ORIGINAL ARTICLE

Expected comparative utility theory: A new theory of rational choice under risk

David Robert

LaSalle, QC, Canada

Correspondence
Email: jeandavidrobert@hotmail.com

Abstract

In this paper, I argue for a new normative theory of rational choice under risk, namely expected comparative utility (ECU) theory. I show that for any choice option, *a*, and for any state of the world, *G*, the measure of the choiceworthiness of *a* in *G* is the comparative utility (CU) of *a* in *G*—that is, the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*. On the basis of this principle, I argue that for any agent, *S*, faced with any decision under risk, *S* should rank his or her decision options (in terms of how choiceworthy they are) according to their comparative expected comparative utility (CECU) and should choose whichever option carries the greatest CECU (or one of them in the event that several alternatives are tied). For any option, *a*, *a*'s CECU is the difference between its ECU and that of whichever alternative(s) to *a* carry the greatest ECU, where *a*'s ECU is a probability-weighted sum of *a*'s CUs across the various possible states of the world. I show that in some ordinary decisions under risk, ECU theory delivers different verdicts from those of standard decision theory.

This article is a revised version of my previously published article "Robert, D. (2021). A restatement of expected comparative utility theory: A new theory of rational choice under risk. *The Philosophical Forum*, 52, 221–243. https://doi.org/10.1111/phil.12299"

1 INTRODUCTION

Standard decision theory, otherwise known as *expected utility* (*EU*) *theory*, counsels agents to rank their choice options (in order of preference) according to their EU and to choose whichever option carries the greatest EU (or one of them in the event that several alternatives are tied). The EU of an option is a probability-weighted sum of each of its possible utilities. EU theory has been the dominant normative theory of rational choice under risk since the 18th century (Bernoulli, 1738), and in more recent times (1920s–), has received foundational support from both economists and philosophers (Bolker, 1966; Jeffrey, 1983; Joyce, 1999; Ramsey, 1931; Savage, 1954; von Neumann & Morgenstern, 1947).[1]

In this paper, I will argue for a new normative alternative to EU theory. I will argue that from the fact that we need a graded, quantitative measure of *choiceworthiness* for decisions under certainty and decisions under risk,[2] it follows that we need a new normative theory of rational choice under risk, namely *expected comparative utility* (*ECU*) *theory*. I will show that for any choice option, *a*, and for any state of the world, *G*, the measure of the choiceworthiness of *a*, in *G*, is the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*. This difference in utility is what I will call the *comparative utility* (*CU*) of *a*. For any agent, *S*, faced with any decision under *certainty*, ECU theory counsels *S* to measure and rank his or her options (in terms of how choiceworthy they are) according to their CU (this is what I will call the *CU principle*) and to choose whichever option carries the greatest CU (or one of them in the event that several alternatives are tied). For any agent, *S*, faced with any decision under *risk*, ECU theory counsels *S* to measure and rank his or her options (in terms of how choiceworthy they are) according to their *comparative expected comparative utility* (*CECU*) (this is what I will call the *CECU principle*) and to choose whichever option carries the greatest CECU (or one of them in the event that several alternatives are tied). For any option, *a*, *a*'s CECU is the difference between its ECU and that of whichever alternative(s) to *a* carry the greatest ECU, where *a*'s ECU is a probability-weighted sum of *a*'s CUs across the various possible states of the world. In this paper, I will show that in some ordinary decisions under risk, ECU theory gives different verdicts from those of EU theory and that EU theory therefore fails as a normative theory of rational choice under risk.

The idea of calculating differences between the utility of an option under consideration and the utilities of its alternatives in the choice situation—idea essentially similar to ECU theory—has been explored in the philosophical literature (Colyvan, 2008; Colyvan & Hájek, 2016)[3] and economic modeling literature (Zhang, 2015; Zhang et al., 2004).

The core of this paper is divided into two sections: In Section 2, I will compare EU theory and ECU theory, revealing how and why they differ, and in Section 3, I will develop a step-by-step

[1] According to Martin Peterson, today, nearly all decision theorists agree that the "principle of maximizing expected value is the appropriate decision rule to apply to decisions under risk […] There are no serious contenders" (Peterson, 2017, p. 66).

[2] A *decision under certainty* is a choice situation where an agent is subjectively certain about which state the world is in and where he or she assigns probability 1 to that state being actual, whereas a *decision under risk* is a choice situation where an agent is *not* subjectively certain about which state the world is in, but where he or she can nevertheless assign probabilities to the different possible states.

[3] Colyvan (2008) has argued for a new decision theory that gives the right verdicts in decision problems where there are an infinite number of states with only finite utilities attached, such as the St-Petersburg game, and where EU theory gives no verdicts whatsoever. Colyvan's new theory, i.e., *relative expectation theory*, states that rational agents rank their choice options on the basis of their *relative expected utility*: for any agent, *S*, and for any two options, *a* and *b*, *S* prefers *a* to *b* if and only if the probability-weighted sum of the differences in utility between *a* and *b* for each possible state is positive, and *S* is indifferent between *a* and *b* if and only if the probability-weighted sum of the differences in utility between *a* and *b* for each possible state is zero. Relative expectation theory gives the same decision advice as EU theory in all decision cases where there are only a finite number of possible states and where the states are probabilistically independent of all choice options. See also Colyvan and Hájek (2016).

argument for ECU theory (and against EU theory).

2 | EU THEORY VS. ECU THEORY

This section will explicate and contrast EU theory and ECU theory.

