

Is risk aversion irrational?

Examining the “fallacy” of large numbers

H. Orri Stefánsson

Received: date / Accepted: date

Abstract A moderately risk averse person may turn down a 50/50 gamble that either results in her winning \$200 or losing \$100. Such behaviour seems rational if, for instance, the pain of losing \$100 is felt more strongly than the joy of winning \$200. The aim of this paper is to examine an influential argument that some have interpreted as showing that such moderate risk aversion is irrational. After presenting an axiomatic argument that I take to be the strongest case for the claim that moderate risk aversion is irrational, I show that it essentially depends on an assumption that those who think that risk aversion can be rational should be skeptical of. Hence, I conclude that risk aversion need not be irrational.

Keywords Risk aversion · Expected utility · Event-Wise Dominance

1 Introduction

There are various types of bets such that although one is not willing to accept a single bet of that type, one might be willing to accept a package consisting of a number of such bets. Suppose that someone makes you the following offer. A symmetric coin is tossed and you get to choose a side. If that side comes up you win \$200, but if the other side comes up you lose \$100. Moderately risk averse¹ people would generally not take the bet: the 50% chance of gaining \$200 is simply not worth a 50% risk of losing \$100. But what if you were

University of Copenhagen
Karen Blixens Vej 4, 2300 Copenhagen, Tel.: +46764281782
E-mail: orri.stefansson@hum.ku.dk

¹ In what follows, I will stick to the convention of calling a person who turns down a bet with a positive monetary expectation *risk averse* w.r.t. money (in the relevant interval). Such betting behaviour can be caused either by money having decreasing marginal utility for the person or by attitudes to risks and chances that are independent of attitudes to quantities of money (Stefánsson and Bradley 2018). Although some (e.g. Hansson 1988) would argue that the behaviour should only be called ‘risk averse’ when it is caused by the

offered a package of such bets on multiple coin tosses (a bet on each toss)? If the number of bets in the package is great enough, then since each bet has a positive monetary expectation (and since probabilistic independence holds between the coin tosses), the package might seem to be worth the risk, even though a single bet from the package would not. For instance, a package consisting of 100 bets of this kind has an expectation of \$5,000 and comes with only a 1/2,300 chance (about 0.04%) of leading to a loss of money. Daniel Kahneman is in no doubt about what one should do if offered such a package of bets: “the offer of a hundred [bets] is so attractive that no sane person would reject it” (2012: 337).

Paul Samuelson (1963) famously claimed that the above reasoning is a fallacy. He called it ‘The fallacy of large numbers’, since he thought that it consisted in a misapplication of the law of large numbers. The law entails that a package consisting of an *infinite* number of bets of the above kind is *guaranteed* to result in a payoff equal to its expectation, and that the probability that the payoff equals the expectation approaches certainty as the number of bets approaches infinity. But as Samuelson rightly pointed out, the 100 bet package has a (tiny) chance of resulting in a huge loss, which it would be unreasonable to completely ignore. Moreover, Samuelson proved that if an *expected utility maximiser* turns down a single bet of a particular kind, and would moreover do so for all wealth levels, then she will not accept a package consisting of multiple (but finite²) independent bets of that kind. Therefore, if expected utility theory is the correct theory of practical rationality, it follows that one cannot rationally turn down the single bet for all levels of wealth while accepting the package.

It is worth emphasising that although one might never be offered a coin-toss-bet like the one under discussion, the above conclusion—if sound—has far-reaching practical implications. A ‘bet’, as I am using the term, can be any prospect that has some chance of the ‘bettor’ getting a good result and some chance of her getting a bad result. Thus understood, one is very often faced with ‘bets’ with positive monetary expectations. For instance, suppose that you have just bought a laptop and are considering whether to buy a two year warranty. Such warranties generally have negative monetary expectations for buyers—otherwise sellers wouldn’t offer them! So the ‘bet’ *not* to buy the warranty has a positive monetary expectation (compared to the alternative of buying a warranty). When deciding whether to buy the warranty, one might consider whether one would buy a warranty if one were buying 100 laptops (say as a firm). If we suppose that the probability of each laptop breaking down is independent of another laptop breaking down,³ I conjecture that most people would *not* buy the warranty if they were buying 100 laptops, since the ‘wins’ of

second type of attitude mentioned above, I will, for the sake of simplicity, call any person who turns down a bet with a positive monetary expectation risk averse w.r.t. money.

² From now on, whenever I talk about packages of bets, I will assume that they consist of finite numbers of bets.

