B.

Table 1: Deciding whether to take drug A			
		Possible states of the world (and their probabilities)	
		Drug A is safe (p)	Drug A is not safe (1-p)
Possible Actions	Take drug A	$x+b_1$	$x-c_1$
	Don't take drug A	x	x

Your initial decision to take drug A can be analysed in more detail by considering the utilities and probabilities represented in Table 1. These details illustrate why we think the term 'interpolation' is apposite. From the entries in this table, you can compute the expected utilities of the two actions as follows:

> EU(take drug A) = $p(x+b_1) + (1-p)(x-c_1)$
> EU(don't take drug A) = x.

Here 'x' denotes the utility of maintaining the *status quo* – of not taking drug A. From these expected utilities, the following criterion can be derived (assuming that c_1 and b_1 are both positive and p<1):

(CRIT$_A$)   EU(take drug A) > EU(don't take drug A) precisely when $\frac{p}{(1-p)} > \frac{c_1}{b_1}$.

You need information about costs, benefits, and probabilities to see whether the inequality on the right-hand side of the CRIT$_A$ biconditional is satisfied. Notice that CRIT$_A$ conveniently separates probabilities from utilities; the former are on the left-hand side of an inequality, while the latter are on the right [Ramsey 1931].

Table 2: Deciding whether to take drug B	
	Possible states of the world (and their probabilities)

		Drug B is safe (q)	Drug B is not safe (1-q)
Possible	Take drug B	$y+b_2$	$y-c_2$
Actions	Don't take drug B	y	y

After you begin taking drug A, you confront a new question—whether you should take drug B. The relevant quantities are displayed in Table 2. As before, 'y' denotes the utility of maintaining the *status quo* – of not taking drug B. Your new problem involves the following criterion (assuming that c_2 and b_2 are both positive and that $q<1$):

(CRIT$_B$) EU(take drug B) > EU(don't take drug B) precisely when $\frac{q}{(1-q)} > \frac{c_2}{b_2}$.

You may or may not know the probabilities, costs, and benefits mentioned in CRIT$_B$, but there is another way to figure out whether to take drug B. Suppose you now believe that your taking drug A was rational, which entails the truth of

(SOURCE) $\frac{p}{(1-p)} > \frac{c_1}{b_1}$.

You can use this inequality to help solve your new decision problem. If

(CONJ) $\frac{q}{(1-q)} \geq \frac{p}{(1-p)}$ and $\frac{c_1}{b_1} \geq \frac{c_2}{b_2}$,

you should take drug B. Four points about this form of reasoning bear emphasis.

First, notice that CONJ (so-called because it is a conjunctive criterion) doesn't require you to compare the ratio of probabilities and the ratio of utilities displayed in CRIT$_B$. According to CONJ, you just need to compare one ratio of probabilities with another (which simplifies to the question of whether $q \geq p$) and compare one cost/benefit ratio with another. If you know that the probability of drug B's being safe is greater than the probability that drug A is safe ($q > p$), and that drug B has a smaller cost/benefit ratio than drug A ($\frac{c_1}{b_1} > \frac{c_2}{b_2}$), then you should take drug B.

Second, notice that the conclusion reached about drug B rests on two components: the conditional premise that if CONJ holds you should take drug B, and the separate, unconditional premise SOURCE. Committing to interpolating from your prior decision about drug A to a new decision about drug B, as such, says nothing whatsoever about the rationality of the decision about drug A. In this respect, interpolation-reasoning resembles Bayes's Rule, which is a procedure for updating your doxastic state given your prior probabilities, rather than a guide to which prior probabilities should be accepted [Urbach and Howson 1993].[5] (The resemblance is only partial, however, in that interpolation doesn't involve the assignment of prior probabilities.)

---

[5] Of course, accepting this reasoning cannot foreclose the possibility of acquiring new evidence that might influence your view about the unconditional premise SOURCE—including the very observation that, given your other commitments, accepting SOURCE would license the conclusion that you should take drug B. As the saying goes, one time-slice's *modus ponens* is another's *modus tollens*.