In what follows, I will assume that for any agent, *S*, and for any choice option, *a*, for *S*, *a*'s *utility* is a *cardinal* indicator of 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 conditions or axioms of *rational preference*, one of which is 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 analyzed as satisfying the IIA. Note however that the IIA has been challenged (Sen, 1993; Wedgwood, 2013, pp. 2668–2670).[4]

According to EU theory, the EU of an option, *a*, in a decision problem with *n* states is formally defined as:

$$\mathrm{EU}(a) = \sum_{i=1}^{n} U(a, s_i)P(s_i)$$

where $U(a, s_i)$ denotes the utility of option *a* when state s_i is actual, and $P(s_i)$ denotes the probability assigned to state s_i. In other words, for any number of alternative options, *a*, *b*, *c*, *d*, and *e*, one calculates the EU of *a* as follows: for each state of the world, one calculates *a*'s utility and one multiplies the result by the probability that one assigns to that state; finally, one sums the totals for every state.

According to EU theory, for any agent, *S*, faced with any decision under *certainty* or any decision under *risk* and for any number of alternative options, *a*, *b*, *c*, *d*, and *e*, for *S*, it is rational for *S* to prefer *a* to *b*, and *a* is more choiceworthy than *b* for *S*, if and only if *a*'s EU is greater than *b*'s; it is rational for *S* to be indifferent between *a* and *b*, and *a* is just as choiceworthy as *b* for *S*, if and only if *a*'s EU is equal to *b*'s; the extent to which *S* rationally prefers *a* to *b*, and the extent to which *a* is more choiceworthy than *b* for *S*, is the difference in EU between *a* and *b*; finally, it is rational for *S* to weakly prefer[5] *a* over the alternative options available to *S*, and *a* is choiceworthy for *S*, if and only if *a* maximizes EU within the set of alternatives available to *S*.

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

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

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

[4] I follow Gustafsson (2021) in taking the independence axiom of expected utility theory to be a requirement of rational preference and in taking the Allais and Ellsberg preferences to be irrational.

[5] For any agent, *S*, and for any two choice options, *a* and *b*, for *S*, if *S weakly prefers a* to *b*, then *S* either prefers *a* to *b* or is indifferent between *a* and *b*.

the highest utility in state s_i), and $P(s_i)$ denotes the probability assigned to state s_i. 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 (within the set of alternatives *b*, *c*, *d*, and *e*), and one multiplies the result by the probability that one assigns to that state; finally, one sums the totals for every state.[6]

More precisely, for any agent, *S*, faced with any decision under *certainty* 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 utility* (*CU*) of *a* in the state of the world to which *S* assigns probability 1. Let us call this principle the *CU principle*. For any choice option, *a*, and for any state of the world, *G*, the choiceworthiness or CU of *a*, in *G*, is the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*.[7] (Henceforth, *c-utiles* are defined as units of CU.) For any agent, *S*, faced with any decision under *risk* 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 comparative utility* (*CECU*) of *a*, that is to say, the difference in ECU between *a* and whichever alternative(s) to *a* carry the greatest ECU. 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 risk). According to ECU theory, for any agent, *S*, faced with any decision under *certainty*, *S* should choose whichever option carries the greatest CU (or one of them in the event that several alternatives are tied) within the set of alternatives available to *S*, and for any agent, *S*, faced with any decision under *risk*, *S* should choose whichever option carries the greatest CECU (or one of them in the event that several alternatives are tied) within the set of alternatives available to *S*.

To demonstrate how to apply EU theory and 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 1500 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 c-utile − [−1 c-utile] = 2 units of CECU). According to EU theory, *S* should also choose option *a*, since its EU (15 utiles) is greater than that of every other option in the decision situation (i.e. the EU of *b* is 14 utiles).

The EUs and ECUs of options *a* and *b* are given by the two equations below. The following notation is used: *A* 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 × 0.98 = 0.9702), *B* 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 × 0.02 = 0.0198), *C* denotes the state "If *S* chooses *a*, then *S* will win the prize (1500 utiles) and if *S* chooses *b*, then *S* will not win the prize (0 utiles)" (probability: 0.01 × 0.98 = 0.0098), *D* denotes the state "If *S* chooses *a*, then *S* will win the prize (1500 utiles) and if *S* chooses *b*, then *S* will win the prize (700 utiles)" (probability: 0.01 × 0.02 = 0.0002), P(*A*) denotes the probability of state *A*, and U(*a*, *A*) denotes the utility of option *a* when state *A* is actual. (See Table 1).

[6] ECU theory only applies to decision problems where the states of the world are probabilistically independent of the agent's choices. The theory therefore fails to give any verdicts in Newcomb decision problems, in which options and states are probabilistically dependent.

[7] *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* (Henderson, 2008).

TABLE 1 Decision matrix

	A (0.9702)	*B* (0.0198)	*C* (0.0098)	*D* (0.0002)
a	0	0	1500	1500
b	0	700	0	700

$$\mathrm{EU}(a) = \mathrm{U}(a, A) \times \mathrm{P}(A) + \mathrm{U}(a, B) \times \mathrm{P}(B) + \mathrm{U}(a, C) \times \mathrm{P}(C) + \mathrm{U}(a, D) \times \mathrm{P}(D) = 15 \text{ utiles}$$
$$\mathrm{EU}(b) = \mathrm{U}(b, A) \times \mathrm{P}(A) + \mathrm{U}(b, B) \times \mathrm{P}(B) + \mathrm{U}(b, C) \times \mathrm{P}(C) + \mathrm{U}(b, D) \times \mathrm{P}(D) = 14 \text{ utiles}$$

$$\mathrm{ECU}(a) = (\mathrm{U}(a, A) - \mathrm{U}(b, A)) \times \mathrm{P}(A) + (\mathrm{U}(a, B) - \mathrm{U}(b, B)) \times \mathrm{P}(B) + (\mathrm{U}(a, C) - \mathrm{U}(b, C)) \times \mathrm{P}(C) + (\mathrm{U}(a, D) - \mathrm{U}(b, D)) \times \mathrm{P}(D) = 1 \text{ c-utile}$$
$$\mathrm{ECU}(b) = (\mathrm{U}(b, A) - \mathrm{U}(a, A)) \times \mathrm{P}(A) + (\mathrm{U}(b, B) - \mathrm{U}(a, B)) \times \mathrm{P}(B) + (\mathrm{U}(b, C) - \mathrm{U}(a, C)) \times \mathrm{P}(C) + (\mathrm{U}(b, D) - \mathrm{U}(a, D)) \times \mathrm{P}(D) = -1 \text{ c-utiles}$$