³ This assumption is reasonable if we assume thorough quality control, in which case a computer braking down could be treated as a random event.

this package (i.e., of *not* buying warranty for the 100 computers) will almost certainly outweigh the ‘losses’. But then, by the above argument, it might seem that one should not buy warranty for the single laptop either.

Let us for the sake of simplicity however focus on Samuelson’s coin-toss-bets. Since a package of 100 such bets seems so attractive, some have taken Samuelson’s argument to show that it would simply be irrational to turn down the single bet (see, for instance, Benartzi and Thaler 1999). But the preference against accepting the single bet can be entirely due to what seems to be moderate aversion to risk. So if the lesson from Samuelson’s argument is that it would be irrational to turn down the single bet, then it seems we have to conclude that moderate risk aversion is irrational. Others have taken Samuelson’s argument to be a *reductio* of the decision theory he assumed: expected utility (EU) theory (see e.g. Lopes 1996, Rabin 2000). Still others think that there must be something wrong with Samuelson’s argument, since they think that there is neither anything wrong with turning down the single bet while accepting the package nor with EU theory as a theory of practical rationality (e.g. Aloysius 2007).

The aim of this paper is to examine more carefully than has previously been done the argument that one cannot rationally turn down the single bet while accepting the package—or, as I shall call it, the argument for “Samuelson’s claim”⁴—and the more general claim that moderate risk aversion is irrational. After formulating an axiomatic argument that I take to be the strongest case for Samuelson’s claim, I show that the argument essentially depends on the principle of *Event-Wise Dominance*, which those who think that risk aversion can be rational have a good reason to reject. Hence, I conclude, the argument is question begging if intended as a general argument against the possibility of rational risk aversion.⁵

2 Samuelson’s argument

The starting point of Samuelson’s discussion was the observation that one of his colleagues at MIT—a “distinguished scholar”, he says—did in fact turn down Samuelson’s offer of a ‘coin-toss-bet’⁶ like the one discussed above, while at the same time confessing his willingness to accept a package consisting of 100 such bets. Samuelson’s colleague, whom I will henceforth refer to as SC, justified his attitude by saying that although the single bet was not enough to make it “reasonably sure that the law of averages will turn out in my favor”,

⁴ As a referee correctly points out, it is not clear whether Samuelson (1963) was concerned with showing that turning down the single bet while accepting the package would be *irrational*. Rather, he seems to have been more concerned with showing that such behaviour would be inconsistent with expected utility maximisation. Nevertheless, for the sake of simplicity, and since Samuelson did endorse expected utility theory in much of his writings, I will occasionally refer to the claim in question as “Samuelson’s claim”.

⁵ It should be noted that this is not how Samuelson himself interpreted his result.

⁶ In what follows, a ‘coin-toss-bet’ is a bet that pays bettors \$200 if their chosen side of a coin comes up on a particular toss but has them paying \$100 otherwise.

the law of large numbers would make a package consisting of 100 such bets very attractive. “I am, so to speak, virtually sure to come out ahead in such a sequence, and that is why I accept the sequence while rejecting the single [bet]” (Samuelson, 1963: 51).

To examine this pair of preferences, Samuelson assumed that his colleague would still turn down an offer of the single bet if his wealth increased or decreased by the maximum amount that he could gain (\$20,000) or lose (\$10,000) from the package of 100 bets⁷ (an innocent assumption, he thought, given what distinguished MIT scholars typically earn).⁸ On the basis of this assumption, Samuelson proved that if his colleague were an expected utility maximiser, as utility is typically defined,⁹ then he would not accept the package consisting of 100 bets, since he did not accept the single bet.

The logic behind Samuelson’s argument is very simple (1963: 53):

If you will not accept one toss, you cannot accept two—since the latter could be thought of as consisting of the (unwise) decision to accept one plus the open decision to accept a second. Even if you were stuck with the first outcome, you would cut your further (utility) losses and refuse the terminal throw. By extending the reasoning from 2 to $3 = 2 + 1$, ... , and from $n-1$ to n , we rule out any sequence [i.e., any package of bets] at all.

In other words, the conclusion is that since SC preferred his status quo to accepting the single coin-toss-bet, and assuming that he would continue to do so for all wealth levels, he should prefer his status quo to any package of such bets.

Samuelson’s (informal) illustration of his argument is rather unfortunate, since it makes it seem as if the argument depends on the assumption that his colleague would turn down the *first* coin-toss-bet in a sequence of such bets. But foreseeing the sequence, the colleague might accept the first bet. As Lara Buchak (2013: 219) puts it:

By assuming that an agent will reject [the first bet], the argument ignores the fact that just as the act of taking a bet cannot be fully described, in the framework of decision theory, without knowing an agent’s current holdings—which bets an agent accepted in the past—the act of taking a bet also cannot be fully described without knowing what choices an agent will be offered in the future.