Third, though in this example a decision problem at t_2 is solved by interpolating into it information about a decision made at t_1, for some $t_2 > t_1$, there's nothing essentially diachronic about interpolation. This form of reasoning is equally good in cases where $t_2 = t_1$, and even when $t_2 < t_1$.[6]

Fourth, though this example involves a decision made by a single individual, interpolation can also be useful in interpersonal settings. For example, if *you* are trying to convince me to reach a decision (say, to take drug B), it may be helpful to point out to me that this decision can be reached by interpolating a solution to a different decision problem that *I* regard as rational (say, to take drug A). (This is, of course, an extremely common everyday rhetorical strategy.) That the comparison between the new and the old strategy comes from an individual numerically distinct from the one making the new decision changes nothing about the decision procedure. Another interpersonal home for interpolation lies in cases of collective decision making. If a committee faces a decision problem (say, whether to recommend drug B for some population), it may be helpful for the committee to notice that this decision can be reached by interpolating information from a solution to a different decision problem it has already come to (say, to recommend drug A for the same population). This is just to say that interpolation can be employed by groups as well as by individuals.

3. Can interpolation be helpful in practice?

One might worry that the alternative decision strategy offered here, involving the use of CONJ and SOURCE, is not a helpful advance over the standard strategy involving $CRIT_B$. Evaluating CONJ appears to require values for p, q, b_1, c_1, b_2, and c_2; but if one knows those values, one could apply $CRIT_B$ directly. How, then, could having the alternative strategy possibly be helpful?

The principal difference between using $CRIT_B$ and using CONJ together with SOURCE is that $CRIT_B$ involves comparing a ratio of probabilities with a cost/benefit ratio, whereas SOURCE requires comparing probabilities with each other and cost/benefit ratios with each other. We suggest that there can be circumstances in which agents find the former comparison difficult and the latter one easy, and that reasoning with CONJ and SOURCE will be advantageous in such cases. Of course, the frequency with which such cases arise in the actual world is an empirical psychological matter, separate from our normative claim that interpolation is a rational strategy. Note that we are assuming that maximizing expected utility is itself rational (see footnote 3).

---

[6] Interpolation is different from the controversial idea that it is a requirement of rationality that agents ought to abide by commitments they made earlier [McClennen 1985]. For example, does the promise you made at time t_1 rationally bind you to abide by that promise at time t_2? We take no stand on this issue. Notice that in our account of interpolation, it isn't a promise made at t_1 that undergirds your decision at t_2 to perform an action (like taking drug B) that licenses interpolating, but the belief that an action performed at t_1 maximized expected utility. Our account of interpolation is compatible with (but does not require) a narrowly instrumentalist conception of rationality that is ethically unconstrained.

We concede that there is no need for the interpolation procedure in the idealized setting where decision theory is often located. Here probabilities and utilities are subjective, and agents are assumed to have precise introspective access to their quantitative values. However, for agents who are committed to basing their probabilities and utility assignments on objective information, or for whom the mind is not an open book, lack of information is a real possibility, and workarounds are helpful when seemingly vital items of information are not at hand. This provides an opening for interpolation to be useful as well as valid.

To illustrate the opening just described, consider a person who commits to the goal of having her subjective probability that drug B is safe be equal to the actual frequency of people who have no bad side effects among the population of people who have taken the drug. Suppose further that this person commits to using enhanced life expectancy as her measure of the benefit that would come to her if she takes the drug and it is safe. Interpolation can be useful to this person by providing her with a tool for deciding whether to take drug B (in a way that respects her doxastic commitment) when she doesn't have that frequency information at hand, and also lacks information on enhanced life expectancy.

4. More than one interpolation source

In the example described above, your question of whether to take drug B has one interpolation source, namely your earlier decision to take drug A. However, it's easy to imagine situations in which an agent has several such sources. Having multiple sources might strengthen your conviction that you should take the new drug; it might also weaken that conviction. Strengthening occurs if interpolating from numerous sources points to taking drug B; if even one of those earlier decisions was rational, you have an interpolation argument, and if all of them are rational, your decision about taking drug B is robustly guided by what you take to be numerous prior rational decisions. Of course, having multiple sources for *negative* interpolation will robustly indicate *not* taking drug B. And if distinct and equally (objectively) rational instances of interpolation pull in opposite directions, the collection may amount to conflicting evidence that, taken as a whole, fails to give any clear guidance about the new decision.

