EXPECTED COMPARATIVE UTILITY THEORY: A NEW THEORY OF RATIONAL CHOICE*

DAVID ROBERT

I. INTRODUCTION

Standard rational choice theory, otherwise known as *Expected Utility* (*EU*) *Theory*, counsels agents to rank their choice options (from least to most choice-worthy) according to their *EU*.[1] The EU of an option is a probability-weighted average of each of its possible utilities. EU Theory has been the dominant model of rational choice since the late 17th century,[2] and in more recent years (1920s–), has received foundational support from both economists and philosophers.[3] In

*Thanks to the editor, Douglas Lackey, and several anonymous reviewers for very helpful comments. I dedicate this paper to my parents, Guy Robert and Josephine Moffa, whom I love dearly and whose loving support made writing this paper possible.

[1] For helpful introductions to EU Theory, see Rachael Briggs, "Normative Theories of Rational Choice: Expected Utility," *The Stanford Encyclopedia of Philosophy* (Spring 2017 edition), ed. Edward N. Zalta, https://plato.stanford.edu/archives/spr2017/entries/rationality-normative-utility/; Katie Steele and H. Orri Stefánsson, "Decision Theory," *The Stanford Encyclopedia of Philosophy* (Winter 2016 edition), ed. Edward N. Zalta, https://plato.stanford.edu/archives/win2016/entries/decision-theory/.

[2] Antoine Arnauld and Pierre Nicole, *The Art of Thinking; Port-Royal Logic*. Translated, with an introduction by James Dickoff and Patricia James, and a foreword by Charles W. Hendel (Indianapolis: Bobbs-Merrill, 1964 [1662]).

[3] Frank P. Ramsey, "Truth and Probability," *The Foundations of Mathematics and other Logical Essays*, ed. Richard B. Braithwaite (London: Routledge and Kegan Paul, 1931) 156–84; John Von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton UP, 1944); Leonard Savage, *The Foundations of Statistics* (Wiley, 1954); Ethan Bolker, "Functions Resembling Quotients of Measures," *Transactions of the American Mathematical Society* 2 (1966): 292–312; Richard Jeffrey, *The Logic of Decision*, 2nd ed. (Chicago: U of Chicago P, 1983).

this paper, I will argue for a novel alternative to EU Theory, namely *Expected Comparative Utility* (*ECU*) *Theory*. ECU Theory aims to improve upon rational choice models based on *regret minimization*.[4] These models are founded on the idea that rational decision-makers strive to minimize choices that elicit regret—the disappointment that occurs when they fail to choose the available option they most prefer (i.e., the available option which carries the greatest utility). Regret minimization models grew out of the seminal work of mathematicians Abraham Wald and John Milnor,[5] and psychologists Daniel Kahneman and Amos Tversky.[6] In the late 1970s, Kahneman and Tversky demonstrated deviations between EU-maximizing choices and people's actual choices, and proposed an alternative (descriptive) decision theory, *Prospect Theory*, which introduced for the first time the notion of a *reference point* (a benchmark of sorts)—that is, an outcome that partitions the set of decision outcomes into perceived gains and losses.

Like Prospect Theory (but unlike EU Theory), the basic regret minimization choice model makes use of a *benchmark* (or zero point of choiceworthiness). According to this model, for any choice option, *a*, and for any state of the world, *G*, the extent of *a*'s "*regret*," in *G*, is the extent to which *a*, in *G*, falls short of whichever available option carries the greatest utility in *G*, and the degree of choiceworthiness of *a*, in *G*, is the degree to which the choice of *a* minimizes that "regret." Thus, the degree of choiceworthiness of *a*, in *G*, is the difference in utility, in *G*, between *a* and *whichever available option carries the greatest utility* in *G* (i.e., the benchmark). Understood in this way, maximizing expected choiceworthiness always coincides with maximizing EU.

ECU Theory also makes use of a benchmark, one that is importantly different from that employed in the regret minimization choice model: for any choice option, *a*, and for any state of the world, *G*, the degree of choiceworthiness of *a*, in *G*, is the difference in utility, in *G*, between *a* and *whichever alternative to a carries the greatest utility* in *G* (i.e., the benchmark). This difference in utility is what I will call the *comparative utility* (*CU*) of *a*. Roughly speaking, ECU

---

[4] See in particular Ronald R. Yager, "Generalized Regret Based Decision Making," *Engineering Applications of Artificial Intelligence* 65 (2017): 400–05. For early models taking into account anticipated regret, see David E. Bell, "Regret in Decision Making Under Uncertainty," *Operations Research* 30 (1982): 961–81; Peter C. Fishburn, "Nontransitive Measurable Utility," *Journal of Mathematical Psychology* 26 (1982): 31–67; Graham Loomes and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal* 92 (1982): 805–24.

[5] Abraham Wald, *Statistical Decision Functions* (New York: Wiley, 1950); John Willard Milnor, "Games Against Nature," *Decision Processes*, ed. R. M. Thrall, C. H. Coombs and R. L. Davis (New York: Wiley, 1954).

[6] Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47 (1979): 263–92; idem, "The Psychology of Preferences," *Scientific American* 246 (1982): 160–73.

Theory counsels agents to rank their options (in terms of how choiceworthy they are) according to their ECU. For any choice option, *a*, the ECU of *a* is the probability-weighted average of the comparative utilities of *a* across the various states of the world. In this paper, I will show that in a number of ordinary decision cases, ECU Theory gives different verdicts from those of EU Theory.

From start to finish, this paper is predicated upon two basic assumptions:

First assumption: For any agent, *S*, and for any choice option, *a*, for *S*, *a*'s *utility* is a cardinal indicator of (*S*'s) preference and is derived from *S*'s preferences as in standard decision theory, that is, via a representation theorem. This requires that *S*'s preferences obey a series of axioms of *rational preference*, one of which (i.e., the independence axiom) entails the *Independence of irrelevant alternatives* (*IIA*) (for preferences): if an option, *a*, is preferred over some alternative option, *b*, then introducing a third option, *c*, in the choice situation will not change the preference ordering between *a* and *b*. For the present purposes, *rational preference* is defined, by stipulation, as satisfying the IIA.