⁷ Strictly speaking, the argument only requires the assumption that SC would turn down the single bet for any level of wealth in the range that the first 99 bets could result in. Hence, when this assumption is formalised, as Assumption 1 in the axiomatic argument presented in section 3, the stated range is from a loss of \$9,900 to a gain of \$19,800.

⁸ You can test this assumption on yourself, assuming that you share SC’s preference against the single bet but for the package, by considering whether your preference would change if you had, for instance, paid \$10,000 less or \$20,000 more off your mortgage, car loan or student loan.

⁹ That is, as a function of terminal wealth, in this case (or more generally as a function of terminal outcomes), rather than, say, as a function of changes in wealth.

The questionable assumption is, however, by no means indispensable to the argument. Rather than starting the inductive argument from the first bet, we can start from the last bet and work backwards. We are supposing that SC is, on the one hand, evaluating a single bet, and, on the other hand, a package of 100 bets. Now under the assumption that SC would turn down the single bet for the wealth range in question, it is clear that if offered the 100 bets sequentially, he would turn down the last bet, irrespective of the outcome from the first 99 bets (i.e., irrespective of how the first 99 bets affect his wealth). For recall that it is assumed that SC has a preference against a single bet whenever his wealth is in the range that could result from these 99 bets. But if he knows that he will turn down the last bet, then he will also turn down the second-to-last bet, and so on, for all bets in the package. Hence, it may seem that Samuelson's claim is true, since his colleague's attitudes do seem to be internally incoherent. A more formal version of this argument will be provided in the next section.

This version of the argument might raise the worry that while SC was considering, on one hand, a single bet, and, on the other hand, a package of bets, the above argument has turned the decision-problem SC was faced with into a sequential one with the option to stop after having accepted any number of bets (Aloysius 2007). In particular, it might seem that Samuelson is falsely assuming that one can rationally accept a package of bets made at a single time just in case one accepts a sequence of the same bets. However, as we shall see in the next section, when the argument is axiomatised, a dominance principle, which is entailed by expected utility theory, ensures that attitudes to the elements of the sequence constrain attitudes to packages of bets in such a way that the backward-inductive argument does show SC to be irrational, on the assumption that the dominance principle in question is a requirement of rationality.

Some have taken Samuelson's argument to be *reductio* of expected utility theory as a theory of practical rationality.¹⁰ The single bet seems quite risky and doesn't really have the potential of a significant reward. However, the package—which, recall, only has a 1/2,300 chance of resulting in monetary loss and could result in a win of up to \$20,000—doesn't seem that risky and has the potential of a large reward. So, it seems quite sensible to turn down the single bet while accepting the package. Lola A. Lopes (1996: 184) goes so far as confessing that:

one of the strongest arguments I can find against the principle of expected utility maximization is that it considers such an eminently sensible decision [i.e., the decision to turn down a single coin-toss-bet while accepting a package of 100 such bets] to be irrational.

Mathew Rabin is even more blunt when describing how damaging he takes Samuelson's result to be for EU theory, calling the implications of the theory

¹⁰ It should be noted, however, that Thoma (ta.) argues that most alternatives to expected utility theory will also have great trouble accommodating SC's preferences, for reasons that I shall briefly discuss in the concluding section.

that Samuelson derived “ludicrous”. And he contends that the reason Samuelson’s conclusion seems counterintuitive to many people is not that they have poor intuitions about near-zero probabilities, as Samuelson suggested. Rather, it “is counterintuitive because it is crazy” (2000: 206).

However, as we shall see in the next section, Samuelson’s conclusion follows from assumptions that are strictly weaker than those of expected utility theory. But it does not take much to get the whole EU theory once one grants a crucial assumption of the argument: the principle of Event-Wise Dominance. Moreover, those who reject EU theory due to its treatment of risk, could take the argument of this paper to illustrate that it is this dominance principle, rather than the strictly (but only slightly) stronger *Sure Thing Principle* (Savage 1954), that is the reason why the theory does badly in this respect. More on this in section 4, after the axiomatic argument has been presented in the next section.