Let a *finite decision* be a decision problem where there are only finitely many states and no infinite utilities. In all finite decisions under risk requiring a choice between *only two* alternative options, ECU theory delivers the same verdicts as EU theory. However, in some finite decisions under risk requiring a choice between *more than two* alternative options, ECU theory gives different verdicts from those of EU theory.[8]

Let us consider the following two examples:

The Walk:[9] Alice is going for a long walk. She knows that within the next hour, there is a 50% chance of sunny skies (state *A*) and a 50% chance of rain (state *B*). She is faced with a choice between five options: bring a rain poncho and wear rain boots[10] (option *a*), bring an umbrella and wear rain boots (option *b*), bring an umbrella and wear running shoes (option *c*), not bring an umbrella and wear rain boots (option *d*), and not bring an umbrella and wear running shoes (option *e*). Each possible outcome of Alice's choice corresponds to the experience of taking a walk, and the utilities indicate Alice's preferences between those possible outcomes. Which option should Alice choose? Should she lug around a poncho or an umbrella and wear heavy rain boots in case it rains, should she forego the poncho and the umbrella and wear running shoes in case the skies are sunny, or should she go for the middle ground: bring an umbrella, but wear running shoes, or not bring an umbrella, but wear rain boots?

The Restaurant: Fred goes to his local restaurant for dinner. When seated, he sees that there are two menus on his table (one to his left and one to his right). Fred notices that each menu shows the same five items, but that the prices for all the items differ between the two menus. Fred asks his waiter which menu shows the correct prices for the food items. His waiter tells him that "the menu to his left shows the correct prices, whereas the menu to his right does not" (state *A*). However, from past experience, Fred knows that his waiter is reliable at giving correct information 50% of the time. So it is possible that "the menu to Fred's right shows the correct prices, whereas the menu to his left does not" (state *B*). Which food item should Fred order? Should he order the pizza (option *a*), the spaghetti (option *b*), the omelet (option *c*), the sandwich (option *d*), or the salad (option *e*)? Each possible outcome of Fred's choice is a mixture of the experience of eating the food and the price of

---

8 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 (Colyvan, 2008; Colyvan & Hájek, 2016, pp. 838–839).

9 This example is inspired from Briggs (2019).

10 From the outset, Alice rules out bringing a rain poncho and wearing running shoes because her feet will get too wet if it rains.

the food, and the utilities indicate Fred's preferences between those possible outcomes.

The above two decision problems can be stated more formally as follows: 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* were realized, then *S* would assign the following utilities to the set of options (see Table 2):

TABLE 2 Decision matrix

	***A* (0.5)**	***B* (0.5)**	**EU**	**CECU**
a	2	10	6	0
b	5	8	6.5	−0.5
c	6	6	6	−1
d	8	4	6	−1
e	10	2	6	0

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. By contrast, according to ECU theory, *a* is more choiceworthy than *b*, for *S*, since the CECU of *a* is greater than that of *b*. In fact, according to ECU theory, *a* is choiceworthy tout court, since its CECU is equal to or greater than zero. (ECU(*a*) = −3, ECU(*b*) = −3.5, ECU(*c*) = −4, ECU(*d*) = −4, ECU(*e*) = −3)

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*.

Let us consider again the previous decision situation. In that 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 (and according to EU theory), as shown below (see Table 3):

TABLE 3 Decision matrix

	***A* (0.5)**	***B* (0.5)**	**EU**	**CECU**
a	2	10	6	−1
b	5	8	6.5	1

According to ECU theory, *b* is choiceworthy tout court since its CECU is equal to or greater than zero. (ECU(*a*) = −0.5, ECU(*b*) = 0.5)

Here is another example where ECU theory violates the IIA and gives different verdicts from those of EU theory[11]: 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* were realized, then *S* would assign the following utilities to the set of options (see Table 4):

[11] Thanks to an anonymous reviewer for giving this example.

TABLE 4 Decision matrix

	A **(0.001)**	***B*** **(0.999)**	**EU**	**CECU**
a	1000	1	1.999	0.002
b	0	2	1.998	−0.002

According to ECU theory, *a* is choiceworthy tout court, since its CECU is equal to or greater than zero (ECU(*a*) = 0.001, ECU(*b*) = −0.001). And according to EU theory, *a* is also choiceworthy tout court, since *a* maximizes EU. Let us now introduce a third choice option (*c*) in the decision situation, all else being the same (see Table 5):

TABLE 5 Decision matrix

	A **(0.001)**	***B*** **(0.999)**	**EU**	**CECU**
a	1000	1	1.999	−0.898
b	0	2	1.998	0.898
c	900	0	0.9	−2.097

In this new decision situation, *b* is choiceworthy tout court according to ECU theory, since *b*'s CECU is equal to or greater than zero. (ECU(*a*) = −0.899, ECU(*b*) = −0.001, ECU(*c*) = −2.098) By contrast, according to EU theory, *a* is choiceworthy tout court, since *a* maximizes EU. 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. ECU theory thus violates the *Irrelevance of statewise dominated alternatives* (*ISDA*) (Quiggin, 1994).