Second assumption: Let a *decision under certainty* be a choice situation where an agent is subjectively certain about which state the world is in (and assigns probability 1 to that state of the world being actual), and let a *decision under uncertainty* be a choice situation where the agent is *not* subjectively certain about which state the world is in. (States of the world, or states, are defined here as possible, mutually exclusive states of affairs.) For any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, *a* is *choiceworthy* for *S* (or [i.e.] it is *rational* for *S* to choose *a*) if and only if *a* maximizes utility, over the space of all alternatives, in the state of the world to which *S* assigns probability 1. I will refer to this as the *Utility Maximization* (*UM*) *Principle*. The UM Principle defines a binary measure of choiceworthiness for decisions under certainty (i.e., whether an option is choiceworthy tout court or *un*choiceworthy tout court).

Now, in addition to a binary measure, we also require a graded measure of choiceworthiness. For any number of alternative choice options, *a*, *b*, *c*, *d*, and *e*, we want to say that *a* (utility: 100) is more choiceworthy than *b* (utility: 5) even if *a* is not choiceworthy tout court (i.e., *a* does not maximize utility). We also want to say that the *extent* to which *a* is more choiceworthy than *b* is greater than the extent to which *c* (utility: 10) is more choiceworthy than *b*. In order to say that *a* is more choiceworthy than *b* (and to what extent), we cannot rely on a binary measure of choiceworthiness. Whether (and to what extent) *a* is more choiceworthy than *b*, and by implication, whether (and to what extent) any option is more choiceworthy than any other within a set of alternatives is necessarily a function of *how* choiceworthy each of the two options is within the set of alternatives (and not necessarily a function of one being choiceworthy tout court and the other *un*choiceworthy tout court). To ask how choiceworthy an option is to ask how desirable or worthy of being chosen that option is, (more or less[7]) how imperative it is to choose that option. Such a question is well-formed and meaningful. In order to answer the question, we require a graded (quantitative) measure of how choiceworthy options are. In this paper, I will develop such a measure.[8] In particular, I will show that utility and EU, as well as measures of regret minimization are inadequate as measures of choiceworthiness. I will argue that taking seriously the notion of a measure of choiceworthiness points to the need for a new theory of rational choice, ECU Theory.

Section II of this paper will develop a graded choiceworthiness measure for decisions under certainty (i.e., CU). On the basis of this measure, section III will explicate and defend ECU Theory. Finally, section IV will assess the cogency of EU Theory in light of ECU Theory.

II. THE CU PRINCIPLE

Upon first consideration, one might suppose that for any agent, *S*, faced with any decision under certainty and for any number of alternative options, *a*, *b*, *c*, *d*, and *e*, available to *S*, the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which *S* (rationally) prefers *a* to *b*, or perhaps the extent to which *S* (rationally) prefers *a* to *b* *more* than *S* (rationally) prefers *b* to *a*. However, intuitively, that is a mistake. Even though we are comparing *a* to *b*, we want to see how *a* and *b* measure up to the *very best alternative options on offer*, in the following way: the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which *S* (rationally) prefers *a* to the most preferred alternative to *a* (either *b*, *c*, *d*, or *e*) *more* than *S* (rationally) prefers *b* to the most preferred alternative to *b* (either *a*, *c*, *d*, or *e*), or so I will argue in what follows.

The present section is devoted to explicating a graded, quantitative choiceworthiness measure for decisions under certainty. Let us begin with a short argument for such a measure. For any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, *a* is choiceworthy for *S*, by definition, if and only if *a* is

[7] Let *a* and *b* denote two mutually exclusive and jointly exhaustive choice options. *a* is more choiceworthy than *b* if and only if choosing *a* is more imperative than choosing *b*, and *a* is just as choiceworthy as *b* if and only if choosing *a* is just as imperative as choosing *b*. However, if *a* is just as choiceworthy as *b*, then both *a* and *b* are choiceworthy, whereas if choosing *a* is just as imperative as choosing *b*, then neither choosing *a* nor choosing *b* is imperative.

[8] "Choiceworthiness", as expressing the degree to which a decision option meets an agent's ends, should be distinguished from the moral notion of "rightness", which arguably cannot be measured on a graded scale (see Thomas Hurka, "More seriously wrong," Manuscript, https://www.brown.edu/academics/philosophy/ethical-inquiry/sites/brown.edu.academics.philosophy.ethical-inquiry/files/uploads/seriously%20wrong%202016.pdf; Neil Sinhababu, "Scalar Consequentialism the Right Way," *Philosophical Studies* Forthcoming; Kristian Olsen, "Subjective Rightness and Minimizing Expected Objective Wrongness," *Pacific Philosophical Quarterly* Forthcoming).

worthy of being chosen by *S* over whichever alternative to *a* carries the greatest utility. It follows that the *extent* to which *a* is choiceworthy for *S* is the extent to which *a* is worthy of being chosen by *S* over whichever alternative to *a* carries the greatest utility. But the extent to which *a* is worthy of being chosen by *S* over whichever alternative to *a* carries the greatest utility is the difference in utility between *a* and whichever alternative to *a* carries the greatest utility. Therefore, the extent to which *a* is choiceworthy for *S* (or [i.e.] the measure of *how* choiceworthy *a* is for *S*) is the difference in utility between *a* and whichever alternative to *a* carries the greatest utility (henceforth, the *CU Principle*).

Let us now consider a longer argument. The simplest attempt at defining a quantitative choiceworthiness measure for decisions under certainty is as follows: for any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, a measure of the choiceworthiness of *a* for *S* is the utility of *a* (or [i.e.] the measure of the utility of *a*) (in the state of the world to which *S* assigns probability 1). I will refer to this as the *Utility Principle*. The UM Principle is true if (*but not* only if) the Utility Principle is true. If the Utility Principle were true, then EU Theory would be vindicated, since measures of choiceworthiness would be interchangeable with measures of (rational) preference (for further elaboration, see section IV). The Utility Principle is, however, untenable.

First, *measures* of quantities, for example, 20°C for temperature, are meaningful (and *only* meaningful) relative to a given zero point and unit of measurement. (Let us call this the *Measurement Principle*.) In the case of temperature, the measure (e.g., 20°C) is defined in relation to the zero point and unit of measurement (i.e., the measure itself presupposes a given temperature unit and zero point of temperature). That is not the case for utility. In accordance with the Measurement Principle, the measure of *a*'s utility (e.g., 20 units of utility [or *utiles*]) is meaningful (and *only* meaningful) relative to a given utility unit and zero point of utility. However, the measure (e.g., 20 units of utility) is *not* defined in relation to the unit and zero point (i.e., the utility measure itself does *not* presuppose a given utility unit and zero point of utility).[9] These values must be explicitly specified. Hence, the Utility Principle is at best underspecified.