3 The argument axiomatised

Tversky and Bar-Hillel (1983) state an axiomatic version of Samuelson’s argument. Their formulation can however be criticised for the reasons raised by Buchak (as quoted above). The axiomatic backward-inductive argument I formulate in this section avoids that problem, and contains some details that Tversky and Bar-Hillel ignore. As we shall see, Tversky and Bar-Hillel’s axiomatic argument was somewhat misleading, since it might have suggested that the (mostly uncontroversial) State-Wise Dominance principle sufficed for the argument, when the (much more controversial) Event-Wise Dominance principle is actually needed.

The axiomatic argument presented below is, as far as I can tell, the strongest argument one could give for what I previously called “Samuelson’s claim”, i.e., for the claim that one cannot rationally turn down the single bet while accepting the package of 100 bets. But it should be noted that this argument shares with other backward-inductive arguments an (implicit) assumption that could be questioned; namely, the assumption that the agent of interest is certain that at each point in time, she will act on her preferences at that time. Since the focus of this paper is on the rational permissibility of risk aversion, I shall not discuss the assumption in question any further (but see e.g. McClennen 1990 for a thorough discussion). Instead, the focus will be on a premise in the argument that turns out to specifically have to do with rationally permissible attitudes to risk.

The first assumption of the axiomatic argument is based on the observation that SC would turn down the last coin-toss-bet he expects to be offered, provided that previous bets did not take his wealth outside the specified range. And crucially, the same holds when SC is considering a particular bet if he is certain that he will not *accept* any further coin-toss-bets. Therefore, he would turn down the 99th bet, having realised that he will not accept the 100th (i.e., the last) bet. However, rather than stating the assumption by referring to bets

SC expects to accept, I will state it more generally as an assumption that SC turns down any single bet when his wealth is within a particular range; but I will first apply the assumption to the last bet and work backwards to the first bet.

To state the assumption more formally, let B_i denote an individual coin-toss-bet; i.e., a bet that either results in bettors winning \$200 or losing \$100 depending on whether their chosen side of a coin comes up or not. The index denotes the order in which the bet is offered, such that, for instance, B_{i+1} is offered next after B_i . $\{B_n\}$ is a package of n such bets. ' \succ ' represents SC's preference, such that $A \succ B$ denotes that SC prefers A to B . X is a random variable in the interval $[w - 9, 900, w + 19, 800]$, where $\$w$ is the dollar amount corresponding to SC's pre-bet wealth. $\$X + B_i$ denotes that SC's wealth is X dollars and that he in addition holds a coin-toss-bet.

The first assumption can now be stated as follows:

Assumption 1 *The agent prefers to turn down bet B_i for all wealth levels between and including $\$w - 9.9k$ and $\$w + 19.80k$. More formally, for any value of X in $[w - 9, 900, w + 19, 800]$:*

$$\$X \succ \$X + B_i$$

The above assumption is an empirical claim. The assumption is that SC's preferences do as a matter of fact have this characteristic. And as discussed, this is clearly true when SC is certain that he will not accept any bet that is offered after B_i .

The next assumption, DOMINANCE, is however a normative claim. The role of the axiom in the present argument is to connect rational attitudes to outcomes of sequences of bets to attitudes to packages of bets. Informally, the axiom, in the form used here, states that if for any possible outcome that a package $\{B_n\}$ of bets can result in, a decision-maker prefers that outcome to that same outcome plus an additional bet B_{n+1} , then the decision-maker prefers $\{B_n\}$ to the package $\{B_{n+1}\}$ that is the result of enlarging $\{B_n\}$ by bet B_{n+1} .

Before stating a general version of DOMINANCE, it may be instructive to consider a simplified version of the principle where the two packages in question are, on one hand, the 'package' consisting only of bet B_1 , and, on the other hand, the package consisting of bets B_1 and B_2 . Suppose that B_1 will result either in outcome x if a coin comes heads up on its first toss—an event I will refer to as H_1 —or in outcome y if the coin comes tails up on its first toss—an event I will refer to as T_1 . (I assume that $T_1 = \neg H_1$.) The outcome of B_2 , however, depends on how the coin lands on the second toss, which is assumed to be independent of how it lands on the first toss. Table 1 represents the two alternatives on offer—i.e., to take the first bet only and to take both bets—where the state space has been partitioned according to how the coin lands on its first toss.

DOMINANCE entails that if x is preferred to $x + B_2$ and y is preferred to $y + B_2$, then the first bet only—i.e., package $\{B_1\}$ —is preferred to the

	H_1	T_1
$\{B_1\}$	x	y
$\{B_2\}$	$x + B_2$	$y + B_2$

Table 1 Take one bet or two?

package consisting of the first and the second bet—i.e., package $\{B_2\}$. And that may seem plausible, as a condition on rational attitudes to bets; and could be seen as a rationality constraint on how sequential preferences in the case under consideration relate to ‘static’ preferences between packages. If no matter what outcome the first bet results in, the agent prefers to not add the second bet, then it might seem that she has no good reason for preferring the combination of the first and the second bet to the first bet only. (As will become apparent, I think this justification of DOMINANCE is mistaken. But it certainly has some intuitive appeal.)