This gives rise to a worry. Without the IIA (and ISDA), it is possible to make up alternatives in any choice set and these manufactured alternatives would be altering the degrees of choiceworthiness of reasonable options.[12] This opens the door to strategic manipulation in the decision process. The worry can be 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 (Chalmers, 2017; Monton, 2019; Smith, 2014, 2016). As such, the *very improbable* outcomes of manufactured alternatives 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 (see Table 6):

TABLE 6 Decision matrix

	1	**2**	**3**	**4**
a	5	5	5	5
b	1	1	1	1
c	1	2	2	2
d	1	3	3	3
e	1	3	5	8

The difference in CU between *a* and *b* is greater in situation 1 ((5 – 1) – (1 – 5) = 8 c-utiles) than

[12] Thanks to Douglas Lackey for raising this point and for wording suggestions.

in situation 2 ((5 – 3) – (1 – 5) = 6 c-utiles), and is greater in situation 2 than in 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 intervals between *a* and *b* in situations 1 to 4. What's more, there are not any contrary cases where CU (or CECU) gives a *less* differentiated picture than does utility (or EU).

3 | THE ARGUMENT FOR ECU THEORY

This section will argue for a new normative theory of rational choice under risk, namely ECU theory. The argument can be broken down into 15 steps, which are numbered below.

Let us begin with a preliminary argument (i.e., the *instrumental rationality argument*). First, note that, in what follows, I will use the words *choiceworthy* and *choiceworthiness* in a non-moral sense. Therefore, what is *choiceworthy* for *S* should be distinguished from what it is *morally good* or *morally right* for *S* to do, and *choiceworthiness* should be distinguished from *moral goodness* and *moral rightness*.[13]

1. For any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, *a* is *choiceworthy* for *S* if and only if *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world, where *S*'s *rational preferences* are preferences that obey the series of rationality conditions or axioms of standard decision theory. Moreover, the *degree* to which *a* is *choiceworthy* for *S*, or (i.e.) the *choiceworthiness* of *a* for *S*, is the degree to which *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world.
2. *S*'s choosing *a* is *instrumentally rational* if and only if *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences),[14] and the *degree* to which *S*'s choosing *a* is *instrumentally rational* is the degree to which *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
3. *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world if and only if *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences), and the degree to which *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world is the degree to which *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
4. Therefore, *a* is *choiceworthy* for *S* if and only if *S*'s choosing *a* is instrumentally rational, and the *choiceworthiness* of *a* for *S* is the *degree* to which *S*'s choosing *a* is instrumentally rational. (4 follows from 1 to 3.)

The argument for ECU theory proceeds as follows:

1. For any agent, *S*, and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a is maximally choiceworthy* for *S* over the space of all alternatives in the choice set.
2. For any agent, *S*, and for any option, *a*, for *S*, *a is maximally choiceworthy* for *S* over the space of all alternatives in the choice set if and only if *a maximizes choiceworthiness* for *S* over the space of all alternatives in the choice set.
3. For any agent, *S*, and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a maximizes*

[13] *Moral rightness* arguably cannot be measured on a graded scale (Hurka, n.d., 2019; Olsen, 2018; Peterson, 2022; Sinhababu, 2018).

[14] According to Kolodny and Brunero (2020), "Someone displays instrumental rationality insofar as she adopts suitable means to her ends."

choiceworthiness for *S* over the space of all alternatives in the choice set (i.e., the *choiceworthiness maximization* (*CM*) *principle*). (3 follows from 1 and 2.)

4. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* is its CECU, that is, the difference between its ECU and that of whichever alternative(s) to *a* carry the greatest ECU (i.e., the *CECU principle*). *a*'s ECU is a probability-weighted sum of *a*'s CUs across the various states of the world, where, for any state of the world, *G*, *a*'s *CU* in *G* is the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*.
5. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a* maximizes CECU. (5 follows from 3 and 4.)
6. For any agent, *S*, faced with any decision under *risk* and for any number of alternative options, *a*, *b*, *c*, *d*, and *e*, for *S*, it is rational for *S* to prefer *a* to *b* if and only if *a*'s EU is greater than *b*'s, it is rational for *S* to be indifferent between *a* and *b* if and only if *a*'s EU is equal to *b*'s, and the extent to which *S* rationally prefers *a* to *b* is the difference in EU between *a* and *b*.
7. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, it is rational for *S* to weakly prefer *a* over the alternative options in the choice set if and only if *a* maximizes EU. (7 follows from 6.)
8. In decisions under *risk*, what option(s) maximize CECU sometimes differ from what option(s) maximize EU.
9. In decisions under *risk*, what option(s) are choiceworthy sometimes differ from what option(s) it is rational to weakly prefer over the alternative options in the choice set. (9 follows from 5, 7 and 8.)
10. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in terms of how choiceworthy they are for *S* (i.e., how worthy of being chosen by *S* they are in light of *S*'s rational preferences within each of the various possible states of the world).
11. It is *not* the case that for any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in order of rational preference. (11 follows from 3, 9 and 10.)
12. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in terms of how choiceworthy they are for *S*, that is, according to their CECU, rather than in order of rational preference, that is, according to their EU. (12 follows from 4, 6, 10 and 11.)
13. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) are *choiceworthy* for *S* (i.e., what option(s) are worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world).
14. It is *not* the case that for any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set. (14 follows from 9 and 13.)
15. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) are choiceworthy for *S* (i.e., what option(s) maximize CECU), even in cases where what option(s) are choiceworthy for *S* (i.e., what option(s) maximize CECU) differ from what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set (i.e., what option(s) maximize EU). (15 follows from 5, 7, 9, 13 and 14, as well as from 3, 5, 7, 9 and 10.)

I will now discuss the different steps in the argument:

1. For any agent, *S*, and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a is maximally*

choiceworthy for *S* over the space of all alternatives in the choice set.

The question whether a given option is more (or less) choiceworthy than (or just as choiceworthy as) another option within a set of alternatives is well-formed and meaningful. Therefore, the question whether a given option is *maximally choiceworthy* within a set of alternatives is also well-formed and meaningful. A given option is *maximally choiceworthy* within a set of alternatives if and only if it is at least as choiceworthy as each of the other options within the set of alternatives. I will assume that Step 1 is true without further argument.

2. For any agent, *S*, and for any option, *a*, for *S*, *a is maximally choiceworthy* for *S* over the space of all alternatives in the choice set if and only if *a maximizes choiceworthiness* for *S* over the space of all alternatives in the choice set.