Second, even *relative to an explicitly given utility unit and zero point of utility*, the measure of the choiceworthiness of *a* for *S* is not necessarily its utility. In accordance with the Measurement Principle, for any given decision situation (under certainty) and for any specified utility unit and zero point of utility (for that situation), the measure of the choiceworthiness of any available option is its

[9] "The zero point and the unit in an expected utility representation are arbitrary; utility values become meaningful only once they have been fixed." (Mark Colyvan and Alan Hájek, "Making Ado Without Expectations," *Mind* 125 (2016): 834).

utility value if and only if it is possible to ascertain how choiceworthy any available option is (in that situation) by solely considering its utility value in relation to that specified utility unit and zero point of utility. In practical terms, what this means is that, for any given *decision setup* (i.e., any decision situation combined with any explicit specification of a utility unit and zero point of utility), the measure of the choiceworthiness of any available option is its utility value if and only if (1) any available option is choiceworthy just in case its utility value is equal to or greater than zero (and not choiceworthy otherwise) and (2) the degree of choiceworthiness of any available option is its utility value. Now, it is straightforward to come up with decision situations where it is possible to select a specific zero point of utility (and a specific utility unit) such that it is *not* the case that any available option is choiceworthy (in that situation) if and only if its utility value is equal to or greater than zero. Per the UM Principle, there are possible decision setups where an option has a positive utility value and is nevertheless *un*choiceworthy, namely setups where that option does *not* maximize utility over the space of all available alternatives, and there are possible decision setups where an option has a negative utility value and is nevertheless choiceworthy, namely setups where that option *does* maximize utility over the space of all available alternatives. Therefore, per the Measurement Principle, there are possible decision setups such that it is *not* the case that the measure of the choiceworthiness of any available option (in that setup) is its utility value.

In light of the preceding considerations and in accord with the Measurement Principle, it is necessarily the case that for any agent, *S*, faced with <u>any</u> decision situation under certainty and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* depends on a unit of measurement of choiceworthiness as well as a zero point of choiceworthiness (or a *benchmark*) in the following way: the measure of the choiceworthiness of *a* for *S* (*relative to* <u>*any*</u> *explicitly given utility unit and zero point of utility*) is the *difference in utility* between *a* and some *benchmark* for *a*, such that (1) *a* is choiceworthy for *S* if and only if the difference in utility between *a* and the benchmark for *a* is equal to or greater than zero (and not choiceworthy otherwise), and (2) the degree of choiceworthiness of *a* for *S* is the difference in utility between *a* and the benchmark for *a*. In other words, the measure of the choiceworthiness of *a* for *S* is the degree to which *a* is worthy of being chosen over the benchmark for *a*. The benchmark for *a* can be, for example, some option in the set of available options, such as the status quo,[10] or some average of the utilities of the available options. Choiceworthiness is thus

[10] Jacob M. Nebel, "Status Quo Bias, Rationality, and Conservatism about Value," *Ethics* 125 (2015): 449–76.

a relative concept.[11] As will become clear in what follows, the concept of choiceworthiness itself presupposes a given benchmark (or zero point of choiceworthiness).

If there are any alternatives to *a* which carry a greater utility than does *a*, then the benchmark for *a* is whichever alternative to *a* carries the greatest utility. Indeed, if there are any alternatives to *a* with a greater utility than *a*, then, in accordance with the UM Principle, *a* is not choiceworthy for *S*. But if *a* is not choiceworthy for *S*, then how choiceworthy *a* is for *S* is simply how *a* compares to whichever alternative is choiceworthy for *S* (or, per the UM Principle, whichever alternative to *a* carries the greatest utility).[12] This is consistent with the basic "regret minimization" rational choice model introduced in section I.

If there are *not* any alternatives to *a* which carry a greater utility than does *a*, then, according to the regret minimization model, the benchmark for *a* is *a*. Against this model, I will now argue that if there are *not* any alternatives to *a* which carry a greater utility than does *a*, then the benchmark for *a* still has to be whichever alternative to *a* carries the greatest utility.[13]

Let us consider two decision situations (or setups): 1 and 2. In each situation, *S* is faced with the same three options: *a*, *b*, and *c*. What's more, in each situation, *S* assigns probability 1 to a given state of the world (and not the same state for both situations). If that state of the world is realized, then *S* assigns the following utilities to the set of options:

$$1 : a(100), b(-10,000), c(-10,000)$$
$$2 : a(100), b(99), c(99)$$

Per the UM Principle, *a* is choiceworthy for *S* in both situations 1 and 2. *a* is also more choiceworthy for *S* in 1 than in 2—that is to say, it is more imperative for *S* to choose *a* if she is in situation 1 than if she is in situation 2. In 2, *S* misses out

[11] Ralph Wedgwood ("Must Rational Intentions Maximize Utility?" *Philosophical Explorations* 20 (2017): 1–20) relies on considerations of incommensurability to argue for the same idea: "the choiceworthiness of options is relative to choice situations". Larry S. Temkin (*Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford UP, 2012)) also addresses this idea: what he calls the "Essentially Comparative View."

[12] See the discussion of the notion of "shortfall from optimality" in Ralph Wedgwood, "Choosing Rationally and Choosing Correctly," *Weakness of Will and Practical Irrationality*, ed. Sarah Stroud and Christine Tappolet (Oxford UP, 2003). For related discussions of the notion of "degrees of wrongness" in ethics, see (e.g.) J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin, 1977), 205; Bjorn Eriksson, "Utilitarianism for Sinners," *American Philosophical Quarterly* 34 (1997): 218–19; Todd Calder, "Kant and Degrees of Wrongness," *The Journal of Value Inquiry* 39 (2005): 229; and Johan E. Gustafsson, "Consequentialism with Wrongness Depending on the Difficulty of Doing Better," *Thought* 5 (2016): 108–09.

[13] As far as I know, this idea has not been explored in the published literature.

on only 1 utile by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *b* or *c*), whereas in 1, *S* misses out on 10,100 utiles by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *b* or *c*). Another way of putting it is that *a* is more choiceworthy in 1 than in 2 because *a* is more worthy of being chosen over the best alternative to *a* in 1 than in 2.

Let me now briefly introduce Ralph Wedgwood's *Benchmark Theory* (*BT*).[14] The basic idea of BT is to rank choice options (in terms of how choiceworthy they are) according to their *expected comparative value*, where the comparative value of an option is its *value* (broadly construed) in some state of the world compared to a benchmark for that state of the world. Wedgwood identifies the benchmark as an average of the options' values within a given state of the world. He emphasizes that all statewise dominated options and more generally, "all the options that *do not deserve to be taken seriously*" (2664) should be excluded from consideration at the outset.[15] Wedgwood explicitly rejects the idea that the value of an option is its utility. Nevertheless, it is interesting to see how BT (henceforth, *BT**) fairs when the value of an option is understood to be its utility.