We can state the axiom more formally as follows:

Assumption 2 (Dominance) *If for all possible outcomes X from $\{B_i\}$, $X \succ X + B_j$, then:*

$$\{B_i\} \succ \{B_i\} + B_j = \{B_{i+1}\}$$

As we have seen, Assumption 1 clearly holds when $B_i = B_{100}$ (i.e., when B_i is the last bet on offer). Together with DOMINANCE the assumption entails:

$$\{B_{99}\} \succ \{B_{100}\} \tag{1}$$

In other words, from Assumption 1 and DOMINANCE, it follows that even if the agent has accepted a package of 99 bets, she prefers not to increase the size of the package by one additional bet. But recall that we are assuming that the agent is certain that she will act on her preferences at each point in time. Hence, from (1) it follows that Assumption 1 also holds when $B_i = B_{99}$. So, we have:

$$\{B_{98}\} \succ \{B_{99}\} \tag{2}$$

By iterating this reasoning:

$$\begin{aligned} \{B_{97}\} &\succ \{B_{98}\} \\ &\vdots \\ \{B_0\} &\succ \{B_1\} \end{aligned}$$

To connect this sequence of pair-wise comparisons, we only need to assume that the preference is *transitive*. I will not provide any argument for this assumption, which is sometimes justified on the grounds that a person who violates it is vulnerable to being ‘money pumped’.

Assumption 3 (Transitivity) *For any alternatives A , B , and C :*

$$\text{If } A \succ B \text{ and } B \succ C \text{ then } A \succ C$$

Transitivity entails, in the presence of Assumptions 1 and 2, that the agent prefers no bet to the package consisting of 100 bets (and, indeed, to any package):

Conclusion: If a person satisfies Assumptions 1 to 3, then she will turn down a package consisting of up to and including 100 coin-toss-bets, since $\{B_0\} \succ \{B_n\}$ whenever $1 \geq n \leq 100$.

So if Transitivity and DOMINANCE are requirements of rationality, and given that Assumption 1 is, as a matter of fact, true of Samuelson's colleague, it would seem that Samuelson's colleague is indeed irrational.

4 Event-Wise and State-Wise Dominance

I will, for the sake of the argument, grant that Transitivity is a requirement of rationality. But what about DOMINANCE? As previously mentioned, Tversky and Bar-Hillel provide a similar axiomatisation of Samuelson's argument to the one provided above. Moreover, they suggest that the axiomatic formulation shows that the argument depends only on assumptions that, first, are weaker than those of expected utility theory, and, second, seem "unobjectionable on normative grounds" (Tversky and Bar-Hillel 1983: 716).

In response to Tversky and Bar-Hillel's first claim, it should be pointed out that while it is, strictly speaking, true that the argument depends on assumptions that are strictly weaker than those of EU theory, it nevertheless contains an assumption, DOMINANCE, that I think is not—and certainly *should not be*—accepted by those who reject EU theory due to its treatment of risk aversion.¹¹ Following Tversky and Bar-Hillel, I have, so far, not distinguished between *State-Wise Dominance* and *Event-Wise Dominance*. As I will explain below, DOMINANCE is an instance of the latter; and, as is well known, one has come a long way towards an expected utility representation once one assumes Event-Wise Dominance. But more importantly, unlike State-Wise Dominance, Event-Wise Dominance is far from being unobjectionable as a requirement of rationality, and will (as I will explain) presumably be rejected by anyone who rejects the (very controversial) *Sure Thing Principle*.

As is conventional, I will use the term 'state' for a maximally specific contingency that leaves no uncertainty as to how a prospect will result. An 'event', on the other hand, is a union of states, meaning that an event may

¹¹ Moreover, there are reasons for rejecting DOMINANCE that have nothing to do with risk aversion, as a referee points out; in particular, as Dietrich and List (2005) show, the so-called *Two Envelope Paradox* calls the principle into question, as a general requirement of rationality.

not fully determine the result of a risky prospect. State-Wise Dominance can be stated as follows:¹²

State-Wise Dominance 1 (SWD) *For any two bets B_i and B_j , if you weakly prefer B_i to B_j conditional on every possible state of the world, then you should weakly prefer B_i to B_j . If, in addition, there is some state conditional on which you strictly prefer B_i to B_j , then you should strictly prefer B_i to B_j .*