5. Comparability: Apples and oranges

We hope that the interpolation strategy described in section 1 will see familiar, plausible—even obvious—once considered. We have observed informally that it is widely enlisted as a rhetorical strategy for motivating/defending everyday decisions. However, when applied in such ordinary cases, it invites a worry about incomparability between decisions.

The two decision problems about drug A and drug B from section 1 are similar in many respects. Both concern whether a drug is safe, and both concern the health benefits of taking a pill if it is safe and the health costs of taking a pill if it is unsafe. The similarity between the two cases encourages the thought that the costs and benefits in the two decisions are straightforwardly comparable, as apparently required by CONJ. But there are certainly pairs of decisions with respect to which the comparability of costs and benefits is less obvious. Indeed, with a more detailed description of the two problems from section 1 (the first problem concerning a better balance of good and bad cholesterol, the second involving your chance of getting the flu), one might bring the comparability of costs and benefits in this pair of cases into question. Whatever one makes of that particular pair, we can ask more generally: can

interpolation be useful in pairs of cases involving apples and oranges, where the costs and benefits in the interpolating problem look to fall along substantially different dimensions from the costs and benefits in the interpolated problem?

Table 3: Deciding whether to watch movie M			
		Possible states of the world (and their probabilities)	
		Movie M is enjoyable (r)	Movie M is not enjoyable (1-r)
Possible	Watch movie M	$z+b_3$	$z-c_3$
Actions	Don't watch movie M	z	z

Though, in general, this will depend on how much interdimensional aggregation one is prepared to accept, we believe there are situations where information from a decision about apples can be interpolated into a decision about oranges.[7] To see this, consider a case in which you are deciding whether to watch movie M this evening. The probabilities and utilities of this problem are represented in Table 3. The expected utilities of the two actions are

$EU(\text{watch movie M}) = r(x+b_3) + (1-r)(x-c_3)$
$EU(\text{don't watch movie M}) = x,$

from which it follows that

$(\text{CRIT}_{\text{movie}})$ $EU(\text{watch movie M}) > EU(\text{don't watch movie M})$ precisely when $\frac{r}{(1-r)} > \frac{c_3}{b_3}$.

Suppose you do the calculation, since you have numbers for the quantities described in $\text{CRIT}_{\text{movie}}$, and this leads you to watch the movie. Weeks later, you are considering whether to take pill A (where the details of this decision problem are described in Table 1). You can use your decision to watch movie M as a source for interpolation. Since you now believe that it was rational to watch the movie weeks earlier, you know that

$(\text{SOURCE})\ \frac{r}{(1-r)} > \frac{c_3}{b_3}$.

The unconditional premise SOURCE can be used to decide whether to take pill A if you allow enough comparability to accept

(CONJ*) $\frac{p}{(1-p)} \geq \frac{r}{(1-r)}$ and $\frac{c_3}{b_3} \geq \frac{c_1}{b_1}$.

Accepting CONJ* means that you think (i) the probability that pill A is safe (p) is greater than the probability that watching movie M will be enjoyable (r), and (ii) that the cost/benefit ratio of

---

[7] We do not intend what we say here to be an *account* of comparability (aggregation): we use the term as placeholder for whatever it is in the psychology of agents that permits them to use the interpolation strategy we have described (cf. the undefined similarity function in Gilboa's and Schmeidler's [1995] comparisons of new and old actions, discussed in footnote 4.)

watching the movie is greater than the cost/benefit ratio of taking pill A.[8] If so, then interpolating from the prior decision about the movie gives you a *pro tanto* reason to take pill A.

This example shows that apples-and-oranges interpolation is sometimes possible. It does not show that it is always possible. One's verdict about the applicability of interpolation to particular pairs of cases will presumably turn on one's views about the comparability of the cost/benefit ratios occurring in conjunctive propositions like CONJ*. (For what it is worth, this consequence of our description of interpolation seems to us to reflect reasonably accurately our own reactions to cases; this provides modest support for the approach.)