Coming back to our example, we can see that BT* agrees with the verdict that *a* is choiceworthy for *S* in situations 1 and 2, but *not* with the verdict that *a* is more choiceworthy for *S* in 1 than in 2. According to BT*, *a* is equally choiceworthy for *S* in situations 1 and 2 since *b* and *c* are strictly dominated by *a* in both 1 and 2 and are therefore excluded from consideration at the outset. If *b* and *c* are *not* excluded from consideration and the benchmark is identified as an average of the values (or utilities) of all the available options, then this alternative approach agrees with our verdict: *a* is more choiceworthy for *S* in 1 than in 2.

Here is a different example:

$$1 : a(100), b(-100), c(-100)$$
$$2 : a(100), b(-500), c(100)$$

Per the UM Principle, *a* is choiceworthy for *S* in both situations 1 and 2. *a* is also more choiceworthy for *S* in 1 than in 2—in other words, it is more imperative for *S* to choose *a* if she is in situation 1 than if she is in situation 2. In 2, *a* is merely optional—*S* misses out on *zero* utiles by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *c*)—whereas in 1, *a* is *not* optional—*S* misses out on 200 utiles by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *b*

[14] Ralph Wedgwood, "Gandalf's Solution to the Newcomb Problem," *Synthese* 190 (2013): 2643–75.
[15] For critiques of BT, see Robert Bassett, "A Critique of Benchmark Theory," *Synthese* 192 (2015): 241–67 and Rachael Briggs, "Decision-Theoretic Paradoxes as Voting Paradoxes," *Philosophical Review* 119 (2010): 1–30.

or *c*). Again, *a* is more choiceworthy in 1 than in 2 because *a* is more worthy of being chosen over the best alternative to *a* in 1 than in 2.

BT* agrees with both verdicts: *a* is choiceworthy for *S* in 1 and 2, and *a* is more choiceworthy for *S* in 1 than in 2. However, if the benchmark is defined as an average of the values (or utilities) of all the available options (whether strictly dominated or not), then this alternative approach does *not* agree with our verdict: *a* is more choiceworthy for *S* in 1 than in 2. The two examples just laid out, when taken together, make for an effective counterexample to BT*.

Another very similar example:

$$1 : a(100), b(-100), c(-100)$$

$$2 : a(100), b(-500), c(99)$$

Per the UM Principle, *a* is choiceworthy for *S* in both situations 1 and 2. *a* is also more choiceworthy for *S* in 1 than in 2—that is, it is more imperative for *S* to choose *a* if she is in situation 1 than if she is in situation 2. In 2, *S* misses out on only 1 utile by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *c*), whereas in 1, *S* misses out on 200 utiles by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *b* or *c*). Once again, *a* is more choiceworthy in 1 than in 2 because *a* is more worthy of being chosen over the best alternative to *a* in 1 than in 2.

BT* agrees with the verdict that *a* is choiceworthy for *S* in 1 and 2, but *not* with the verdict that *a* is more choiceworthy for *S* in 1 than in 2. According to BT*, *a* is equally choiceworthy for *S* in situations 1 and 2 since *b* and *c* are strictly dominated by *a* in both 1 and 2 and are therefore excluded from consideration at the outset. If *b* and *c* are *not* excluded from consideration and the benchmark is identified as an average of the values (or utilities) of all the options, then *a* is more choiceworthy for *S* in 2 than in 1. I take this to be a further counterexample to BT*.

One final example:

$$1 : a(100), b(-100), c(100)$$

$$2 : a(100), b(99), c(100)$$

Per the UM Principle, *a* is choiceworthy for *S* in both situations 1 and 2. *a* is also equally choiceworthy for *S* in both situations—that is to say, it is just as imperative for *S* to choose *a* if she is in situation 1 as it is if she is in situation 2. In both situations, *a* is merely optional—*S* misses out on *zero* utiles by not choosing *a*, but instead choosing the best alternative to *a* (i.e., *c*). To put it another way, *a* is just as choiceworthy in 1 as it is in 2 because *a* is just as worthy of being chosen over the best alternative to *a* in 1 as it is in 2. (BT* agrees with both verdicts.)

These four examples serve to illustrate that if there are *not* any alternatives to *a* with a greater utility than *a*, then how choiceworthy *a* is depends on how much utility *S* would miss out on by not choosing *a*, but instead choosing the best alternative to *a*. The greater the amount of utility *S* would miss out on by not choosing *a*, but instead choosing the best alternative to *a*, the more choiceworthy *a* becomes. Thus, the benchmark for *a* must be whichever alternative to *a* carries the highest utility.

What follows is that whether or not there are any alternatives to *a* which carry a greater utility than does *a*, the benchmark for *a* has to be whichever alternative to *a* carries the greatest utility.[16] Therefore, for any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* (relative to any explicitly given utility unit and zero point of utility) is the *CU* of *a* (in the state of the world to which *S* assigns probability 1). The CU of *a* is the difference in utility between *a* and whichever alternative to *a* carries the greatest utility (or one of them in the event that several alternatives are tied). As previously indicated, I will refer to this principle as the *CU Principle*.[17] Like the Utility Principle, the CU Principle entails the UM Principle. Henceforth, *c-utiles* are defined as units of CU.

In light of the CU Principle, the Utility Principle can be falsified. If the Utility Principle were true, then in accordance with the Measurement Principle, it would be the case that for *any* given decision situation, there is at least one specification of a utility unit and zero point of utility such that it is possible to ascertain how choiceworthy any available option is (for *S*) by solely considering its utility value in relation to that specification of a utility unit and zero point of utility. In other words, it would be the case that for *any* given decision situation, there is at least one specification of a utility unit and zero point of utility such that (1) any available option is choiceworthy (for *S*) if and only if its utility value is equal to or greater than zero (and not choiceworthy otherwise) and (2) *the degree of choiceworthiness of any available option (for S) is its utility value*. As we will now see, that is not the case. Let us consider the following decision setup: *S* is faced with three options: *a*, *b*, and *c*. What's more, *S* assigns probability 1 to a given state of the world. If that state of the world is realized, then *S* assigns the following utilities to the available options: *a* (0), *b* (–100), *c* (–1,000). Therefore, no matter what zero point of utility is selected, *S* assigns the following utility intervals between the available options: between *a* and *b*, *S* assigns a positive interval of

[16] This means that there is no unique benchmark for a given choice situation. Instead, the benchmark is relative to a specific choice option. The benchmark for *a* may be some alternative, *b*, and the benchmark for *b* may be *a*.