Event-Wise Dominance, on the other hand, says that:

Event-Wise Dominance 1 (EWD) *For any two bets B_i and B_j , if there is some partition of the state-space into events such that you weakly prefer B_i to B_j conditional on every event in that partition, then you should weakly prefer B_i to B_j . If, in addition, there is some non-empty event in the partition, conditional on which you strictly prefer B_i to B_j , then you should strictly prefer B_i to B_j .*

EWD is stronger than SWD. If B_i is weakly preferred to B_j conditional to each event in the maximally fine partition of the state space—i.e., the partition where each ‘event’ is a state—then the first bet should, according to EWD, be weakly preferred to the latter. So EWD entails SWD. But the converse does not hold, since SWD only says that dominance reasoning should be used on the maximally fine partition.

DOMINANCE is an instance of EWD where the partition that generates the dominance relationship concerns how the coin lands on the first i tosses. However, SWD does not suffice for Samuelson’s argument, since there are some states of the worlds conditional on which Samuelson’s colleague would clearly prefer to add an additional bet; namely, any state of the world where that extra bet is a winning bet. So SWD does not entail that SC should, given his preference against the single bet, prefer a smaller package to a larger one. In other words, SC’s preferences are not irrational, if we accept SWD but not EWD as a requirement of rationality.

While Tversky and Bar-Hillel are right, of course, that we do not get expected utility theory from EWD and Transitivity, we do come very close. For notice that EWD implies the Sure Thing Principle (STP)—the cornerstone of Savage’s (1954) influential version of expected utility theory—if we simply assume that the relation of preference-conditional-on-event is complete.¹³ This is important since the STP is the focus of the main disagreement between EU

¹² Recall that by ‘bet’ I simply mean a risky prospect that will turn out favourably to the decision-maker in some event but unfavourably otherwise.

¹³ Here is a simple demonstration. First a formal statement of the Sure Thing Principle. Suppose we have four prospects, call them f , g , f' and g' , let $f(E)$ denote the outcome received when f is chosen in event E , and suppose: $f(\neg E) = g(\neg E)$, $f'(\neg E) = g'(\neg E)$, $f(E) = f'(E)$, $g(E) = g'(E)$. Then the STP says that: $f \succsim g \Leftrightarrow f' \succsim g'$. To derive the STP from EWD, suppose first that $f \succsim g$. Now by reflexivity of the indifference relation over outcomes, we have $f(\neg E) \sim g(\neg E)$, $f'(\neg E) \sim g'(\neg E)$, $f(E) \sim f'(E)$, $g(E) \sim g'(E)$. By completeness of the conditional preference relation, we have either $f(E) \succsim g(E)$ or $g(E) \succsim f(E)$ (or both). Now by EWD, at least the first of these *must* hold, because

theorists and those who think that (Savage-style) EU theory does not adequately capture rational attitudes to risk (see e.g. Buchak 2013). Hence, while the axiomatic version of Samuelson’s argument relies on assumptions that are strictly weaker than those of EU theory, it nevertheless contains an assumption that comes very close to the typical object of discontent amongst those who dislike EU theory.

Moreover, those who think that risk aversion can be rational—and, in particular, those who think that expected utility theory is too demanding when it comes to what attitudes to risk it allows—have special reasons for objecting to Event-Wise Dominance.¹⁴ As an illustration, consider a choice between A and B in Table 2 (which is similar to two of the lotteries in Allais’ 1953 paradox), where the prize of each bet depends on which ticket, numbered 1 to 100, is drawn.

	1	2 – 34	35 – 100
A	\$0	\$3500	\$0
B	\$2400	\$2400	\$0

Table 2 Half Allais choice problem

Now let E be the event that one of tickets 1 to 34 is drawn, $\neg E$ the event that one of tickets 35 to 100 is drawn. This means that we have a partition of the state space into two events such that in one event (E), the first bet has a much worse minimum prize than the second, whereas in the second event ($\neg E$), both bets have the same outcome, which equals the minimum prize of the first bet in the first event. Then a risk averse person might prefer the second bet given E : in this event A has a much worse minimum prize, and while A has a higher expectation than B (assuming that the draw is random), the risk averse might not be willing to risk ending up with \$0 when she can guarantee herself \$2400. But since the two bets have the same outcome in $\neg E$, one should of course be indifferent between the bets given this event. So, a risk averse person might (at least) weakly prefer B to A given E and also given $\neg E$. Nevertheless, she might strictly prefer A to B when she does not know whether E or $\neg E$ occurs, since they then have the same minimum prize, which might make her willing to incur the slightly higher risk of ending up with nothing by going for the bet with the considerably higher expectation.¹⁵

otherwise we would have $g \prec f$. But then $f'(E) \lesssim g'(E)$, which by EWD entails $f' \lesssim g'$. Assuming $f' \lesssim g'$ would have allowed us to derive $f \lesssim g$ in a similar manner. Hence, EWD and completeness of the conditional preference relation entail the STP.