6. Relaxing two simplifications

So far, we've discussed interpolation by considering two decision problems in which there are two states of the world and two possible actions, and in which one of those actions has the same utility in both states of the world. Both these simplifications are dispensable.

Table 4: a 3x3 decision problem				
		States of the world (and their probabilities)		
		State 1 (p)	State 2 (q)	State 3 (r)
Possible Actions	A1	u_{11}	$u_{\cdot 2}$	$u_{\cdot 3}$
	A2	$u_{\cdot 1}$	u_{22}	$u_{\cdot 3}$
	A3	$u_{\cdot 1}$	$u_{\cdot 2}$	u_{33}

Consider the decision problem depicted in Table 4, where neither of the above simplifications is deployed. We assume that for each state of the world, one of the actions is best, and the two inferior actions in that state of the world are equal in utility. For convenience, suppose that Ai is the best action when the world is in state i (i = 1,2,3). Here's a two-part criterion for A1 to be the best of the three actions:

$$EU(A1) > EU(A2) \text{ precisely when } p/q > \frac{u_{22} - u_{\cdot 2}}{u_{11} - u_{\cdot 1}} \text{ and}$$
$$EU(A1) > EU(A3) \text{ precisely when } p/r > \frac{u_{33} - u_{\cdot 3}}{u_{11} - u_{\cdot 1}}$$

Here terms of the form '$u_{ii} - u_{\cdot 1}$' represent the gain in utility that accrues from doing the best action in state of the world i, rather than one that is inferior.

If A1 is in fact the best of the three, we have our source for interpolation:

$$(\text{3-ACT SOURCE}) \quad p/q > \frac{u_{22} - u_{\cdot 2}}{u_{11} - u_{\cdot 1}} \text{ and } p/r > \frac{u_{33} - u_{\cdot 3}}{u_{11} - u_{\cdot 1}}$$

In this case, the agent knows the point values of all four of the ratios in the above conjunction.

Table 5: a 3x3 decision problem	
	States of the world (and their probabilities)

---

[8]Note that one might be in a position to endorse the second conjunct of $CONJ_1$ even when one is not committed to the pairwise comparability of c_1 with c_3, or b_1 with b_3.

		State 4 (w)	State 5 (x)	State 6 (y)
Possible Actions	D1	v_{11}	$v_{\cdot 2}$	$v_{\cdot 3}$
	D2	$v_{\cdot 1}$	v_{22}	$v_{\cdot 3}$
	D3	$v_{\cdot 1}$	$v_{\cdot 2}$	v_{33}

The agent then confronts a new decision problem that concerns three new actions and three new states of the world, as shown in Table 5. D1 is the best of these three actions precisely when

$$EU(D1) > EU(D2) \text{ precisely when } w/x > \frac{v_{22} - v_{\cdot 2}}{v_{11} - v_{\cdot 1}} \text{ and}$$

$$EU(D1) > EU(D3) \text{ precisely when } w/y > \frac{v_{33} - v_{\cdot 3}}{v_{11} - v_{\cdot 1}}$$

The unconditional premise 3-ACT SOURCE helps you decide whether D1 is best. If you know that

$$w/x > p/q \text{ and that } \frac{u_{22} - u_{\cdot 2}}{u_{11} - u_{\cdot 1}} > \frac{v_{22} - v_{\cdot 2}}{v_{11} - v_{\cdot 1}} \text{ and that}$$

$$w/y > p/r \text{ and that } > \frac{u_{33} - u_{\cdot 3}}{u_{11} - u_{\cdot 1}} > \frac{v_{33} - v_{\cdot 3}}{v_{11} - v_{\cdot 1}} \text{ ,}$$

you can conclude that D1 is the best action to perform in the second 3x3 problem.

This reasoning can be extended to nxm problems involving comparisons between n distinct possible actions in m states of the world. As such, interpolation is a general and powerful technique for thinking about rational decisions. To be fair, interpolation is obviously more complicated in the 3x3 problem (and in yet larger problems) than it was in the 2x2 problem, and this may make it less useful as a practical device for dealing with lack of information concerning the point values of probabilities and utilities. However, in terms of the normative theory, these problems are on the same footing as far as interpolation is concerned.