[17] The CU Principle is also true if utilities are taken as primitive (rather than derived from a representation theorem), provided that the UM Principle is true when utilities are understood in those terms.

100 utiles, between *b* and *c*, *S* assigns a positive interval of 900 utiles and between *a* and *c*, *S* assigns a positive interval of 1,000 utiles. Per the CU Principle, the degrees of choiceworthiness of the available options are as follows: *a* (100), *b* (–100), *c* (–1,000). Therefore, the differences between the degrees of choiceworthiness of the available options are as follows: between *a* and *b*, the difference is 200 c-utiles, between *b* and *c*, the difference is 900 c-utiles and between *a* and *c*, the difference is 1,100 c-utiles. Since the utility intervals and the differences in degrees of choiceworthiness are at variance, we have a decision situation where no matter what zero point of utility (and what utility unit) is selected, it is *not* the case that the degree of choiceworthiness of any available option is its utility value.

In closing this section, it should be mentioned that the idea of calculating differences between the utility of an option under consideration and the utilities of its alternatives in the choice situation has been explored in the economic modeling literature.[18] Incidentally, CU should be distinguished from the purely descriptive economic concept of *opportunity cost*. For any agent, *S*, let *a* be the highest-valued choice option available to *S*. The CU of *a,* for *S*, is the value of whatever *additional benefit S* would enjoy by choosing *a* over the highest-valued alternative to *a*. By contrast, the opportunity cost of *a*, for *S*, is the value of whatever *cost S* would incur by choosing *a* over the highest-valued alternative to *a*, where this includes the *total value* of the highest-valued alternative to *a*.[19]

III. ECU THEORY

As I showed in section I, we require a graded (quantitative) measure of how choiceworthy options are. When we move beyond decision-making under certainty, it is most natural, in light of the CU Principle, to identify the measure of an option's choiceworthiness as expressing that option's *expected choiceworthiness*, or *ECU*, that is to say, the *expected value*, or the probability-weighted average of all possible values, of that option's choiceworthiness, or CU, *in the actual state of the world*. This roughly encapsulates ECU Theory.

As a first approximation, therefore, ECU Theory says that for any agent, *S*, and for any choice option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* (relative to any explicitly given utility unit and zero point of utility) is the ECU

---

[18] Zhang et al., "Modeling Traveler Choice Behavior Using the Concepts of Relative Utility and Relative Interest," *Transportation Research Part B* 38 (2004): 215–34.

[19] Robert H. Frank, "The Opportunity Cost of Economics Education," *The New York Times*, 2005, http://www.nytimes.com/2005/09/01/business/the-opportunity-cost-of-economics-education.html and David R. Henderson, "Opportunity Cost," *The Concise Encyclopedia of Economics*, Library of Economics and Liberty, 2008, http://www.econlib.org/library/Enc/OpportunityCost.html.

of a. The *ECU* of an option, a, in a decision problem with n states is formally defined as:

$$\mathrm{ECU}(a) = \sum\nolimits_{i=1}^{n} (U(a,\ s_i) - U(bm(a), s_i))P(s_i)$$

where $U(a, s_i)$ denotes the utility of option a when state s_i is actual, $U(bm(a), s_i)$ denotes the utility of the benchmark for a when state s_i is actual (i.e., the utility in state s_i of whichever alternative(s) to a have the highest utility in state s_i), and $P(s_i)$ denotes the probability assigned to state s_i.[20] The CU Principle is straightforwardly entailed by ECU Theory. Furthermore, ECU Theory presupposes that the states of the world in any decision problem are probabilistically independent of the agent's choices.

ECU Theory, as formulated above, is not quite right though. In accordance with the Measurement Principle, if the measure of the choiceworthiness of options is their ECU, then only options with ECU equal to or greater than zero can be choiceworthy. However, as I will illustrate in section IV, there will always be cases (regardless of what utility unit and zero point of utility are specified) where every option in a decision situation *under uncertainty* has negative ECU. Since at least one option in a decision situation must be choiceworthy—the one with the highest degree of choiceworthiness (or one of them in the event that several alternatives are tied)—ECU Theory, as defined above, is false in decision cases *under uncertainty*.

By the same line of reasoning as employed in section II, we reach the following conclusion: for any agent, S, faced with any decision *under uncertainty* and for any choice option, a, for S, the measure of the choiceworthiness of a for S (relative to any explicitly given utility unit and zero point of utility) is the *comparative expected choiceworthiness*, or *comparative expected comparative utility* (*CECU*), of a, that is to say, the difference in ECU between a and whichever alternative to a carries the greatest ECU (or one of them in the event that several alternatives are tied). Let us call this principle the *CECU Principle*. For any two alternative options, a and b, a's CECU is greater than b's if and only if a's ECU is greater than b's, and a's CECU is equal to b's if and only if a's ECU is equal to b's. We are now in a position to precisely define ECU Theory: ECU Theory is the conjunction of the CU Principle (for decisions under certainty) and the CECU Principle (for decisions under uncertainty).

[20] In other words, for any number of alternative options, a, b, c, d, and e, one calculates the ECU of a as follows: for each state of the world, one subtracts a's utility from the utility of b, c, d, or e, whichever of b, c, d, and e maximizes utility in that state, and one multiplies the result by the probability that one assigns to that state; finally, one sums the totals for every state.

To demonstrate how to apply this new decision rule (i.e., ECU Theory) to a concrete decision problem, let us consider the following case: An agent, *S*, is faced with a choice between two independent options or gambles: one option, *a*, offering a 0.01 probability of winning a prize worth 1,500 utiles (and nothing otherwise), and one option, *b*, offering a 0.02 probability of winning a prize worth 700 utiles (and nothing otherwise). According to ECU Theory, *S* should choose option *a*, since its CECU is equal to or greater than zero (1– [–1] = 2).