¹⁴ Critics of expected utility theory are faced with well-known dynamic arguments in favour of the theory theory, that can in fact be formulated as arguments in favour of Event-Wise Dominance. For a detailed discussion of ways in which those who are critical of expected utility theory can respond to these arguments, see McClennen (1990).

¹⁵ Another related way in which a person might be risk averse, is if she, in addition to caring about the expected values of prospects, also cares about the dispersion in the monetary values of the possible outcomes (McClennen 1990). Again, such a person might prefer B to A conditional on E (since the former has minimal dispersion but the latter has quite a lot),

A more general worry about EWD, that can be illustrated by again considering the bets in Table 2, concerns the demand that if there is *any* partition of the state space into events such that B is weakly preferred to A conditional on every event, then B should be preferred to A . Notice that there is a refinement of the $\{E, \neg E\}$ partition according to which the dominance relation no longer holds: further partition E according to whether ticket 1 is drawn (event F) or not ($\neg F$). Then it is no longer the case that B event-wise dominates A , since the latter is preferred conditional on event $E \wedge \neg F$. Recall that the same point holds for the axiomatic version of Samuelson’s argument: the smaller package does not event-wise dominate the larger one once the partition is fine enough to determine the outcome of the extra bet that differentiates the large package from the smaller one. So, one might reasonably wonder why the $\{E, \neg E\}$ partition should be decisive when it comes to determining the rational preference between A and B , rather than, say, the maximally fine partition. Similarly, in the case of Samuelson’s colleague’s preferences, one might wonder why the partition according to how the coin lands on the first i tosses should be decisive.

The problem with EWD raised in the preceding paragraph shows that what seems to be the main intuition underlying SWD does not justify the claim that violations of EWD are irrational. It seems that the reason for preferring one prospect over another has to be found in how the two prospects turn out in some state of the world. So, one prospect should be preferred to another only if there is some possible way the world might be (i.e., some state of the world) such that the one prospect turns out better than the other. This principle does not quite deliver SWD (which makes the stronger claim about when one alternative should be preferred to another), but it cannot be violated by someone who satisfies SWD. If E and $\neg E$ in the above example (similarly for the events in the discussion of SC’s preference) were states of the world, then the risk averse decision-maker discussed above (and SC) would be violating this very plausible principle. That is, if E and $\neg E$ were two states of the world, then it would seem that since B is weakly preferred to A conditional on both E and $\neg E$, there could be no reason to prefer A to B . But E and $\neg E$ are not maximally specific, and despite B event-wise dominating A with respect to the $\{E, \neg E\}$ partition, there *is* a state where A turns out better than B , namely, any state in event $E \wedge \neg F$. Correspondingly, there is *a* reason for preferring A over B —in particular, that A turns out better in any state in $E \wedge \neg F$ —despite B event-wise dominating A relative to the $\{E, \neg E\}$ partition. More generally, it is false that when a bet event-wise dominates another, one can have *no* reason for preferring the dominated bet.

In sum, Event-Wise Dominance is clearly *not* “unobjectionable on normative grounds”, contrary to what Tversky and Bar-Hillel (1983: 716) claim. For if the type of risk-aversion often displayed in the Allais paradox is rationally permissible, as many people argue (e.g. Buchak 2013, Stefánsson 2016, and

but would be indifferent between the two bets conditional on $\neg E$; but nevertheless prefer A to B unconditionally, since both bets have quite a lot of dispersion, but the former has a considerably higher monetary expectation.

Bradley and Stefánsson 2017), then EWD is not rationally required. Hence, since the strongest case for Samuelson's claim that one cannot rationally accept a package of bets of a particular kind while turning down a single bet of that kind essentially depends on EWD, the claim should be rejected by anyone who thinks that risk aversion is rationally permissible. More generally, Samuelson's argument does not show risk aversion to be irrational.

