

Risk Aversion and the Long Run*

** I am very grateful to Sergio Tenenbaum, Jonathan Weisberg, Joseph Heath, Julia Nefsky, Michael Bratman, R.A. Briggs, Amelia Hicks, Luc Bovens, Richard Bradley and three anonymous referees for feedback on earlier drafts of this paper. I also thank audiences at the University of Konstanz, the Workshop on the Ethics of Social Risk in Montreal, the London School of Economics, the Varieties of Agency Workshop at Stanford University, the Meeting of the Society for Exact Philosophy in Miami, the Saint Louis Annual Conference on Reasons and Rationality, the Cambridge Philosophy of Science Seminar, and the London Judgment and Decision Making Group at University College London for valuable discussion and suggestions.*

Abstract: This paper argues that Lara Buchak’s risk-weighted expected utility (REU) theory fails to offer a true alternative to expected utility theory. Under commonly held assumptions about dynamic choice and the framing of decision problems, rational agents are guided by their attitudes to temporally extended courses of action. If so, REU theory makes approximately the same recommendations as expected utility theory. Being more permissive about dynamic choice or framing, however, undermines the theory’s claim to capturing a steady choice disposition in the face of risk. I argue this poses a challenge to alternatives to expected utility theory more generally.

1. Introduction

In our lives we face a long series of choices with uncertain consequences. We decide on careers, partners or on what city to live in not knowing if they will be right for us. Most everyday decisions also have uncertain consequences. On a smaller scale, we make decisions on which charity to give to not knowing for sure where our money will do the most good, on whether to buy concert tickets not knowing if we’ll be able to find a date, and on whether to go to the concert not knowing if we’ll like the music.

We face these choices repeatedly. We also often display risk aversion when we make these decisions. Even when only small stakes are involved, we often go for a safe option when another one would have had a higher expected payoff. Risk aversion for the small stakes gambles we face repeatedly in our lives, including the ethically significant ones,¹ is common and seems sensible. At the same time, we often do not display the same risk aversion when we decide on a number of gambles together in a larger stakes gamble, and that seems equally sensible. Yet, the orthodox theory of choice under uncertainty, expected utility theory, does not seem to be able to account for this pattern of behaviour. This not only calls into question this prominent theory of practical rationality. Insofar as expected utility theory is incorporated in a variety of ethical theories, most notably subjective consequentialism and contractarianism, challenges to expected utility theory may call for revisions to those theories, too.

Alternatives to expected utility theory promise to offer a superior treatment of everyday risk aversion, and could potentially find fruitful application in ethical theorizing.² I here argue that one prominent such alternative, rank-dependent utility theory in the form of Buchak's risk-weighted expected utility (REU) theory³ ultimately fails to do so. The core problem is that once an agent realizes that any small stakes risky choice currently under her consideration is just one of a large number of similar choices she faces in her life, then under some plausible additional assumptions, REU theory itself seems to recommend not being risk averse. I will argue that this in fact poses a challenge to alternatives to expected utility theory more generally.

As normative theories of rational choice, formal decision theories purport to tell us what it means to make rational decisions when facing risky choices. In particular, these theories are presented as theories of instrumental rationality, that is, theories about what it means to take the best means to one's ends,

¹ For evidence on small-scale risk aversion in charitable giving, for instance, see Lata Gangadharan, Glen Harrison, and Anke Leroux, "Preferences over Risky Aid," *Monash University Working Paper* (2015).

² In fact, Buchak has already applied her theory to decision-making behind the Rawlsian veil of ignorance in her "Taking Risks behind the Veil of Ignorance," *Ethics* 127 (2017): 610–644.

³ Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2013).

whatever they may be. At the same time, these theories also aim to be explanatory of the choices rational agents make. However, formal decision theories can be applied only once the choices we face have been described in a certain way, namely as formal decision problems. Formal decision theories describe how we (rationally) solve such decision problems. But the choices we face do not come neatly packaged into formal decision problems.

The most common answer to the question of how we should specify decision problems is that we should include in our description of the decision problem everything that is relevant to the decision at hand, where a factor is relevant when its inclusion in the decision problem could change the agent's preferences over which act to choose. If this is our notion of relevance, then what is relevant will in part depend on the nature of our decision theory. In the case of REU theory, that is the root of the problem I wish to point out. According to REU theory, and under the theory of dynamic choice defended by many critics of expected utility theory, namely resolute choice, any future risky choice an agent expects to take is relevant for the evaluation of a present gamble. But when an agent takes into account all the future risky choices she expects to face when formulating her decision problem, REU theory does not allow for risk aversion for each small stakes gamble anymore.

The resulting problem for REU theory is this: REU theory is introduced as a vindication of the ordinary decision-maker. But the ordinary decision-maker displays small stakes risk aversion. REU theory can only account for her preferences insofar as she specified her decision problem more narrowly than seems practically rational. REU theory thus seems to fail to rationally vindicate her behaviour. Moreover, since most of the examples that motivate REU theory involve risk aversion for small stakes gambles, REU theory seems to fail to make sense of the examples that motivated it in the first place.

REU theory can avoid this conclusion by being more permissive about how agents formulate their decision problems, and by relaxing the assumption that agents need to be strictly resolute. I argue, however, that even if such permissiveness is compatible with rationality, this would call into question the purpose of REU theory as a normative and explanatory theory of choice under uncertainty. Such

permissiveness severely limits the extent to which REU theory can be action-guiding. Moreover, insofar as we are interested in REU theory as an explanatory theory, such permissiveness also means that the predictions of REU theory end up being very sensitive to the way in which agents choose to frame their decision problems, or choose to act in dynamic choice contexts. This means that in the kinds of cases centrally motivating REU theory, the formalism cannot justifiably capture any steady choice tendencies in the face of risk, which limits the explanatory value of the theory.

2. Samuelson's Colleague

This paper is concerned with decision theories as normative and explanatory theories of practical rationality. Expected utility theory is a label applied to a number of theories about how to make one-off decisions under conditions of uncertainty, that is, when we are unsure about what the consequences of each of our potential acts are.⁴ Expected utility theory instructs agents to choose an act with maximum expected utility. To determine which acts maximize expected utility, we must be able to assign a probability $p(o)$ to the n possible outcomes o of each act available. We then also assign a utility $u(o)$ to each possible outcome. These are then weighted by their probability to arrive at the expected utility associated with the act.⁵ Critics of expected utility theory argue that expected utility theory does not offer a satisfactory treatment of risk aversion. A common way of doing so is to present examples of preferences that are both common and do not seem intuitively irrational, and to argue that expected utility theory

⁴ 'Uncertainty' is here used in a wide sense and is meant to include both what is commonly called 'risk' (where objective probabilities are known to an agent) and 'uncertainty' (where they are not).