The ECUs of options *a* and *b* are given by the two equations below. The following notation is used: s_1 denotes the state "If *S* chooses *a*, then *S* will not win the prize (0 utiles) and if *S* chooses *b*, then *S* will not win the prize (0 utiles)" (probability: $0.99 \times 0.98 = 0.97$), s_2 denotes the state "If *S* chooses *a*, then *S* will not win the prize (0 utiles) and if *S* chooses *b*, then *S* will win the prize (700 utiles)" (probability: $0.99 \times 0.02 = 0.0198$), s_3 denotes the state "If *S* chooses *a*, then *S* will win the prize (1,500 utiles) and if *S* chooses *b*, then *S* will not win the prize (0 utiles)" (probability: $0.01 \times 0.98 = 0.0098$), s_4 denotes the state "If *S* chooses *a*, then *S* will win the prize (1,500 utiles) and if *S* chooses *b*, then *S* will win the prize (700 utiles)" (probability: $0.01 \times 0.02 = 0.0002$), $P(s_1)$ denotes the probability of s_1, and $U(a, s_1)$ denotes the utility of option *a* when s_1 is actual.

$$\begin{aligned}\mathrm{ECU}(a) = {} & (\mathrm{U}(a,s_1)-\mathrm{U}(b,s_1)) \times \mathrm{P}(s_1) + (\mathrm{U}(a,s_2)-\mathrm{U}(b,s_2)) \\ & \times \mathrm{P}(s_2) + (\mathrm{U}(a,s_3)-\mathrm{U}(b,s_3)) \times \mathrm{P}(s_3) + (\mathrm{U}(a,s_4)-\mathrm{U}(b,s_4)) \\ & \times \mathrm{P}(s_4) = 1 \ \text{c}-\text{utile}\end{aligned}$$

$$\begin{aligned}\mathrm{ECU}(b) = {} & (\mathrm{U}(b,s_1)-\mathrm{U}(a,s_1)) \times \mathrm{P}(s_1) + (\mathrm{U}(b,s_2)-\mathrm{U}(a,s_2)) \\ & \times \mathrm{P}(s_2) + (\mathrm{U}(b,s_3)-\mathrm{U}(a,s_3)) \times \mathrm{P}(s_3) + (\mathrm{U}(b,s_4)-\mathrm{U}(a,s_4)) \\ & \times \mathrm{P}(s_4) = -1 \ \text{c}-\text{utiles}\end{aligned}$$

One specific line of argument in support of ECU Theory is that (contrary to EU Theory) ECU Theory agrees with (and entails) Wedgwood's *Gandalf's Principle*: the choiceworthiness of an option *in a given state of the world* should be measured relative only to the values of the other options *in that state*, and not to the values of the options in other states. According to Wedgwood,

> to make a rational choice in [cases involving uncertainty], one *does not need to consider* whether one is in a nice state of nature or a nasty one. All that one needs to consider are the *degrees* to which each of the available options is better (or worse) than the available alternatives *within* each of the relevant states of nature. Admittedly, when one is uncertain which state of nature one is in, one must make *some* comparisons across the states of nature. But since one does not even need to know whether one is in a nice state of nature or a nasty one, it seems that the only relevant comparisons are comparisons of the *differences* in levels of goodness between the various options *within* each

> state of nature with the *differences* between those options within each of the other states of nature – not any comparisons of *absolute* levels of goodness across different states of nature.[21]

Although Wedgwood uses terms such as "better," "worse," and "levels of goodness" in his explication of Gandalf's Principle, the principle can be expressed equally well using replacement terms such as "preferred," "dispreferred," and "levels of utility."

Gandalf's Principle is an eminently reasonable principle.[22] In a paper critiquing Wedgwood's BT, Robert Bassett concurs: "Gandalf's Principle strikes me as an eminently sensible principle to incorporate into rational decision-making."[23] There is, however, one alternative decision theory which agrees with (and entails) both the CU Principle and Gandalf's Principle and which has some *prima facie* plausibility—*Maximum Likelihood Comparative Utility* (*MLCU*) *Theory*: for any agent, *S*, and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* (relative to any explicitly given utility unit and zero point of utility) is the *most likely value* of *a*'s choiceworthiness (or CU) in the actual state of the world. We require a further argument to rule out this decision theory.

This brings me to the following decision case: Let us suppose that an agent, *S*, is faced with three choice options: *a*, *b*, and *c*. *S* assigns probability 0.51 to a state of the world, *A*, and 0.49 to a state of the world, *B*. If state *A* or state *B* is realized, then *S* assigns the following utilities to the set of options:

$$A : a(110), b(80), c(100)$$
$$B : a(-1,000), b(110), c(100)$$

According to MLCU Theory, *a* is uniquely choiceworthy for *S*, since state *A* is more likely to obtain than state *B* and the CU of option *a* in state *A* is greater than that of any other available option. Yet, it is clear that choosing option *a* is a mistake, since state *B* is almost as likely to obtain as state *A* and the comparative *dis*-utility of option *a* in state *B* is very high (–1,110 c-utiles). I take this to be an effective counterexample to MLCU Theory.

IV. THE FAILURE OF EU THEORY

In this final section, I will assess the cogency of standard rational choice theory (i.e., EU Theory) in view of ECU Theory. Let us begin by giving a precise definition of EU Theory: for any agent, *S*, and for any number of alternative choice

[21] Wedgwood (2013): 2654.
[22] For an illuminating discussion of Gandalf's principle, see Wedgwood (2013): 2652–55.
[23] Bassett (2015): 241–67.

options, *a*, *b*, *c*, *d*, and *e*, for *S*, (1) *a* is more choiceworthy than *b*, for *S*, if and only if *a*'s EU is greater than *b*'s, (2) *a* is just as choiceworthy as *b*, for *S*, if and only if *a*'s EU is equal to *b*'s, (3) (given that utility is assumed to be cardinal) the extent to which *a* is more choiceworthy than *b*, for *S*, is the difference in EU between *a* and *b*, and (4) *a* is choiceworthy for *S* if and only if *a* maximizes EU over the space of all alternatives.

Decision-theoretic representation theorems—such as those of von Neumann and Morgenstern, Savage, and Bolker and Jeffrey[24]—show that if an agent fails to prefer choice options with higher EU, then that agent violates at least one of several axioms of *rational preference*,[25] one of which entails the Independence of irrelevant alternatives (IIA) (as I indicated in section I). These representation theorems are supposed to provide arguments for EU Theory (as a theory of rational *choice*), the underlying assumption being that "rational preference" talk can be translated into "choiceworthiness" talk. In what follows, I will demonstrate that such an assumption is false. I will do so by showing that ECU Theory's verdicts sometimes deviate from those of EU Theory.