7. Beyond maximizing expected utility

We have assumed for the sake of argument that an agent's decision is rational precisely when it maximizes expected utility. However, we think that interpolation can be rationally valuable even when that assumption is relaxed. The cases we have in mind are ones in which rational agents take into account both the variance and the expected utilities of the actions they are considering.[9]

---

[9] The following two bets make the contrast between variance in utility and expected utility clear:

Wager 1: win $0 with probability ½ and win $10 with probability ½.
Wager 2: win $4 with probability ½ and win $6 with probability ½.

The two wagers have the same expected utility (= $5), but Wager 1 has the greater variance, meaning the average deviation from the expected value. For much more on this topic, cf. Hacking [2001], Buchak [2013].

Decision making where variance and expected value both matter

Variance

B

C

A

Expected Utility

Consider the three actions depicted in the accompanying figure. The rule of maximizing expected utility entails that rational agents should be indifferent in the choice between actions A and B. It is not obvious that this recommendation is correct. Plausibly, it is rationally permissible to choose A over B; reduction in variance can be a tiebreaker. The choice between A and C is more complicated for agents who value variance reduction; here rational agents need to embrace a policy that describes how increasing expected utility should trade off against reducing variance. Here again, it is plausibly rational permissible to adopt a policy that solves this problem. We doubt that there is a uniquely rational policy of this kind, and we take no stand on the question of whether some policies are beyond the pale. Likewise, we leave it open that risk-takers can be rational just like those who are risk averse.

The point of relevance here is that the strategy of interpolation can be rationally valuable both to agents who ruthlessly maximize expected utility and to agents who take both variance and expected utility into account: a rational agent can benefit from interpolating information about an old decision that they think was rational into a new decision problem they face if they care about expected utility alone, or about both variance and expected utility.

8. Other interpolation problems: Calibration
Our protocol for interpolation reasoning involves attention to both cost/benefit ratios and to probabilities. Interestingly, however, there is a special type of interpolation reasoning that is useful in cases in which there is no uncertainty about the state of the world.

One type of case with this structure can arise when new currencies are introduced into extant systems of exchange, leaving consumers unable to interpret price signals as they normally would. Consider, for example, the introduction of the URV (the Unidade Real de Valor, or Unit of Real Value) in Brazil in 1994, as part of a plan to solve the country's longstanding hyperinflation problem. The Brazilian government introduced the URV not by releasing newly

denominated bills and coins, but by insisting that incomes, taxes, and prices be described in URV (as well as cruzeiros) across the country.[10]

Suppose you are a Brazilian encountering URV-denominated prices for the first time in a restaurant; suppose further, for the sake of the example, that you don't have access to either the cruzeiro-denominated price list or the day's official cruzeiro-URV conversion rate.[11] You see a sandwich priced at 1 URV. Should you strike the deal? Presumably you should if the benefits exceed the cost. However, given your epistemic situation, it's hard for you to make sense of the cost associated with the transaction proposed. Crucially, your difficulty does not stem from uncertainty about the probabilities of possible outcomes; it comes, rather, from your inability to assess costs in a meaningful way. Even so, you can constrain your decision about this transaction by comparing it to other actual or possible transactions on offer. Perhaps you observe that the neighbouring market offers a quart of milk at 1 URV. You know that the objective benefits of a quart of milk are no greater than those of the contemplated sandwich. (Assume you are otherwise indifferent between the sandwich and the milk, that the cost of travel between restaurant and market is negligible, etc.) Comparing these two offers/possible decisions, you reason that the decision to buy the restaurant's sandwich at 1 URV is rational if the decision to buy the market's milk at 1 URV is rational. Suppose you additionally believe that the market's offer is fair (the proprietor is an old friend who is always generous with customers), hence that deciding to buy its quart of milk at 1 URV is rational. Then you should regard the restaurant's proposed transaction involving the sandwich (whose cost/benefit ratio is equal to that of the proposed transaction in the market you've already deemed rational) as rational as well.