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 unchoiceworthy tout court). To ask how choiceworthy an option is is to ask how desirable or worthy of being chosen that option is, how imperative it is to choose that option.[15] 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—i.e., we require a graded, quantitative measure of *choiceworthiness*.

3. For any agent, *S*, and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a maximizes choiceworthiness* for *S* over the space of all alternatives in the choice set (i.e., the CM principle). (3 follows from 1 and 2.)
4. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* is its CECU, that is, the difference between its ECU and that of whichever alternative(s) to *a* carry the greatest ECU (i.e., the *CECU principle*). *a*'s ECU is a probability-weighted sum of *a*'s CUs across the various states of the world, where, for any state of the world, *G*, *a*'s *CU* in *G* is the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*.

3.1 The CU principle

In order to establish the CECU principle, I first need to argue for a graded, quantitative measure of choiceworthiness for decisions under *certainty* (i.e. the *CU principle*). According to the CU principle, 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* is its *comparative utility* (*CU*). For any choice option, *a*, and for any

[15] A qualification is in order: 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.

state of the world, *G*, *a*'s *CU* in *G* is the difference in utility, in *G*, between *a* and whichever alternative(s) to *a* carry the greatest utility in *G*. In what follows, I will provide three arguments for the CU principle.

To that end, I will assume that for any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, *a* is *choiceworthy* for *S* 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 principle 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 unchoiceworthy tout court).

3.1.1 | Argument 1

It is generally accepted 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 equivalently 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 (rationally) preferred alternative to *a* (either *b*, *c*, *d*, or *e*) *more* than *S* (rationally) prefers *b* to the most (rationally) preferred alternative to *b* (either *a*, *c*, *d*, or *e*).[16] After all, if *S* must choose an alternative to *a*, then *S* should choose the most (rationally) preferred alternative to *a* (either *b*, *c*, *d*, or *e*) (or one of them in the event that several alternatives are tied), and not necessarily the option to which *S* is comparing *a* (i.e., option *b*). The same goes for option *b*. Therefore, the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which {the difference in utility between *a* and whichever alternative(s) to *a* carry the greatest utility (i.e., *b*, *c*, *d*, or *e*)} is greater than {the difference in utility between *b* and whichever alternative(s) to *b* carry the greatest utility (i.e., *a*, *c*, *d*, or *e*)}. It follows that the extent to which *a* is choiceworthy for *S* (or [i.e.] the measure of *how* choiceworthy *a* is for *S*) is *a*'s *CU*, that is to say, the difference in utility between *a* and whichever alternative(s) to *a* carry the greatest utility (i.e., *b*, *c*, *d*, or *e*). The same goes for option *b*. This is what I have referred to as the *CU principle*.

An alternative approach is to say that the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which *S* (rationally) prefers *a* to the most (rationally) preferred option (or options) (i.e., *a*, *b*, *c*, *d*, or *e*) *more* than *S* (rationally) prefers *b* to the most (rationally) preferred option (or options) (i.e., *a*, *b*, *c*, *d*, or *e*). In other words, the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which {the difference in utility between *a* and whichever option(s) carry the greatest utility (i.e., *a*, *b*, *c*, *d*, or *e*)} is greater than {the difference in utility between *b* and whichever option(s) carry the greatest utility (i.e., *a*, *b*, *c*, *d*, or *e*)}. If that is the case, then the extent to which *a* is choiceworthy for *S* (or [i.e.] the measure of *how* choiceworthy *a* is for *S*) is *a*'s *CU**, i.e., the difference in utility between *a* and whichever option(s) carry the greatest utility (i.e., *a*, *b*, *c*, *d*, or *e*). The same goes for option *b*. I will refer to this as the *CU* principle*.[17]

[16] More precisely, the extent to which *a* is more choiceworthy than *b*, for *S*, is the extent to which *S* (rationally) prefers *a* to the most (rationally) preferred alternative (or alternatives) to *a* (i.e., *b*, *c*, *d*, or *e*) *more* than *S* (rationally) prefers *b* to the most (rationally) preferred alternative (or alternatives) to *b* (i.e., *a*, *c*, *d*, or *e*).

[17] For any choice option, *a*, and for any state of the world, *G*, *a*'s *CU** in *G* is the difference in utility, in G, between *a* and whichever option(s) carry the greatest utility in *G* (i.e., *a*, *b*, *c*, *d*, or *e*). The rule of maximizing *expected CU** (or *ECU**) counsels agents to choose whichever option in the choice set has the greatest ECU* (or one of them in the event that several alternatives are tied), where ECU* is a probability-weighted sum of an option's CUs* across the various states of the world. The rule of maximizing ECU* is equivalent to the rule of maximizing EU (i.e., EU theory), which means

The CU* principle is however untenable, since it results in a double standard. It entails that the degrees of choiceworthiness of all the option(s) that do *not* carry the greatest utility depend on what other options are available in the choice set—those degrees of choiceworthiness may be different negative numbers, but never 0—whereas the degrees of choiceworthiness of all the option(s) that *do* carry the greatest utility do *not* depend on what other options are available in the choice set—those degrees of choiceworthiness are 0 no matter what the utilities of the other options are. Moreover, the latter standard is implausible. It's as if the degrees of choiceworthiness of all the option(s) that do *not* carry the greatest utility did *not* depend on what other options are available in the choice set—it's as if those degrees of choiceworthiness were the same negative number, e.g., −1, no matter what the utilities of the other options are. Contrary to the CU* principle, the original CU principle does not suffer from these problems.