⁵ This follows the version of expected utility theory presented in John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton: Princeton University Press, 1944), which will suffice for our purposes. Formally, act a 's utility is given by:

$$EU(a) = \sum_{i=1}^n p(o_i) \cdot u(o_i)$$

cannot accommodate these preferences. Examples of such preferences are those displayed in the Allais and Ellsberg Paradoxes.⁶ This section will introduce another such case.

Paul Samuelson reports having had the following lunch-time conversation with a colleague, who claimed to be a proponent of expected utility theory.⁷ Samuelson asked his colleague, henceforth SC, to consider a bet which gave him a 50% chance of winning \$200 and a 50% chance of losing \$100. SC refused on the grounds that this bet is too risky for him. However, he added that he would accept a series of one hundred such bets, since in the series of bets, he is almost certain to come out ahead. And indeed, the chance of making a loss on the series is below 0.5%, while the expected return is \$5,000. The series of bets thus seems like an extremely attractive offer. At the same time, SC's preference to reject the individual gamble merely seems to display a reasonable level of risk aversion. Samuelson shows that expected utility theory cannot easily make sense of these preferences, and thus seems to imply that SC is irrational. At least this is so when SC's decision problem is framed in the way it is by Samuelson.

Samuelson takes the relevant outcomes in SC's case to be adequately described by the different amounts of total wealth SC may reach in the gambles he is offered. He also assumes that SC takes the probabilities he is given at face value. If he is an expected utility maximizer, he thus maximizes the sum of his utilities over wealth levels weighted by these probabilities. Given this characterization of the problem, Samuelson shows that SC's preferences are incompatible with expected utility theory, under the plausible assumption that SC would prefer to reject each individual gamble at any wealth level that he might reach in the series of gambles (that is, at any wealth level between \$9,900 below and \$19,800 above his current wealth). In fact, for any gamble, and any number of repetitions, given this assumption, it is irrational to reject a single gamble that one would accept as part of a compound of many such gambles. What Samuelson shows is

⁶ Maurice Allais, "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine," *Econometrica* 21 (1953): 503–546, and Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics* 75 (1961): 643–669.

⁷ Paul Samuelson, "Risk and Uncertainty: a Fallacy of Large Numbers," *Scientia* 98 (1963): 108–113.

that in expected utility theory, under the assumptions he makes, gambles that are unfair in the utility metric, that is, have a negative expectation of utility, cannot be made fair by compounding them.⁸

I will not reproduce Samuelson's proof here. What is important for us is the feature of expected utility theory that is driving this and similar results.⁹ And that is that in expected utility theory, risk averse behaviour, such as turning down a free bet with a positive expected gain in terms of money, is explained by the concavity of the utility function with respect to the good in question. When a utility function is concave, the marginal utility derived from a good is decreasing: Any additional unit of the good adds less utility the more of the good the agent already has. If that is so, then setting the status quo at utility 0 , the expected utility of a gamble can be negative even if the expected money value of a gamble is positive. For the gamble in question here, this is so if the utility loss from losing \$100 is bigger than the utility gain

⁸ If we drop the assumption that SC will reject the single gamble at all possible wealth levels that he could be at during a sequence of gambles, then, as Stephen Ross, "Adding Risks: Samuelson's Fallacy of Large Numbers Revisited," *Journal of Financial and Quantitative Analysis*, 34 (1999): 323–339 shows, it is possible for an expected utility maximizer to reject the single gamble and accept the compound. However, I take it that SC's preferences are sensible even if the original assumption is met.

⁹ The case of SC illustrates a more general divergence between what seems reasonable and what expected utility maximization demands when it comes to large and small stakes gambles involving monetary pay-offs. Matthew Rabin and Richard Thaler, "Anomalies: Risk Aversion," *Journal of Economic Perspectives*, 15 (2001): 219–232 show that, for an expected utility maximizer, any significant risk aversion for small stakes gambles implies implausibly high levels of risk aversion for larger stakes gambles. Conversely, plausible levels of risk aversion for large stakes gambles imply near risk neutrality for small stakes gambles. For instance, an expected utility maximizer who turns down a 50/50 bet of losing \$10 and winning \$11 will turn down any 50/50\$ bet involving a loss of \$100, no matter how large the potential gain. And Matthew Rabin, "Risk Aversion and Expected Utility: A Calibration Theorem," *Econometrica* 68 (2000): 1281-1292 shows that under expected utility maximization, SC should also reject a 50/50 gamble of losing \$200 or winning \$20,000.

from winning \$200. For instance, suppose the agent's utility in money is given by $u(m) = \sqrt{m}$, and the agent's current wealth is \$100. Then the expected utility of accepting SC's gamble is $0.5 * \sqrt{0} + 0.5 * \sqrt{300} \approx 8.66$. Rejecting the gamble yields $\sqrt{100} = 10$, which is higher. And so an agent with such a utility function should reject SC's gamble, and will thereby display risk aversion. Figure 1 illustrates this.




Figure 1: A Concave Utility Function

If we assume expected utility maximization, the risk aversion SC displays for the low stakes gamble must mean that the utility function we assign to him is concave in a way that makes a difference for small amounts of money. And I show here that this in turn implies extreme levels of risk aversion for larger stakes gambles, under the assumption that the agent would turn down SC's small stakes gamble at any relevant wealth levels.¹⁰ Let current wealth be w . For SC to reject the small stakes gamble, the utility added by each dollar between w and $w + 200$ can be at most $1/2$ of the utility added by each dollar

¹⁰ The following illustration draws on Rabin and Thaler's "Anomalies: Risk Aversion".

between w and $w - 100$. But given that we assumed that SC would also turn down the gamble if his initial wealth level was $w + 200$, the added utility of each additional dollar between $w + 200$ and $w + 400$ can be at most $1/2$ of the added utility of dollar $w + 200$, which is itself at most $1/2$ of that of dollar $w - 100$. When we iterate this to larger sums of money, it turns out that dollar $w + 2,000$ can add at most 0.2% of the utility of dollar $w - 100$.

Hence, risk aversion for small amounts of money at wealth level w means that the utility gain from each additional dollar is only a fraction of what it is at w when SC is only \$2,000 richer. An extra dollar at $w + 5,000$ adds virtually no utility compared to the utility gains at $w - 100$. This is why a rejection of the gamble SC is offered implies extreme levels of risk aversion when it comes to larger stakes gambles, such as the compound gamble SC considers. Such an agent must be assigned a utility function with a rapidly decreasing marginal utility of money. If he rejects the small stakes gamble, SC should end up rejecting the large stakes gamble, because the larger sums of money he stands to gain from the larger stakes gamble add little utility for him. So if we want to insist that having preferences like SC's is rational, but hold on to the assumption that the wealth levels that can be achieved in the gambles are the only relevant outcomes, then we have to abandon expected utility theory. In particular, we have to argue that the way expected utility theory treats risk aversion is inadequate: Risk aversion cannot be merely a matter of a concave utility function. Alternatives to expected utility theory like REU theory hold out the promise of offering a treatment of risk aversion which can make better sense of the ordinary decision maker's attitudes to risk.