In contrast, suppose you know that the neighbouring hotel asks 1 URV for *two* sandwiches—i.e., the hotel offers twice the benefits available in the restaurant, for the same monetary cost. (Assume you are indifferent between the individual sandwiches, that the costs of travel between restaurant and hotel is negligible, etc.). Comparing the possible decision in the hotel against the proposed transaction in the restaurant, you reason that a decision to buy the restaurant's sandwich at 1 URV is irrational if the decision to buy the hotel's two sandwiches at 1 URV is irrational. Suppose you know that the hotel's prices are always set at irrationally high objective cost/benefit ratios so as to prey on unsuspecting tourists. Then you should regard the proposed transaction at the restaurant (whose cost/benefit ratio is twice that of the proposed transaction at the hotel you've already deemed irrationally high) as irrational as well.

---

[10] The intention was that URV-denominated prices would remain stable (the URV was fixed at 1-1 parity with the US dollar) even as the cruzeiro continued to inflate, that this would make Brazilians accustomed to thinking about prices in URV, and so would prepare them to make use of the new stable currency once it became non-virtual (whereupon it was aptly named the 'Real') on 1 July 1994. See https://en.wikipedia.org/wiki/Plano_Real.
For a popular account of the episode, see https://www.npr.org/sections/money/2010/10/01/130267274/the-friday-podcast-how-four-drinking-buddies-saved-brazil.
[11] Either of these pieces of information could be used to interpolate into the current decision in a different way— viz., each would allow you to interpolate information from decisions about actual or possible cruziero-denominated exchanges (which you might more confidently assess as rational or not) into your current decision about the proposed URV-denominated.

Here is a second example involving the use of interpolation for calibration. Suppose you hope to incentivize service in your academic department by distributing points toward teaching relief to faculty in exchange for carrying out various administrative tasks (graduate placement, running the colloquium series, refereeing submissions to the undergraduate journal, etc.). Once again, you face a calibration problem: you must fix how much teaching relief credit is earned by fulfilling each task. Unfortunately, there is no antecedently published conversion rate to rely on. However, you can make headway on your problem by interpolation. Suppose it is a matter of longstanding practice in your department (predating the introduction of the new and more general points system) that the Director of Graduate Studies always earns one course of teaching relief per year. Then you can fix the benefits earned by other service roles by comparison to that unit value. If the role of graduate placement director is (say) one-half as onerous as that of the DGS job, then a decision to assign .5 courses per year of teaching relief to the placement director is rational if the decision to assign 1 course per year of teaching relief to the DGS was rational.

9. Conclusion

The conditional rule we've used to describe interpolation has more going for it than its intuitive plausibility. It is entailed by a dominance rule that is universally endorsed in theories of how to reason about imprecise probabilities. This rule says that you should choose action X over action Y if on every probability function p consistent with your evidence, the p-expectation of the utility of X is greater than that of Y. This rule licenses interpolation. If you are deciding whether to take drug A and you don't have enough information to form precise probabilities and utilities, but you are able to judge that the probability ratio $\frac{p}{1-p}$ is greater than the utility ratio $\frac{c_1}{b_1}$, this fact translates into a constraint on the set of probability and utility functions that represent your imprecise state of mind. Taking drug A is rational, in this framework, because its expected utility is higher than that of not taking it on every pair of probability and utility functions in this set. That is, taking the drug EU-dominates not taking it. If you then face the decision about whether to take drug B and you judge that $q > p$ and that $\frac{c_1}{b_1} > \frac{c_2}{b_2}$, that judgment implies further constraints on the set of functions representing your state of mind; these constraints entail that taking drug B EU-dominates not taking it. In other words, if your judgmental state meets the set of constraints concerning drug A, then taking drug B is rational if your judgmental state meets the second set of constraints as well. In this sense, it is possible to view interpolation reasoning as outlining a way of applying the dominance rule to an interesting class of cases.[12]

Decision theory describes how agents should decide whether to perform a given action if they know the costs, benefits, and probabilities that pertain to that action. However, when agents commit themselves to thinking about some or all of these three considerations as objective quantities, ignorance about some or all of them becomes possible. In that situation, familiar workarounds (dominance reasoning, the minimax rule, etc.) become useful; we suggest that interpolation can be added to this list. Interpolating information from a decision problem that the agent thinks she has already solved rationally into the decision problem at hand provides a rational way forward.