Let a *finite decision case* be a decision problem where there are only finitely many states and no infinite utilities. In all finite decision cases requiring a choice between *only two* alternative options, ECU Theory delivers the same verdicts as EU Theory. However, in a number of finite decision cases requiring a choice between *more than two* alternative options, ECU Theory gives different verdicts from those of EU Theory.[26] Let us consider the following example: an agent, *S*, is faced with five choice options: *a*, *b*, *c*, *d*, and *e*. *S* assigns probability 0.5 to a state of the world, *A*, and 0.5 to a state of the world, *B*. If state *A* or state *B* is realized, then *S* assigns the following utilities to the set of options:

[24] John Von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton UP, 1944); Leonard Savage, *The Foundations of Statistics* (Wiley, 1954); Ethan Bolker, "Functions Resembling Quotients of Measures," *Transactions of the American Mathematical Society* 2 (1966): 292–312; Richard Jeffrey, *The Logic of Decision*, 2nd ed. (Chicago: U of Chicago P, 1983).

[25] Following the formulation of Briggs ("Normative Theories of Rational Choice: Expected Utility," *The Stanford Encyclopedia of Philosophy*).

[26] In a number of decision cases where there are infinitely many states with only finite utilities attached (e.g., the St. Petersburg game), ECU Theory inherits the advantages of Mark Colyvan's *Relative Expectation Theory* over EU Theory. More specifically, in such (infinite) decision cases, ECU Theory delivers the intuitively correct verdicts, whereas EU Theory delivers none (see Mark Colyvan, "Relative Expectation Theory," *Journal of Philosophy* 105 (2008): 37–44; Colyvan and Hájek (2016): 838–39).

$$A : a(2), b(5), c(6), d(8), e(10)$$
$$B : a(10), b(8), c(6), d(4), e(2)$$

According to EU Theory, *b* is more choiceworthy than *a*, for *S*, since the EU of *b* (6.5 utiles) is greater than that of *a* (6 utiles). In fact, according to EU Theory, *b* is choiceworthy tout court since its EU is greater than that of every other option.

$$\mathrm{EU}(a) = (\mathrm{U}(a, A) \times \mathrm{P}(A)) + (\mathrm{U}(a, B) \times \mathrm{P}(B)) = 6 \text{ utiles}$$
$$\mathrm{EU}(b) = (\mathrm{U}(b, A) \times \mathrm{P}(A)) + (\mathrm{U}(b, B) \times \mathrm{P}(B)) = 6.5 \text{ utiles}$$
$$\mathrm{EU}(c) = (\mathrm{U}(c, A) \times \mathrm{P}(A)) + (\mathrm{U}(c, B) \times \mathrm{P}(B)) = 6 \text{ utiles}$$
$$\mathrm{EU}(d) = (\mathrm{U}(d, A) \times \mathrm{P}(A)) + (\mathrm{U}(d, B) \times \mathrm{P}(B)) = 6 \text{ utiles}$$
$$\mathrm{EU}(e) = (\mathrm{U}(e, A) \times \mathrm{P}(A)) + (\mathrm{U}(e, B) \times \mathrm{P}(B)) = 6 \text{ utiles}$$

By contrast, according to ECU Theory, *a* is more choiceworthy than *b*, for *S*, since the ECU of *a* (–3 c-utiles) is greater than that of *b* (–3.5 c-utiles). In fact, according to ECU Theory, *a* is choiceworthy tout court, since its CECU is equal to or greater than zero ([–3] – [–3] = 0).

$$\mathrm{ECU}(a) = ((\mathrm{U}(a, A) - \mathrm{U}(e, A)) \times \mathrm{P}(A)) + ((\mathrm{U}(a, B) - \mathrm{U}(b, B)) \times \mathrm{P}(B)) = -3 \text{ c}-\text{utiles}$$
$$\mathrm{ECU}(b) = ((\mathrm{U}(b, A) - \mathrm{U}(e, A)) \times \mathrm{P}(A)) + ((\mathrm{U}(b, B) - \mathrm{U}(a, B)) \times \mathrm{P}(B)) = -3.5 \text{ c}-\text{utiles}$$
$$\mathrm{ECU}(c) = ((\mathrm{U}(c, A) - \mathrm{U}(e, A)) \times \mathrm{P}(A)) + ((\mathrm{U}(c, B) - \mathrm{U}(a, B)) \times \mathrm{P}(B)) = -4 \text{ c}-\text{utiles}$$
$$\mathrm{ECU}(d) = ((\mathrm{U}(d, A) - \mathrm{U}(e, A)) \times \mathrm{P}(A)) + ((\mathrm{U}(d, B) - \mathrm{U}(a, B)) \times \mathrm{P}(B)) = -4 \text{ c}-\text{utiles}$$
$$\mathrm{ECU}(e) = ((\mathrm{U}(e, A) - \mathrm{U}(d, A)) \times \mathrm{P}(A)) + ((\mathrm{U}(e, B) - \mathrm{U}(a, B)) \times \mathrm{P}(B)) = -3 \text{ c}-\text{utiles}$$

ECU Theory gives different verdicts from those of EU Theory because ECU Theory, contrary to EU Theory, violates the IIA (for choiceworthiness evaluations). According to this principle, for any decision situation, *T*, and for any choice option, *a*, in *T*, if *a* is choiceworthy in *T*, then *a* is also choiceworthy in *T* if some other option(s) are eliminated from the pool of options in *T*. Likewise, if *a* is not choiceworthy in *T*, then *a* is also not choiceworthy in *T* if some other option(s) are added to the pool of options in *T*.[27] Let us consider

[27] Kenneth Arrow, *Social Choice and Individual Values* (New York: Wiley, 1951); R. Duncan Luce and Howard Raiffa, *Games and Decisions: Introduction and Critical Survey* (New York: Wiley, 1967). For a critique of the IIA, see in particular Amartya Sen, "Internal Consistency of Choice," *Econometrica* 61 (1993): 495–521. Sen refers to this principle as "basic contraction consistency (Property α)". See also Wedgwood's critique (2013: 2668–70).

again the previous decision situation. In such a situation, ECU Theory dictates that a is choiceworthy. However, if options c, d, and e are eliminated from the pool of options, then b is choiceworthy according to ECU Theory, as shown below:

$$A : a(2), b(5)$$
$$B : a(10), b(8)$$
$$\mathrm{ECU}(a) = ((\mathrm{U}(a,A) - \mathrm{U}(b,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(a,B) - \mathrm{U}(b,B)) \times \mathrm{P}(B)) = -0.5 \ \text{c}-\text{utiles}$$
$$\mathrm{ECU}(b) = ((\mathrm{U}(b,A) - \mathrm{U}(a,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(b,B) - \mathrm{U}(a,B)) \times \mathrm{P}(B)) = 0.5 \ \text{c}-\text{utiles}$$