3. Risk-Weighted Expected Utility Theory

Most alternatives to expected utility theory have been introduced as merely descriptive theories of choice under uncertainty, with no claim to capturing rational choice.¹¹ Here, I focus on a recent alternative theory that was introduced explicitly as an alternative theory of instrumental rationality. And that is Buchak's REU theory. It is a rank-dependent theory formally equivalent to that introduced by John Quiggin in 1982.¹² But most of our conclusions will likewise apply for any alternative to expected utility theory that tries to make sense of preferences like those of SC.

In expected utility theory, only two components go into the agent's choice rule, or the representation of the agent's preferences: probabilities and utilities. REU theory adds a third element, namely a risk function $r(x)$. It represents agents as maximizing their risk-weighted expected utility. REU theory requires us to order the various outcomes an act could lead to according to their utility. Let these be o_1, o_2, \dots, o_n , where the index indicates the outcome's rank in terms of utility, starting with the lowest. The REU of the act is then calculated as follows: To the utility of the worst outcome, we add the weighted utility difference between that outcome and the second worst, and then the weighted utility difference between that outcome and the next better one, and so on. The weights are a function $r(x)$ of the probability of receiving at least the utility of the higher of the two outcomes.¹³ Formally, the REU of a gamble g is given by

$$\begin{aligned} REU(g) = & u(o_1) + r\left(\sum_{i=2}^n p(o_i)\right)(u(o_2) - u(o_1)) + r\left(\sum_{i=3}^n p(o_i)\right)(u(o_3) - u(o_2)) \\ & + \dots + r(p(o_n))(u(o_n) - u(o_{n-1})) \end{aligned}$$

¹¹ The most well-known is prospect theory, introduced by Daniel Kahnemann and Amos Tversky, "Prospect Theory: an Analysis of Decision under Risk," *Econometrica* 47 (1979): 263–291.

¹² John Quiggin, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization* 3 (1982): 323–343.

¹³ Buchak, *Risk and Rationality*, 53.

REU maximizers can be risk averse even if we do not assign decreasing marginal utility in money to them, e.g. if we assign them linear utility in money. Agents with linear utility in money turn out to be generally risk averse¹⁴ when $r(p(o))$ is smaller than $p(o)$ for every o , and $r(p(o))$ is convex. This means that, as the probability of receiving at least a particular outcome increases, the weight assigned to it becomes ever larger. An example of such a risk function Buchak uses throughout her book is that of $r(x) = x^2$. Given such a risk averse risk function, outcomes that are very good but unlikely will contribute less to the overall value of a gamble than they would for a less risk averse agent, or indeed for an expected utility maximizer.

How are the components of REU theory to be interpreted? According to Buchak, the utility function is meant to capture the agent's attitudes to outcomes, which pick out the goals she wants to achieve. As in expected utility theory, the probabilities capture her beliefs about how likely the various outcomes, or states that lead to those outcomes, are. And lastly, differences in the risk function represent differences in how agents structure the attainment of their goals, for instance, whether they are prudent or venturesome in their attainment of money. The REU of a gamble then gives us the instrumental value the agent assigns to a particular gamble. According to the theory, two agents with the same utility and probability functions can disagree about the instrumental value of a gamble if they have different risk functions. REU theory is thus a more permissive theory of instrumental rationality than expected utility theory: According to REU theory, there are various different permissible ways of structuring the attainment of the very same goals.

While, according to REU theory, instrumental rationality is permissive about what risk function an agent may have, the risk function we assign to an agent is still meant to pick out a more or less stable character trait that concerns how she acts in the face of risk. Buchak suggests that the risk function describes where

¹⁴ Buchak defines general risk aversion as follows: For any two gambles g and h , where h is a mean-preserving spread of g , the agent prefers g . Ibid., 21-22.

the agent falls between the two virtues of prudence and venturesomeness.¹⁵ It is plausible that such steady character traits could play an important role in explaining our choices in the face of uncertainty.

Moreover, one may wish to argue that there are non-instrumental requirements of rationality to act in accordance with one's character traits, and to be in that sense consistent. Regarding the counter-examples to expected utility theory, the hope is then to show not only how it could be instrumentally rational to have the preferences displayed in them. It is also to explain these preferences in terms of some steady disposition to choose in the face of risk, or a commitment to structuring the attainment of one's goals in a particular way. Buchak in fact applies the theory to various counter-examples to expected utility theory with that aim. Importantly for us, it also seems like REU theory can make sense of preferences like those of SC.

Suppose that SC is an REU maximizer with linear utility in money, such that $u(m) = m$, where m is amount of money, and risk function $r(x) = x^2$. Again, the probabilities presented to SC are taken at face value. Accepting the single gamble then has a REU of -25.¹⁶ Given that this is negative, SC will reject the gamble. But it turns out that the REU of accepting one hundred such gambles is 4004.56 given these utility and risk functions. Since SC's utility function is linear in money, this value can be interpreted as the certainty equivalent of the gamble: The agent is indifferent between receiving \$4004.56 for certain, and the larger compound gamble. In fact, the agent would accept the compound gamble resulting from any series of three or more gambles. Figure 2 shows the certainty equivalent, divided by the number of individual gambles, of compounds resulting from series of gambles of various lengths.

¹⁵ Ibid. 55-56.

¹⁶ Call SC's gamble s . We then have:

$$\text{REU}(s) = -100 + 0.5^2 * 300 = -25$$

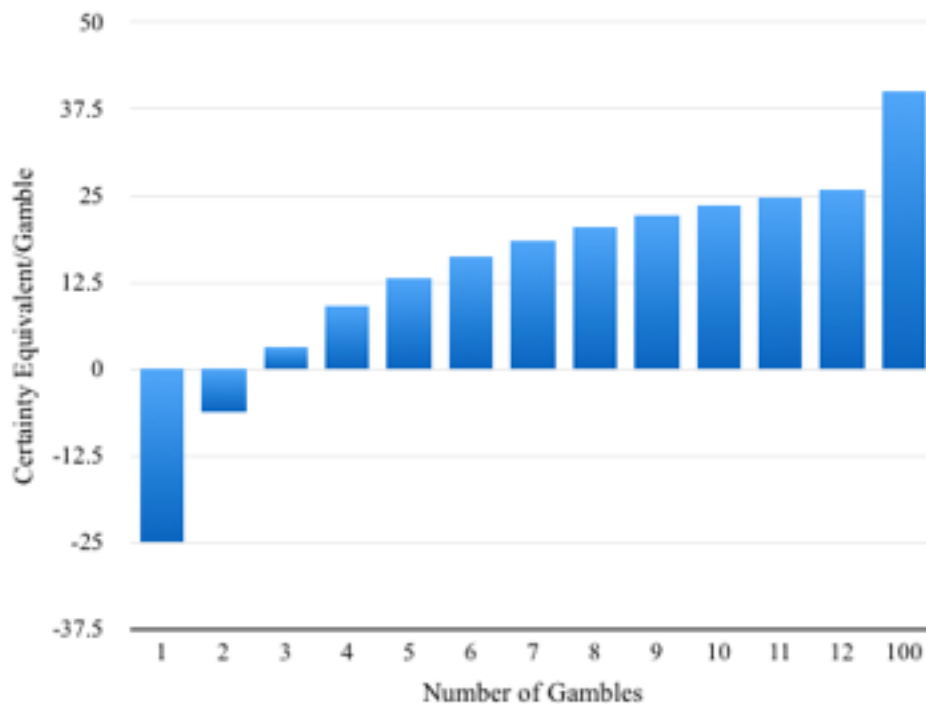


Figure 2: Certainty Equivalent Per Gamble for Repetitions of SC's Gamble

Note that the certainty equivalent per gamble increases with the number of repetitions that are offered. In order for SC's preferences to turn out rational, this must in fact be so. After all, the certainty equivalent of an individual gamble must be negative for SC to reject it. But the certainty equivalent of the compound gamble must be positive for him to accept it, and so the certainty equivalent per individual gamble must also be positive. In the case of REU theory, it can even be shown that for any gamble, under the conditions we specified, as the number of repetitions increases to infinity, the REU certainty equivalent per gamble approaches the expected utility of the gamble.¹⁷

It helps to stop and consider how this is possible. For the risk-weighted expected utility maximizer we considered, the individual gambles are independent in two senses: They are probabilistically independent, such that winning one of the gambles does not make winning the next one more likely. And they are

¹⁷ Ibid. 217-18.

independent in terms of utility, such that winning one does not affect the increase in utility he might gain through the next. This is due to the linearity of the agent's utility function in money. As his wealth decreases or increases, he nevertheless receives the same amount of added utility from winning or losing the next gamble. Under these circumstances, the expected utility of a combination of gambles is just the sum of the expected utilities of the individual gambles. And so for an expected utility maximizer for whom the different gambles are independent both in terms of utility and in terms of probability, each individual gamble would have the same value, whether it was part of a series of gambles or not.

This is not so for our REU maximizer. For him, there is a kind of interdependence or complementarity between the individual gambles that causes the instrumental value of the compound of the gambles to be different from the sum of the values of the individual gambles. Indeed, any theory which hopes to cast SC's preferences as rational must account for the interdependence of the instrumental values of the gambles he is offered. SC's gambles are probabilistically independent, and the failure of expected utility theory to account for SC's preferences suggests that interdependence in utility can't be the whole story: Decreasing marginal utility cannot explain SC's preferences, and it is not clear what other types of interdependence in utility could be involved here.

In the case of REU theory, it is the risk function which accounts for the complementarity in the instrumental value of the individual gambles. Due to the risk function, for an REU maximizer, the position of an outcome in the overall spread of outcomes, as well as the probability of other outcomes matters for how much one outcome contributes to the overall value of a gamble. And this is how compounding can lead to complementarities. It can affect the value contribution each of the outcomes of the individual gambles makes, by affecting the position of an outcome in the overall ranking of outcomes.

4. Risky Choice over Time

We have now seen that REU theory seems to be able to make sense of SC's preferences, while expected utility theory cannot. This seems to count in favour of REU theory. In the following, I want to show that the argument we just gave relies on ignoring an important observation, however. And that is that every individual risky choice we face in our lives is embedded in a long series of risky choices. This will also be so for any real-life case that resembles the choice SC faces. As we will see, given a common commitment, amongst critics of expected utility theory, to a particular way of choosing in dynamic contexts, and a plausible assumption about the specification of decision problems, this means that REU theory cannot in fact accommodate SC's choices.

We often face risky choices that resemble the individual gamble SC faced. We face similar monetary gambles when we decide how to invest our savings. But many non-monetary gambles also have a similar structure. Take my choice of whether to cycle to work or take the train. Cycling always takes the same amount of time, and I can be certain that I will make my first appointment on time. The train is less reliable, such that I may either end up late, or even have time for a coffee before my appointment. Or suppose that I have two options for lunch, one of which reliably produces bland food, whereas food at the other varies, with a 50% probability of either great or unpleasant food. Like in SC's case, it does not seem unreasonable to go for the safe option in these cases. But unlike SC's case, we rarely face a one-off choice involving the corresponding compound gamble representing, for instance, eating at the risky lunch option one hundred times. Seeing that one-off decisions involving such compound gambles are rare, we might think that SC's case is of little relevance for ordinary decision-makers. But I think this is too hasty. I have to decide how to get to work, and what to have for lunch every day. And that means that I will face the relevant compound gamble *over time*, even if I never face it in a one-off choice.

To get a sense of how repeated risk-taking could make a difference for agents with SC-type preferences, suppose that SC faces his gamble one hundred times in a row, on one hundred consecutive days. For

simplicity, assume he only finds out how much he ends up with in the end.¹⁸ How should he decide? Such sequential decision-making is standardly represented through a dynamic choice problem. In dynamic choice problems, the agent is modelled as making a series of time-indexed choices, where the final payoffs she receives depend on all the previous choices she made. Figure 3 shows a decision tree that represents the dynamic choice problem where SC makes only two consecutive choices. The round nodes stand for chance nodes, and we think of the agent as going ‘up’ or ‘down’ with the probability specified in the decision tree. The square nodes are decision nodes, where the agent can decide whether to go ‘up’ or ‘down’.




Figure 3: Dynamic Version of SC’s Decision Problem, Two Consecutive Choices

¹⁸ Whether he finds out in the end or throughout the decision problem doesn’t make a difference for SC if he follows the sophisticated strategy, described below, since he will reject each gamble either way. If he follows the resolute strategy, and finds out how gambles turned out along the way, SC also has to consider conditional plans whereby he accepts each gamble, unless he drops below a certain wealth level. Still, under normal circumstances, resolute SC will accept most gambles.