5 Concluding remarks

The axiomatic argument presented above shows that if both Transitivity and Event-Wise Dominance are requirements of rationality, then one cannot rationally turn down a single coin-toss-bet for all levels of wealth while accepting some package of such bets. So, if one is to vindicate a preference like that of Samuelson's colleague—and, it would seem, if one wants to allow for moderate risk aversion—one has to give up either Transitivity or Event-Wise Dominance. However, those who think that risk aversion can be rational—and, in particular, those who think that the orthodox treatment of risk attitudes in decision theory and economics is flawed—should be skeptical of EWD anyway, for reasons I have discussed.

In fact, theories that have been developed as alternatives to expected utility theory, in particular to allow for a wider range of risk attitudes than those for which EU theory can account, typically allow for violations of EWD (see e.g. Kahneman and Tversky 1979, Buchak 2013, Bradley and Stefánsson 2017). And these theories are indeed consistent with preference patterns like those reported by Samuelson's colleague. However, since these theories allow for violations of EWD—and, in particular, since they allow agents to evaluate bets differently depending on whether they are embedded in a larger bet or not—applications of them can have different implications depending on how a decision-problem is framed.

In particular, as Thoma (ta.) stresses, theories that allow for violations of EWD may recommend the rejection of a bet when it is formulated as a one-off decision, while recommending the acceptance of the same bet when it is formulated as part of, say, the bundle of bets that the decision-maker in question will be faced with for the rest of her life. In contrast, theories that include EWD will deliver the same verdict irrespective of which of these two ways of viewing a bet is adopted, and can thus be safely combined with a 'short-term' perspective. So, a complete vindication of preferences like those of Samuelson's colleague would have to show that he is justified in both violating EWD *and* in taking the short-term perspective. The latter is however beyond the scope of this paper.

Acknowledgements Versions of this paper were presented at the Final Conference of the Franco-Swedish Program for Philosophy and Economics, SCAS Uppsala (2015), and Self Prediction in Decision Theory and Artificial Intelligence, University of Cambridge (2015). I am grateful to the audience for their questions, comments and suggestions. In addition, I would like to thank Richard Bradley and Katie Steele for helpful comments and suggestions.

References

1. 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.
2. Aloysius, John A. (2007), "Decision making in the short and long run: Repeated gambles and rationality," *British Journal of Mathematical and Statistical Psychology*, 60(1): 61-69.
3. Benartzi, Shlomo and Richard H. Thaler (1999), "Risk aversion or myopia? Choices in repeated gambles and retirement investments," *Management Science* 45(3): 364-381.
4. Bradley, Richard, and H. Orri Stefánsson (2017), "Counterfactual Desirability," *British Journal for the Philosophy of Science*, 68(2): 485-533.
5. Buchak, Lara (2013), *Risk and Rationality*, Oxford University Press.
6. Dietrich, Frans and Christian List (2005), "The two-envelope paradox: An axiomatic approach", *Mind* 114(454): 239-248.
7. Hansson, Bengt (1988), "Risk Aversion as a Problem of Conjoint Measurement," in Peter Gärdenfors and Nils-Eric Sahlin, *Decision, Probability and Utility*, Cambridge University Press, pp. 136-158
8. Jeffrey, Richard (1965), *The Logic of Decision*, University of Chicago Press.
9. Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47(2): 263-291.
10. Kahneman, Daniel (2012), *Thinking, Fast and Slow*, Penguin Books.
11. Lopes, Lola A. (1996), "When time is of the essence: Averaging, aspiration, and the short run," *Organizational Behavior and Human Decision Processes*, 65(3): 179-189.
12. McClennen, Edward F. (1990), *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press.
13. Rabin, Matthew (2000), "Diminishing marginal utility of wealth cannot explain risk aversion," in Daniel Kahneman, Amos Tversky (eds.) *Choices, Values and Frames*, Cambridge University Press, pp. 202-208.
14. Samuelson, Paul (1963), "Risk and uncertainty: A fallacy of large numbers," *Scientia* 98: 108-113.
15. Savage, Leonard (1954), *The Foundations of Statistics*, John Wiley & Sons.
16. Stefánsson, H. Orri (2016), "Desirability of conditionals", *Synthese* 193(6): 1967-1981.
17. Stefánsson, H. Orri and Richard Bradley (2018), "What is risk aversion?", *The British Journal for the Philosophy of Science*, in print (<https://doi.org/10.1093/bjps/axx035>).
18. Thoma, Johanna (ta.), "Risk aversion and the long run", *Ethics*, in print (<http://eprints.lse.ac.uk/88363/>).
19. Tversky, Amos and Maya Bar-Hillel (1983), "Risk: The long and the short", *Journal of Experimental Psychology* 9(4): 713-717.