Let us now consider four choice situations involving decisions under certainty (see Table 7): Compared to the *difference in utility* and the *difference in CU**, the *difference in CU* is a more plausible measure of the extent to which *a* is more choiceworthy than *b* in situations 1–4, as explained above. The differences in utility and CU* between *a* and *b* are the same in all four situations (4 units), whereas the differences in CU between *a* and *b* are as follows (in situations 1–4):

TABLE 7 Decision matrix[a]

	1	**2**	**3**	**4**
a	5	5	5	5
b	1	1	1	1
c	1	2	2	2
d	1	3	3	3
e	1	3	5	8

[a] Table 7 is identical to Table 6.

1. (5 – 1) – (1 – 5) = 8 c-utiles
2. (5 – 3) – (1 – 5) = 6 c-utiles
3. (5 – 5) – (1 – 5) = 4 c-utiles
4. (5 – 8) – (1 – 8) = 4 c-utiles

The CU principle is therefore well-supported.

3.1.2 | Argument 2

1. For any agent, *S*, faced with any decision under certainty and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a* is worthy of being chosen by *S* over whichever alternative(s) to *a* are the most choiceworthy for *S*. (True by definition)
2. *a* is choiceworthy for *S* if and only if *a* maximizes choiceworthiness for *S* over the space of all alternatives in the choice set (i.e., the CM principle). (Assumption)
3. The extent to which *a* is choiceworthy for *S* (or [i.e.] the measure of *how* choiceworthy *a* is for *S*) is the extent to which *a* is worthy of being chosen by *S* over whichever alternative(s) to *a* are the most choiceworthy for *S*. (3 follows from 1 and 2.)
4. *a* is choiceworthy for *S* if and only if *a* maximizes utility over the space of all alternatives in the

that both rules deliver the same verdicts in all decision cases.

choice set (i.e., the UM principle). (Assumption)

5. *a* maximizes choiceworthiness for *S* over the space of all alternatives in the choice set if and only if *a* maximizes utility over the space of all alternatives in the choice set. (5 follows from 2 and 4.)
6. *a* maximizes choiceworthiness for *S* over the space of all alternatives in a subset of the choice set if and only if *a* maximizes utility over the space of all alternatives in that subset of the choice set. (6 follows from 5.)
7. Whichever alternative(s) to *a* are the most choiceworthy for *S* are whichever alternative(s) to *a* carry the greatest utility. (7 follows from 6.)
8. The extent to which *a* is worthy of being chosen by *S* over some alternative to *a* is the difference in utility between *a* and that alternative to *a*. (True by conceptual analysis)
9. 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(s) to *a* carry the greatest utility (i.e., the *CU principle*). (9 follows from 3, 7 and 8.)

3.1.3 | Argument 3

Let us now consider a longer argument. The simplest attempt at defining a graded, 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*, the measure of the choiceworthiness of *a* for *S* is 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. 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).[18] 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 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 (a) any available option is choiceworthy just in case its utility value is equal to or greater than zero (and not

[18] "[S]ince the utilities of options, whether ordinal or interval-valued, can only be determined relative to the utilities of other options, there is no such thing as the absolute utility of an option, at least not without further assumptions. The further assumptions would need to relate particular options to particular privileged levels of utility; for instance, one would need to argue that a rational agent's preference ordering should incorporate, say, a privileged zero-utility option, in which case ratios of utility distances from this option would be meaningful." (Steele & Stefánsson, 2020) "The zero point and the unit in an expected utility representation are arbitrary; utility values become meaningful only once they have been fixed." (Colyvan & Hájek, 2016, pp. 838–839)

choiceworthy otherwise) and (b) 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 unchoiceworthy, 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 accordance with the measurement principle, it is necessarily the case that for any agent, *S*, faced with *any* 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 *benchmark*) in the following way: the measure of the choiceworthiness of *a* for *S* (relative to *any* explicitly given utility unit and zero point of utility) is the *difference in utility* between *a* and some *benchmark* for *a*, such that (i) *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 (ii) 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 whichever option has the highest utility, whichever option has the lowest utility, or the status quo, or some average of the utilities of the available options. Choiceworthiness is thus a relative concept.[19] 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 (or one of them in the event that several alternatives are tied). 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(s) are choiceworthy for *S* (or, per the UM principle, whichever alternative(s) to *a* carry the greatest utility). 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 (or one of them in the event that several alternatives are tied). As far as I know, this idea has not been explored by others in the published literature.

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 (but not the same state for both situations). If that state of the world were realized, then *S* would assign the following utilities to the set of options (see Table 8):

[19] Ralph Wedgwood (2017) relies on considerations of incommensurability to argue for the same idea: "the choiceworthiness of options is relative to choice situations". Temkin (2012) also addresses this idea: what he calls the "Essentially Comparative View."

TABLE 8 Decision matrix

	1	2
a	100	100
b	−10,000	99
c	−10,000	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 *S* is in situation 1 than if *S* 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., *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 us now briefly introduce Ralph Wedgwood's *benchmark theory* (*BT*) (Wedgwood, 2013). 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*" (p. 2664) should be excluded from consideration at the outset.[20] 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 (see Table 9):

TABLE 9 Decision matrix

	1	2
a	100	100
b	−100	−500
c	−100	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—that is to say, it is more imperative for *S* to choose *a* if *S* is in situation 1 than if *S* 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* 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

[20] For critiques of BT, see Bassett (2015) and Briggs (2010).

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 a counterexample to BT*.

Another very similar example (see Table 10):

TABLE 10 Decision matrix

	1	**2**
a	100	100
b	−100	−500
c	−100	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 *S* is in situation 1 than if *S* 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 (see Table 11):

TABLE 11 Decision matrix

	1	**2**
a	100	100
b	−100	99
c	100	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 *S* is in situation 1 as it is if *S* 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 (or one of them in the event that several alternatives are tied).

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 (or one of them in the event that several alternatives are tied). 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*. 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(s) to *a* carry the greatest utility. As previously indicated, I will refer to this principle as the *CU principle*. Like the utility principle, the CU principle entails the UM principle.