To solve these dynamic decision problems, it is not enough to provide a theory of how agents decide between probability distributions over outcomes, such as REU theory or expected utility theory. When the agent is at the first choice node, she is not deciding directly between such probability distributions over outcomes, but between two different continuation branches. To decide, she also needs a dynamic decision rule, that is, a rule about how to decide within a decision tree, given her preferences over probability distributions over outcomes. Two such rules are prominent in the literature.¹⁹ The first is resolute choice. Let a plan be a set of choices, one for each decision node the agent could find herself at in a decision tree. Each plan has a probability distribution over outcomes associated with it. A resolute agent considers which such probability distribution she prefers at the outset, picks a plan that leads to it, and then simply carries out that plan. We know that, at the outset of our decision problem, SC prefers taking the gamble one hundred times to never taking it. Plausibly, he also takes this to be his most preferred course of action out of all the ones available to him.²⁰ And so if SC is resolute, and we model his repeated choice as a dynamic choice problem, SC would choose at the beginning to accept the gamble every day, and then stick to that plan.

The second prominent dynamic choice rule is sophisticated choice. A sophisticated agent assumes that she will make rational decisions in the future, and that the set of choices it will be rational for her to make will not depend on the choices she has made in the past: She treats each choice as if it was the first in a series of choices. In terms of decision trees, she treats the continuation of a tree as if it was a new decision tree. The sophisticated agent then makes an optimal choice given her predictions of her own future choices.

¹⁹ For a formal treatment of dynamic decision problems and the dynamic choice rules introduced here, see Edward McClennen, *Rationality and Dynamic Choice: Foundational Explorations* (Cambridge: Cambridge University Press, 1990).

²⁰ At least this is so if he is the kind of REU maximizer we described above. Since for him the average REU per individual gamble is increasing, adding another gamble both has value in itself, and makes all the other gambles he faces more valuable. And so once we have crossed the threshold where the average value per gamble is positive, adding more gambles is always better.

Sophisticated choice is often associated with backward induction reasoning. We start at the end of the tree, and assume that the agent makes a rational decision between the probability distributions over outcomes associated with the final nodes. We then reason backwards from there. We eliminate those branches that the agent would not choose, and consider what the agent would decide at the previous choice nodes given that restriction, and so on. What does this choice strategy imply for agents with SC-type preferences? In the tree in Figure 3, at all the last choice nodes, SC would choose to reject the gamble and go ‘down’, both given the REU function we have been using, and Samuelson’s assumption that SC would reject the gamble at any wealth level he might be at during the series. But that means that at the first choice node, the agent will think of himself as just facing the one-off gamble. Again, he will reject, and choose ‘down’. If we were to scale the problem up to one hundred consecutive choices, we still get the result that the sophisticated agent rejects the gamble on each occasion. He thus, over time, ends up foregoing an almost certain large monetary gain.

Proponents of alternatives to expected utility theory have often defended resolute choice.²¹ We have just seen that sophisticated agents who violate expected utility theory may, over time, forego almost certain gains over time which they would have opted for in a one-off choice. For such agents, sophisticated choice may even result in a sure loss relative to some alternative series of choices they could have engaged in. For instance, sophisticated choosers will often be willing to pay at the beginning of a choice problem to not get to make a choice later on. If he is an REU maximizer with the risk and utility functions described earlier, a sophisticated SC would be willing to pay up to \$4004.56 to bind himself to take all gambles. But there is an alternative series of choices whereby he takes all gambles without paying this

²¹ See McClennen’s *Rationality and Dynamic Choice*, Mark Machina, “Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty,” *Journal of Economic Literature* 27 (1989): 1622–1668, and, to some extent, Buchak, *Risk and Rationality*.

cost. If we take this to show that sophisticated SC is instrumentally irrational,²² then we can preserve the rationality of his preferences only if we give up sophistication. Resolute choice is an alternative choice strategy that enables agents to avoid the sure loss sophisticated agents might be making. For instance, as we have seen, resolute SC can simply take all gambles without paying up to \$4004.56 to bind himself. This makes the strategy attractive to proponents of expected utility theory.

Various authors have been sceptical that resolute choice is rationally permissible, since it may require counter-preferential choice.²³ In the following, I will nevertheless proceed on the assumption that REU maximizers such as SC choose resolutely. First, this is because, for the reasons just given, resolute choice is defended by several proponents of alternatives to expected utility theory. Second, I have argued elsewhere that as long as we take only the utility function over outcomes to capture the agent's goals in action, as Buchak does, resolute choice is compatible with instrumental rationality even though it may call for counter-preferential choice.²⁴ Third, in real-life decision contexts involving repeated risk-taking in which agents tend to have SC-type preferences, agents often manage to choose in accordance with their more long-term preferences, and thus at least emulate resolute choice in the kinds of cases we are interested in. For instance, in investment decisions, we often make decisions on the basis of the long-term

²² For an argument that it is, see, for instance, R.A. Briggs, "Costs of Abandoning the Sure-Thing Principle," *Canadian Journal of Philosophy* 45 (2015): 827–840. For an argument that it needn't be, see Lara Buchak, "Revisiting Risk and Rationality: A Reply to Pettigrew and Briggs," *Canadian Journal of Philosophy* 45 (2015): 841–862.

²³ See, for instance, Isaac Levi, "Consequentialism and Sequential Choice," in Susan Hurley and Michael Bacharach (editors), *Foundations of Decision Theory: Issues and Advances* (Oxford: Blackwell, 1991), Patrick Maher, "Diachronic Rationality," *Philosophy of Science* 59 (1992): 120–141, and Katie Steele, "What are the Minimal Requirements of Rational Choice? Arguments from the Sequential Decision Setting," *Theory and Decision* 68 (2010): 463–487.

²⁴ Johanna Thoma, *Advice for the Steady: Decision Theory and the Requirements of Instrumental Rationality*, (Doctoral Dissertation, University of Toronto, 2017).

returns of an investment, even though we could potentially opt in and out of the investment on a shorter term basis. And it seems sensible to do so.

Fourth, note that in this section we considered only the case where an agent faces the choice of whether to accept a particular gamble exactly one hundred times. Even sophisticated choice may license acting in accordance with more long-term preferences in cases where there is no known end-point to a sequence of choices. If at any point, there is only a chance that the next choice will be the last in the sequence, backward induction reasoning is blocked. If the agent thinks that for any possible future decision, it is sufficiently likely that the gamble will be faced again several times, sophisticated choice, too, will recommend accepting the gamble and going with the more long-term preference. And this, I think, is the more realistic case. The argument I will make in the following will apply to sophisticated agents in this more realistic scenario as well. Lastly, insofar as the arguments presented in the following apply only to resolute agents, they establish that adopting resolute choice does not help alternatives to expected utility theory avoid counter-intuitive results after all. Defenders of sophisticated choice can thus see them as additional arguments against resolute choice.