---

[12] We are grateful to Richard Bradley for pointing this out to us.

References

Al-Ghazali, Abu Hamid Muhammad 1111/1909. *The Alchemy of Happiness.* translated from the Persian by Claud Field. London: John Murray.

Buchak, Lara 2013. *Risk and Rationality*. Oxford: Oxford University Press.

Gilboa, Itzhak and David Schmeidler 1995. Case-Based Decision Theory, *The Quarterly Journal of Economics*, 110/3: 605-639.

Gilboa, Itzhak and David Schmeidler 2001. *A Theory of Case-Based Decisions*. Cambridge: Cambridge University Press.

Hacking, Ian 2001. *An Introduction to Probability and Inductive Logic*. Cambridge: Cambridge University Press.

McClennen, Edward 1985. Prisoner's Dilemma and Resolute Choice, in *Paradoxes of Rationality and Cooperation: Prisoner's Dilemma and Newcomb's Problem*, eds. Richmond Campbell and Lanning Sowden, Vancouver: University of British Columbia Press: 94–104.

Ramsey, Frank 1931. Truth and Probability, in *The Foundations of Mathematics and Other Logical Essays*, ed. R.B. Braithwaite, London: Kegan Paul, Trench, Trubner, & Co.: 156-198.

Rubinstein, Ariel 1988. Similarity and Decision-Making Under Risk (Is there a Utility Theory Resolution to the Allais Paradox?), *Journal of Economic Theory* 46: 145-153.

Shafer, Glenn 1987. Savage Revisited, *Statistical Science* 1/4: 463-485.

Urbach, Peter and Colin Howson 1993. *Scientific Reasoning: The Bayesian Approach*. Peru, Illinois: Open Court.

Vassend, Olav, Elliott Sober, and Branden Fitelson 2017. The Philosophical Significance of Stein's Paradox, *European Journal for the Philosophy of Science* 7/3: 411-433.

Weatherson, Brian ms. Decision making with imprecise probabilities, http://brian.weatherson.org/vdt.pdf

Institutional affiliations: Department of Philosophy, University of California, San Diego;
Department of Philosophy, University of Wisconsin–Madison

tables 1-5:

Table 1: Deciding whether to take drug A			
		Possible states of the world (and their probabilities)	
		Drug A is safe (p)	Drug A is not safe (1-p)
Possible Actions	Take drug A	$x+b_1$	$x-c_1$
	Don't take drug A	x	x

Table 2: Deciding whether to take drug B			
		Possible states of the world (and their probabilities)	
		Drug B is safe (q)	Drug B is not safe (1-q)
Possible Actions	Take drug B	$y+b_2$	$y-c_2$
	Don't take drug B	y	y

Table 3: Deciding whether to watch movie M			
		Possible states of the world (and their probabilities)	
		Movie M is enjoyable (r)	Movie M is not enjoyable (1-r)
Possible Actions	Watch movie M	$z+b_3$	$z-c_3$
	Don't watch movie M	z	z

Table 4: a 3x3 decision problem				
		States of the world (and their probabilities)		
		State 1 (p)	State 2 (q)	State 3 (r)
Possible Actions	A1	u_{11}	$u_{\cdot 2}$	$u_{\cdot 3}$
	A2	$u_{\cdot 1}$	u_{22}	$u_{\cdot 3}$
	A3	$u_{\cdot 1}$	$u_{\cdot 2}$	u_{33}

Table 5: a 3x3 decision problem				
		States of the world (and their probabilities)		
		State 4 (w)	State 5 (x)	State 6 (y)
Possible Actions	D1	v_{11}	$v_{\cdot 2}$	$v_{\cdot 3}$
	D2	$v_{\cdot 1}$	v_{22}	$v_{\cdot 3}$
	D3	$v_{\cdot 1}$	$v_{\cdot 2}$	v_{33}

Figure caption: Decision making where variance and expected value both matter.