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 (a) 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 (b) *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 were realized, then *S* would assign the following utilities to the available options: *a* (0), *b* (−100), *c* (−1000). 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 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 1000 utiles. Per the CU principle, the degrees of choiceworthiness of the available options are as follows: *a* (100), *b* (−100), *c* (−1000). 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 1100 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.

The utility principle, let us recall, states that 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* is the utility of *a* in the state of the world to which *S* assigns probability 1. Since the utility principle is false and since the expected utility of *a* equals the utility of *a* in the state of the world to which *S* assigns probability 1, it follows that for any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, the measure of the choiceworthiness of *a* for *S* *cannot* be the expected utility of *a*.

3.2 | The CECU principle

As I argued in discussing Step 2, we require a graded, quantitative measure of how choiceworthy options are. When we move from decision-making under certainty to decision-making under risk, we can, in light of the CU principle, 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 sum of all possible values, of that option's choiceworthiness, or CU, *in the actual state of the world*. That roughly encapsulates ECU theory.

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 illustrated in Section 2, 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 risk* 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) (i.e., the CM principle)—ECU theory, as defined above, is false in decision cases *under risk*. By the same lines of reasoning as employed in Section 3.1.3, we reach the following conclusion: ECU theory is the conjunction of the CU principle (for decisions under certainty) and the CECU principle (for decisions under risk). Let us recall that according to the CECU principle, for any agent, *S*, faced with any decision under *risk* 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 CECU of *a*, that is to say, the difference in ECU between *a* and whichever alternative(s) to *a* carry the greatest ECU.

One line of argument in support of ECU theory is that, contrary to EU theory, ECU theory entails Wedgwood's *Gandalf's principle*: the choiceworthiness of an option *in a given state of the world* should be measured only relative to the values of the other options *in that state*, and not to the values of the options *in other states*. According to Wedgwood (2013, p. 2654),

> to make a rational choice in [cases involving risk], 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.

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.[21] In a paper critiquing Wedgwood's BT, Robert Bassett (2015) concurs: "Gandalf's principle strikes me as an eminently sensible principle to incorporate into rational decision-making." There is, however, one alternative decision theory which 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, and in cases where there is more than one maximally likely value of *a*'s choiceworthiness (or CU) in the actual state of the word, the measure of the choiceworthiness of *a* for *S* (relative to any explicitly given utility unit and zero point of utility) is *a*'s CECU across the maximally likely states of the world. We require a further argument to rule out MLCU 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* were realized, then *S* would assign the following utilities to the set of options (see Table 12):

[21] For an illuminating discussion of Gandalf's principle, see Wedgwood (2013), pp. 2652–2655.

TABLE 12 Decision matrix

	A **(0.51)**	***B*** **(0.49)**
a	110	−1000
b	80	110
c	100	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 (−1110 c-utiles). I take this to be a counterexample to MLCU theory.

A second line of argument in support of ECU theory is that for the same reasons as those given in Section 3.1.1 (except that we consider here rational preferences within various possible states of the world instead of rational preferences within a decision situation under certainty), compared to the difference in EU, the difference in CECU is a more plausible measure of the extent to which option *a* is more choiceworthy than option *b* in the following decision matrices (Tables 13–16). The differences in EU between *a* and *b* are the same in all four decision matrices (4 units), whereas the differences in CECU between *a* and *b* are as follows:

TABLE 13 Decision matrix

	A **(0.5)**	***B*** **(0.5)**	**EU**	**CECU**
a	5	5	5	8
b	1	1	1	−8
c	1	1	1	−8
d	1	1	1	−8
e	1	1	1	−8

Note: The difference in CECU between *a* and *b* = 16 units.

TABLE 14 Decision matrix

	A **(0.5)**	***B*** **(0.5)**	**EU**	**CECU**
a	5	5	5	4
b	1	1	1	−6
c	2	2	2	−5
d	3	3	3	−4
e	3	3	3	−4

Note: The difference in CECU between *a* and *b* = 10 units.

TABLE 15 Decision matrix

	***A* (0.5)**	***B* (0.5)**	**EU**	**CECU**
a	5	5	5	0
b	1	1	1	−4
c	2	2	2	−3
d	3	3	3	−2
e	5	5	5	0

Note: The difference in CECU between *a* and *b* = 4 units.

TABLE 16 Decision matrix

	***A* (0.5)**	***B* (0.5)**	**EU**	**CECU**
a	5	5	5	−6
b	1	1	1	−10
c	2	2	2	−9
d	3	3	3	−8
e	8	8	8	6

Note: The difference in CECU between *a* and *b* = 4 units.

5. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, *a* is choiceworthy for *S* if and only if *a* maximizes CECU. (5 follows from 3 and 4.)
6. For any agent, *S*, faced with any decision under *risk* and for any number of alternative options, *a*, *b*, *c*, *d*, and *e*, for *S*, it is rational for *S* to prefer *a* to *b* if and only if *a*'s EU is greater than *b*'s, it is rational for *S* to be indifferent between *a* and *b* if and only if *a*'s EU is equal to *b*'s, and the extent to which *S* rationally prefers *a* to *b* is the difference in EU between *a* and *b*.

Decision-theoretic representation theorems—such as those of von Neumann and Morgenstern (1947), Savage (1954), Bolker (1966) and Jeffrey (1983), and Joyce (1999)—show that if an agent fails to prefer choice options with higher EU, then that agent violates at least one of a series of axioms of *rational preference*,[22] one of which is the IIA. A further justification for Step 6 comes from money-pump arguments for the axioms of EU theory (see Gustafsson, 2022).

7. For any agent, *S*, faced with any decision under *risk* and for any option, *a*, for *S*, it is rational for *S* to weakly prefer *a* over the alternative options in the choice set if and only if *a* maximizes EU. (7 follows from 6.)
8. In decisions under *risk*, what option(s) maximize CECU sometimes differ from what option(s) maximize EU.

   For examples, see Section 2.

9. In decisions under *risk*, what option(s) are choiceworthy sometimes differ from what option(s) it is rational to weakly prefer over the alternative options in the choice set. (9 follows from 5, 7 and 8.)