5. Framing Decision Problems

In the last section, I described SC as facing one hundred individual gambles on one hundred consecutive days. We framed this as a dynamic choice problem involving one hundred consecutive choices. In this dynamic choice problem, as a resolute agent, SC will choose to take the gamble on every single day. But suppose SC had framed things such that he faces one hundred separate one-off choice problems, taking into account only current wealth and immediate monetary pay-offs, as we did above. In that case, SC will end up rejecting the gamble on each occasion. Given his resoluteness, framing thus makes a crucial difference for what SC should do. In the following, I want to argue that according to standard views on the framing of decision problems, resolute agents with SC-type preferences are thus not justified in framing their decision problems narrowly by looking at individual gambles in isolation.

There is an older debate within expected utility theory on whether and how agents are justified in making decisions applying expected utility theory to ‘small-world decision problems’. Small-world decision problems are decision problems that are not ‘grand-world’ — where grand-world decision problems are decision problems that individuate all elements of the decision problem so finely that everything that could potentially make a difference to an agent’s evaluation of potential options is included.²⁵ Small-world decision problems are coarser in that they leave out some such detail. Arguably, only small-world problems can be feasibly considered by human decision-makers. Still, it is a common commitment amongst decision theorists that rational agents should try to at least approximate whatever decision they would make in the grand-world context. James M. Joyce expresses this explicitly when he claims that a decision’s rationality depends both on the extent to which the decision obeys rationality axioms within the small-world decision problem the agent is considering, and on the extent to which she is justified in believing that she would come to the same decision, were she to consider her grand-world decision problem, and capable of evaluating it consistently, taking into account all her relevant attitudes.²⁶ This commitment is natural, given the grand-world decision problem is the ‘fully considered’ decision problem, in that it allows for all the agent’s relevant beliefs and desires to find expression. As Joyce shows, Leonard Savage’s version of expected utility theory offers no guarantee that the utility representation of an expected utility maximizer’s attitudes in a small-world decision problem will agree with the utility representation of her attitudes in the corresponding grand-world problem. Richard Jeffrey’s, on the other hand, does.²⁷

This older debate about the justifiability of small-world decision-making concerned only non-dynamic decision problems. These may consider the agent’s own future choices insofar as they affect the potential outcomes of the agent’s present choices and the probability with which they occur, but don’t model them

²⁵ See Leonard Savage, *The Foundations of Statistics*, 2nd revised edition (Hoboken: Wiley, 1972).

²⁶ James M. Joyce, *The Foundations of Causal Decision Theory* (Cambridge: Cambridge University Press, 1999), 74.

²⁷ Richard Jeffrey, *The Logic of Decision*, 2nd edition (Chicago: University of Chicago Press, 1965/1983).

explicitly as choice nodes in a dynamic choice problem. Once we turn to dynamic choice theory, decision problems can be more or less detailed not only in the specification of outcomes (and other standard elements of non-dynamic choice problems). They can also be more or less detailed in which and how many of the agent's decisions are modelled explicitly in the dynamic choice problem. Including more or fewer decisions in a dynamic choice problem introduces another way in which framing may affect what an agent concludes she should do on any one occasion: It may affect an agent's choices by affecting her planning horizon. For a resolute agent, who makes a plan for the entire dynamic decision problem she takes herself to be facing and follows through with it, the length of the dynamic choice problem, and thus her planning horizon, has a very direct impact on her choices.

We can now extend the notion of a grand-world decision problem to dynamic choice problems. We can think of an agent's grand-world dynamic choice problem as the dynamic choice problem that, aside from finely individuating all other elements of the decision problem, explicitly models all decisions an agent may get to make in her life. According to the commitment expressed by Joyce, agents should try to approximate the decision they would make in this grand-world decision problem. Decisions made on the basis of small-world dynamic decision problems, such as those modelling only a subset of an agent's life choices explicitly, are rational only to the extent that the agent is justified in believing that she would come to the same conclusion if she were to consider her more comprehensive grand-world dynamic decision problem.

Looking back at resolute SC now, were he to frame the choice about each individual gamble separately, and come to the conclusion that he should not take the gamble, he would clearly not be justified in believing that this is the conclusion he would come to in the grand-world decision problem. Any resolute agent with SC's preferences will evaluate an individual gamble differently depending on whether she conceives of herself as choosing it in a one-off choice or as part of a series of one hundred similar choices. And presumably, for such an agent, embedding a gamble in an even larger sequence of gambles

will further change her evaluation of the single gamble.²⁸ If SC is resolute, and aims to choose as he would in the grand-world decision problem where he faces his gamble repeatedly, he should not consider himself as facing only a one-off choice. The future choices he faces are actually crucially relevant to any current choice.

6. Risk-weighted Expected Utility Theory and Framing

SC has preferences which make future repetitions of his gamble relevant to a present choice he faces, given that he is resolute. As we have seen in Section 3, REU theory can apparently account for his preferences because it involves a risk function which introduces complementarities between different risky choices the agent faces, even when these are independent in terms of probability and utility. But, as I want to argue here, this feature of REU theory means that, for resolute REU maximizers, it will be true quite generally that they are not justified in thinking that a choice made in the context of a small-world decision problem will concur with what they would choose in the grand-world dynamic choice problem. And in fact, in the grand-world dynamic choice problem, they will choose approximately like expected utility maximizers.

²⁸ This is assuming that moving to a grand-world context doesn't alter the agent's preferences in such a way that the divergence between the agent's attitudes to small stakes and large stakes gambles disappears, that is, that SC-style preferences are not based on some kind of mistake that an agent discovers once she considers her grand-world decision problem. However, first, SC's preferences do not obviously involve a mistake. This is why they form a motivating case for REU theory in the first place. And second, if we had independent reason for thinking SC's preferences are mistaken, then we would have an even more direct reason for rejecting REU theory and other theories that could accommodate SC-type preferences than the one provided in the following.