Here is another example where ECU Theory violates the IIA:[28] An agent, S, is faced with two choice options: a and b. S assigns probability 0.001 to a state of the world, A, and 0.999 to a state of the world, B. If state A or state B is realized, then S assigns the following utilities to the set of options:

$$A : a(1000), b(0)$$
$$B : a(1), b(2)$$

According to ECU Theory, a is choiceworthy tout court, since its ECU is greater than that of any other option.

$$\mathrm{ECU}(a) = ((\mathrm{U}(a,A) - \mathrm{U}(b,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(a,B) - \mathrm{U}(b,B)) \times \mathrm{P}(B)) = 0.001 \ \text{c}-\text{utiles}$$
$$\mathrm{ECU}(b) = ((\mathrm{U}(b,A) - \mathrm{U}(a,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(b,B) - \mathrm{U}(a,B)) \times \mathrm{P}(B)) = -0.001 \ \text{c}-\text{utiles}$$

Let us now introduce a third choice option (c) in the decision situation, all else being the same:

$$A : a(1000), b(0), c(900)$$
$$B : a(1), b(2), c(0)$$

[28] Thanks to an anonymous reviewer for suggesting this example.

In this new decision situation, *b* is choiceworthy according to ECU Theory, since *b*'s ECU is greater than that of any other option.

$$\mathrm{ECU}(a) = ((\mathrm{U}(a,A) - \mathrm{U}(c,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(a,B) - \mathrm{U}(b,B)) \times \mathrm{P}(B)) = -0.899 \ \text{c-utiles}$$

$$\mathrm{ECU}(b) = ((\mathrm{U}(b,A) - \mathrm{U}(a,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(b,B) - \mathrm{U}(a,B)) \times \mathrm{P}(B)) = -0.001 \ \text{c-utiles}$$

$$\mathrm{ECU}(c) = ((\mathrm{U}(c,A) - \mathrm{U}(a,A)) \times \mathrm{P}(A)) + ((\mathrm{U}(c,B) - \mathrm{U}(b,B)) \times \mathrm{P}(B)) = -2.098 \ \text{c-utiles}$$

This example is particularly telling because option *c* is statewise dominated by *a*. Whether state *A* or state *B* is actual, option *a* is strictly preferred to option *c*. Yet, introducing option *c* in the decision situation changes ECU Theory's verdict: *b*, instead of *a*, is uniquely choiceworthy.

This gives rise to a worry. Without the IIA, it is possible to make up alternatives in any choice set—adding on, for example, choosing to worship a (highly improbable) deviant deity in the hopes of future riches—and these manufactured alternatives would be altering the degrees of choiceworthiness of reasonable options.[29] This opens the door to strategic manipulation in the decision process. The worry can be partially overcome, however, if we accept Nicholas Smith's theory of *Rationally Negligible Probabilities*: for any given decision, any outcome with probability $\leq p$, where p is very close to 0, can be rationally excluded from consideration in the decision process.[30] As such, the *very improbable* outcomes of manufactured alternatives (e.g., to become rich as a result of worshiping a deviant deity) cannot alter the degrees of choiceworthiness of the other available options in the choice set.

Just as ECU Theory delivers verdicts which are at odds with EU Theory, ECU Theory also supplies a more discriminating measure of the intervals in rankings of *more than two* choice options. Let us consider four choice situations involving decisions under certainty:

---

29 Thanks to Douglas Lackey for raising this point and for wording suggestions.

30 Nicholas J. J. Smith, "Is Evaluative Compositionality a Requirement of Rationality?" *Mind* 123 (2014): 457–502; idem, "Infinite Decisions and Rationally Negligible Probabilities," *Mind* 125 (2016): 1199–212; Adam Chalmers, *An Offer You Can't (Rationally) Refuse: Systematically Exploiting Utility-Maximisers with Malicious Gambles*, 2017, Thesis, http://hdl.handle.net/2123/16131.

$$1 : a(5), b(1), c(1), d(1), e(1)$$
$$2 : a(5), b(1), c(2), d(3), e(3)$$
$$3 : a(5), b(1), c(2), d(3), e(5)$$
$$4 : a(5), b(1), c(2), d(3), e(8)$$

The difference in CU between *a* and *b* is greater in situation 1 ((5 – 1) – (1 – 5) = 8 c-utiles) than in situation 2 ((5 – 3) – (1 – 5) = 6 c-utiles), situation 3 ((5 – 5) – (1 – 5) = 4 c-utiles) and situation 4 ((5 – 8) – (1 – 8) = 4 c-utiles), whereas the difference in utility between *a* and *b* is the same in all four situations (4 utiles). Therefore, compared to utility, CU is a more discriminating measure of the extent to which *a* is more choiceworthy than *b* in situations 1–4. What's more, there do not appear to be any contrary cases where CU (or CECU) gives a *less* differentiated picture than does utility (or EU).

What the foregoing comparisons (between EU Theory and ECU Theory) show is that rational preference is not a reliable indicator of choiceworthiness. That is because whereas the criterion of rational preference satisfies the IIA (by stipulative definition), the criterion of choiceworthiness (i.e., ECU Theory) violates that principle, as demonstrated above. Thus, not only do representation theorems not support EU Theory, but EU Theory also fails as a theory of rational *choice*. EU Theory's verdicts track what it is rational to *prefer* (in accordance with the IIA), but do not always track what it is rational to *choose*, that is, what is *choiceworthy*.

It is important to emphasize that the proposed criterion of choice (i.e., choiceworthiness) is independent from the standard choice criterion (i.e., rational preference). The latter is not shown here to violate the assumptions, for example, the IIA, which are needed to derive utilities from preferences via a representation theorem. ECU Theory stands vindicated.

In recent years, several alternatives to EU Theory have been proposed, for example, Mark Colyvan's *Relative Expectation Theory* (*RET*), Paul Bartha's *Relative Utility Theory* (*RUT*), and Lara Buchak's *Risk-Weighted Expected Utility* (*REU*) *Theory*.[31] In all finite decision cases, RET and RUT deliver the same verdicts as EU Theory. As for REU Theory, it *can* deliver the same verdicts as EU Theory, depending on the risk attitude of the agent equipped with the REU decision rule. These alternative "rational preference" tracking decision theories are therefore subject to the same objection as that leveled here against EU Theory: they fall short as theories of rational *choice*.

[31] Colyvan, "Relative Expectation Theory"; Paul Bartha, "Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities," *Synthese* 154 (2007): 5–52; idem, "Making do Without Expectations," *Mind* 125 (2016): 799–827); Lara Buchak, *Risk and Rationality* (Oxford UP, 2013).