Apart from the special case where $r(x) = x$ and REU theory reduces to expected utility theory, the risk function creates complementarities between any two risky choices an agent faces in her life. In REU theory, just as in expected utility theory, utilities are assigned to outcomes involving different kinds of goods, occurring at different times and places. Having such a single measure allows us to express the way in which the agent trades off different kinds of goods, or the way in which she evaluates gambles that involve different kinds of goods. The risk function is applied to this single utility measure. And thus it creates complementarities between any two gambles the agent faces.²⁹

We can illustrate this with the transport and restaurant choice scenarios introduced above. Like any risky choices, according to REU theory, these can also be described as choices over utility gambles. Suppose that the probabilities involved in these gambles are also independent for any two occasions where I face the choice, and that the utilities in the transport and restaurant choice are the same as those in SC's original choice. Now if I am a resolute REU maximizer with the risk function we used above, it follows that if I think of my choices of means of transport two at a time, I will come to a different conclusion from when I think of them three at a time. In the first case I will cycle both times; in the second case I will

²⁹ In their "Costs of Abandoning the Sure-Thing Principle," Briggs similarly notes that in REU theory sub-gambles do not have "stable values", that is, their value depends on what larger gamble they are a part of. This explains why REU maximizers may find themselves in problematic dynamic choice problems where they stand to make a sure loss unless they give up sophistication. Briggs also shows that the lack of a stable value of sub-gambles means that even in non-dynamic choice problems, the utilities that describe an agent's attitudes to non-dynamic grand-world problems do not map on straightforwardly to the utilities that would accurately represent her attitudes in corresponding small-world problems, as they do, for instance, in Jeffrey's version of expected utility theory. In her "Reply to Pettigrew and Briggs," Buchak acknowledges this, though slightly qualifies the claim. Small-world reasoning is thus already problematic for REU maximizers. What I show here is that the difference that planning horizon makes for resolute agents introduces an additional problem for small-world decision making in dynamic contexts.

take the train each time. Or suppose I decide to consider my lunch decision together with my decision of what means of transport to take on two consecutive days. This will make me change my choice from cycling to taking the train, along with eating at the riskier restaurant — even though the utilities of the possible outcomes of all these decisions are independent from one another.

The fact that REU theory creates complementarities between all the risky choices an agent faces, together with the assumption that the agent is resolute thus leads to a serious problem for REU theory. It ensures that an agent can never be justified in believing that the solution to a decision problem that takes into account only a single, or a small number of risky choices is going to approximate the solution to the grand-world dynamic decision problem. For the REU maximizer, different dynamic decision frames may result in different recommendations regarding one and the same action, even if that action will influence neither the probability nor the utility of any other outcomes she may end up with in the future. More short-term perspectives on one's choices will not generally agree with the most long-term perspective.³⁰ According to the criterion for framing decision problems introduced in the last section, for resolute REU maximizers, decisions based on dynamic choice problems taking into account only a sub-set of the agent's choices will thus not be rational.

Worse still, there are reasons to believe that in the grand-world dynamic decision problem, a risk averse and resolute REU maximizer will choose approximately like an expected utility maximizer, at least when it comes to small-stakes gambles. We have seen above that for a risk averse REU maximizer, as the

³⁰ Note that expected utility theory does not have this particular problem. For the expected utility maximizer, when she faces a series of gambles which are independent in terms of utility and probability, how she carves up the decision problem is not going to matter. In that case, the expected utility of the compound gamble is the same as the sum of the expected utilities. When she constructs a small world decision problem in which she faces a one-off choice of whether to accept an individual gamble, this will give her the same conclusion as the one she would draw were she to consider the series of gambles as a whole.

number of repetitions of a gamble goes to infinity, the average REU of each gamble tends to its expected utility. And so a resolute agent considering a large number of repetitions of the same gamble will choose just like an expected utility maximizer. Now of course the actual grand-world problem an agent faces will be more complicated than a large number of repetitions of the same gamble. But we can speculate that in the case of large numbers of different small-stakes gambles, too, resolute REU maximizers will behave approximately like expected utility maximizers. For one, Buchak's repetition theorem will also apply, for instance, to the compound gamble resulting from the many risky choices one makes in a typical day, or in a typical month, so that the average REU of a typical day's or month's choices tends towards its expected utility.³¹

If, when facing small-stakes gambles, risk averse and resolute REU maximizers behave approximately like expected utility maximizers in the grand-world dynamic choice problem, and rationality demands that agents aim to make the choices they would make in the grand-world dynamic choice problem, then rationality demands that resolute REU maximizers behave approximately like expected utility maximizers regarding such gambles. Interestingly, this means that such REU maximizers may be rationally permitted to consider small-world decision problems regarding small-stakes gambles after all. Only, they should evaluate them as if they were expected utility maximizers.

This conclusion undermines the central motivation for REU theory, namely that it can account for various counter-examples to expected utility theory. Assuming the agent is resolute, and aims to choose as she would in the grand-world dynamic decision problem, REU theory cannot in fact cast choices like those of SC as rational. When SC says that he would reject the single gamble, he apparently fails to integrate his decision concerning that individual gamble into the grand-world decision problem — or otherwise he

³¹ Another supporting consideration is that, as long as different gambles are reasonably independent, an agent's exposure to risk throughout her life is diversified, and the variance of the overall risk 'portfolio' is lower than the weighted average of the individual gambles' variances. The lifetime gamble is in that sense less risky than the constituent gambles. The risk function makes the agent sensitive to that lower variance.

would accept it. He was probably already playing less favourable gambles every day by holding some of his retirement savings in stocks.³²

In fact, all of the toy examples that are used to motivate and illustrate REU theory, and not only SC's case, involve agents who must have left out relevant detail from the specification of their decision problems, if they are resolute. Given that these examples are usually presented as one-off choices involving outcomes that are described merely as the immediate goods one receives as the consequence of having made that choice, we can no longer be sure that on a specification of the decision problem that includes everything that is relevant to a resolute agent, REU theory can make sense of them. At the very least, then, the burden of proof is on proponents of REU theory to show that the counter-examples can still be made sense of once we think of them in the context of the agent's grand-world dynamic choice problem. But it also follows from the preceding that if these problems involve small stakes gambles and we think that the agent faces many such gambles in her life, then a resolute agent should behave roughly like an expected utility maximizer. And in that case there is no hope of REU theory accommodating the counter-example to expected utility theory.

The conclusion that risk averse, resolute REU maximizers should choose approximately like expected utility maximizers when facing small-stakes gambles also calls into question whether, quite apart from its treatment of these motivating examples, REU theory offers a true alternative to expected utility theory. Indeed, out of the two, expected utility theory is the simpler theory, making it more attractive. There is

³² Shlomo Benartzi and Richard Thaler, "Myopic Loss Aversion and the Equity Premium Puzzle," *Management Science* 45 (1999): 364–381. Benartzi and Thaler refer to risk aversion for small stakes gambles as "myopic loss aversion", suggesting that what explains these preferences is loss aversion in the sense introduced by Daniel Kahnemann and Amos Tversky, "Choices, Values, and Frames," *American Psychologist*, 39 (1984): 341–350, combined with "narrow framing", as discussed in Daniel Kahnemann and Dan Lovallo, "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk-Taking," *Management Science* 39 (1993).

also the related question of what the REU representation would even capture, if rational REU maximizers in fact behave like expected utility maximizers much of the time. It can't capture the agent's observable choices, since they are those of an expected utility maximizer. Instead, the REU representation must be based on some underlying preferences that the agent would display only in hypothetical one-off choices. But it is not clear why we should care about representing such preferences. Most importantly, the REU representation certainly no longer captures what REU theory set out to capture. The distinguishing feature of REU theory was supposed to be that it introduces a risk function that represents the agent's commitment to treating risks in a certain way, or her venturesome or prudent character traits. A resolute REU maximizer who acts like an expected utility maximizer in the grand-world decision problem displays no such character traits or commitments, since the preferences that may be representable as REU maximizing never manifest themselves in choice.

7. Permissiveness About Dynamic Choice

We have seen that REU maximizers will make choices regarding small-stakes gambles roughly like expected utility maximizers given two conditions: First, the agent is resolute, and chooses in dynamic decision problems as she would, were she to make one choice outright. Second, the agent aims to choose as she would in her grand-world decision problem, and that grand-world decision problem includes a long series of risky choices. The same will in fact hold for any alternative to expected utility theory that tries to make sense of SC's preferences. This is due to the fact that for anyone with SC's preferences, there exists a divergence between the evaluation of a gamble in isolation and the evaluation of a series of gambles. Seeing that SC-type violations of expected utility theory are common, any alternative to expected utility theory that tries to make better sense of our ordinary attitudes to choice under uncertainty should try to account for them. And so the problem I have pointed to poses a general challenge to alternatives to expected utility theory.

Alternatives to expected utility theory like REU theory can avoid the conclusion we reached in the last section by either relaxing the assumption that agents should be resolute, or by relaxing the assumption that agents ought to aim to act as they would in their grand-world decision problem, or indeed by relaxing both conditions. If we wanted to relax the latter assumption and be more permissive about the framing of decision problems, perhaps we could say that rationality is silent on the question of what temporal perspective the agent ought to take. We may even hold that it is perspective-dependent in the sense that we can only ever declare an action rational with reference to whatever temporal perspective the agent took, while the choice of perspective itself is arational.

For the sake of argument, suppose the resolute REU maximizer is free to take whatever temporal perspective she likes. The problem now is that almost any course of action could be endorsed by REU theory with reference to the right perspective. In our example, somebody with an REU function that can account for SC's choices in a one-off decision may accept all gambles or reject them all in a series of choices. The agent may even group different occasions in dynamic decision problems of different length, such that she accepts some but not others. Similar claims apply for any more complex series of gambles an REU maximizer may face.

There are two reasons why this is a worrying implication of permissiveness about temporal perspective. First, it means that REU theory gives only very limited action-guidance. Looking to the REU representation of one's preferences is of little help if that representation is compatible with a wide variety of choice behaviours, and rationality gives us no further guidance. Second, remember that the risk function was supposed to capture an agent's steady choice disposition to take risks, or her 'venturesomeness'. What we have found now, however, is that the theory allows for a wide range of different choice behaviours even given some fixed risk function. An agent with an extremely risk averse risk function may still accept an SC-type gamble, namely if she is resolute and adopts a grand-world decision frame. And an agent who has a much more 'venturesome' risk function may still reject the gamble, by thinking of each choice separately. If this is so, it seems like the risk function no longer in fact captures any kind of steady disposition, or character trait. At most, it captures a disposition of prudence or

venturesomeness relative to some temporal perspective. But unless the agent also has a steady disposition to choose particular temporal perspectives, this will not result in any steady choice behaviour in the face of uncertainty independently of temporal perspective.

Relaxing the assumption of resolution is another way in which a proponent of REU theory may wish to avoid the conclusions of the last section. And doing so may be independently plausible. Resolution is advocated by many critics of expected utility theory as if it was the only way to avoid the pragmatic disadvantages associated with sophistication. But in fact those who violate expected utility theory need not be strictly resolute in order to avoid sure loss. In the case of SC, being resolute is one way of avoiding the sure loss of paying (potentially a lot of) money to bind himself to accepting SC-type gambles in the future. But he also avoids sure loss by not binding himself and rejecting the gambles on some or all occasions. As long as there is some chance that the agent would do worse by accepting more gambles, then this alternative response also avoids sure loss. A more permissive attitude to dynamic choice strategies may thus also be endorsed by those who abandoned sophistication to avoid the possibility of sure loss.

What would such permissiveness mean for REU theory? As with permissiveness about temporal perspective, it would mean that the same combination of risk, utility, and probability function could manifest in a variety of different choice behaviours. It could lead to a wide range of different choice behaviours even given a fixed temporal perspective. For example, an REU maximizer with SC's preferences may accept or reject all gambles she is offered over time, or indeed choose to accept only some, even if she adopts a grand-world dynamic decision frame.

What we have found, then, is that relaxing the assumption that REU maximizers are resolute, and/or relaxing the condition that they need to consider themselves as trying to solve their grand-world decision problem makes the predictions of REU theory extremely sensitive to the agents' choice of temporal frame, and/or the way they choose to behave in dynamic decision problems. This sensitivity calls into question REU theory's claim to representing an agent's steady dispositions in the face of risk via the risk

function. This is problematic for the explanatory project of REU theory, since these dispositions were meant to explain rational agents' choices under risk. But it also limits the extent to which REU theory is action-guiding. Moreover, the sensitivity just described does not make the project of coming up with a REU representation in the first place much easier. As in the case of the resolute REU maximizer who aims to choose as she would in the grand-world dynamic choice problem, when we are permissive about dynamic choice and temporal framing, the agent's choice behaviour may still end up being a poor guide to the underlying REU maximizing preferences the theory is presupposing — those that would supposedly manifest themselves in hypothetical one-off choices.

8. Conclusions

The counter-examples to expected utility theory that motivate alternative theories are usually presented as one-off decisions, isolated from the agent's wider decision context. However, any real agent faces many decisions involving uncertainty in her life. I argued that this fact raises deep problems for alternatives to expected utility theory. If agents are resolute, and aim to choose as they would in the grand-world dynamic choice problem, REU theory, as well as any theory that tries to accommodate our ordinary, divergent attitudes to small and large stakes risks, struggles to still accommodate the examples that motivate them in the first place once the dynamic context is made explicit. In fact, REU theory ends up making approximately the same recommendations regarding small-stakes gambles as expected utility theory. It thus offers no real alternative to expected utility theory.

To avoid this conclusion, we need to be more permissive about dynamic choice or the framing of decision problems, or both. But this permissiveness also makes the predictions of REU theory very sensitive to an agent's choice of temporal frame, and her choice about how to treat dynamic decision problems. This, we have argued, undermines the theory's claim to capturing a steady choice disposition in the face of risk, and thus its purpose as a normative and explanatory theory of choice under uncertainty. One alternative,

of course, is to go back to expected utility theory, even if that may entail that we cannot rationalize much of our ordinary risk averse behaviour.