Rational preference is not a reliable measure of choiceworthiness. That is because whereas the

[22] Following the formulation of Briggs (2019).

criterion of rational preference (i.e., EU) satisfies the IIA (as assumed in this paper), the criterion of choiceworthiness (i.e., CECU) violates that principle (see Section 2 for examples). 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.

10. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in terms of *how choiceworthy* they are for *S* (i.e., how worthy of being chosen by *S* they are in light of *S*'s rational preferences within each of the various possible states of the world).

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. (See the discussion of Step 2.)

Let us now consider again the *instrumental rationality argument*:

1. For any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, the *degree* to which *a* is *choiceworthy* for *S*, or (i.e.) the *choiceworthiness* of *a* for *S*, is the degree to which *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world, where *S*'s *rational preferences* are preferences that obey the series of rationality conditions or axioms of standard decision theory.
2. The *degree* to which *S*'s choosing *a* is *instrumentally rational* is the degree to which *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
3. The degree to which *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world is the degree to which *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
4. Therefore, the *choiceworthiness* of *a* for *S* is the *degree* to which *S*'s choosing *a* is instrumentally rational. (4 follows from 1 to 3.)

11. It is *not* the case that for any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in order of rational preference. (11 follows from 3, 9 and 10.)

For any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, the *degree* to which *S*'s choosing *a* is instrumentally rational is the degree to which *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world. (See the *instrumental rationality argument*.) Therefore, for any agent, *S*, faced with any decision under *risk*, it is a *theory* of instrumental rationality that *S* should measure and rank his or her options in order of rational preference. However, it is *not* true *by definition* that for any agent, *S*, faced with any decision under *risk*, it is an actual *requirement* of instrumental rationality that *S* should measure and rank his or her options in order of rational preference.

12. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should measure and rank his or her options in terms of how choiceworthy they are for *S*, that is, according to their CECU, rather than in order of rational preference, that is, according to their EU. (12 follows from 4, 6, 10 and 11.)

I can offer a further argument for Step 12. Compared to the criterion of rational preference (i.e., EU), the criterion of choiceworthiness (i.e., CU/CECU) supplies a more plausible measure of the intervals in rankings of more than two choice options. For decisions under certainty, see Section 3.1.1. For decisions under risk, consider the intervals between options *a* and *b* in decision matrices 13–16.

13. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) are *choiceworthy* for *S* (i.e., what option(s) are worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world).

Let us consider again the *instrumental rationality argument*:

1. For any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, *a* is *choiceworthy* for *S* if and only if *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world, where *S*'s *rational preferences* are preferences that obey the series of rationality conditions or axioms of standard decision theory.
2. *S*'s choosing *a* is *instrumentally rational* if and only if *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
3. *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world if and only if *S*'s choosing *a* is a suitable means to *S*'s ends (i.e., *S*'s rational preferences).
4. Therefore, *a* is *choiceworthy* for *S* if and only if *S*'s choosing *a* is instrumentally rational. (4 follows from 1 to 3.)

14. It is *not* the case that for any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set. (14 follows from 9 and 13.)

For any agent, *S*, faced with any decision under certainty or any decision under risk and for any option, *a*, for *S*, *S*'s choosing *a* is instrumentally rational if and only if *a* is worthy of being chosen by *S* in light of *S*'s rational preferences within each of the various possible states of the world. (See the *instrumental rationality argument*.) Therefore, for any agent, *S*, faced with any decision under *risk*, it is a *theory* of instrumental rationality that *S* should choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set. However, it is *not* true *by definition* that for any agent, *S*, faced with any decision under *risk*, it is an actual *requirement* of instrumental rationality that *S* should choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set.

15. For any agent, *S*, faced with any decision under *risk*, it is a requirement of instrumental rationality that *S* should choose out of what option(s) are choiceworthy for *S* (i.e., what option(s) maximize CECU), even in cases where what option(s) are choiceworthy for *S* (i.e., what option(s) maximize CECU) differ from what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set (i.e., what option(s) maximize EU). (15 follows from 5, 7, 9, 13 and 14, as well as from 3, 5, 7, 9 and 10.)

4 | CONCLUSION

In recent years, several alternatives to EU theory have been proposed, for example Mark Colyvan's

(2008) *relative expectation theory* (*RET*), Paul Bartha's (2007, 2016) *relative utility theory* (*RUT*), and Lara Buchak's (2013) *risk-weighted expected utility* (*REU*) *theory*. In all finite decision cases, RET and RUT deliver the same rankings and recommendations as EU theory. As for REU theory, it *can* deliver the same rankings and recommendations 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 are false theories of *instrumental rationality*.

ACKNOWLEDGMENTS

Thanks to Douglas Lackey, Kenny Easwaran, John Quiggin and several anonymous reviewers for very helpful comments on earlier drafts of this paper. I am also grateful for the helpful comments I received on Reddit (https://www.reddit.com/r/AcademicPhilosophy/comments/p49mxz/a_restatement_of_expected_comparative_utility/ , https://www.reddit.com/r/philosophy/comments/p4vajs/a_restatement_of_expected_comparative_utility/ , https://www.reddit.com/r/academiceconomics/comments/tnnfd5/a_restatement_of_expected_comparative_utility/ , https://www.reddit.com/r/philosophy/comments/ttrhhl/a_restatement_of_expected_comparative_utility/).

REFERENCES

Bartha, P. (2007). Taking stock of infinite value: Pascal's wager and relative utilities. *Synthese*, *154*, 5–52. https://doi.org/10.1007/s11229-005-8006-z

Bartha, P. (2016). Making do without expectations. *Mind*, *125*, 799–827. https://doi.org/10.1093/mind/fzv152

Bassett, R. (2015). A critique of benchmark theory. *Synthese*, *192*, 241–267. https://doi.org/10.1007/s11229-014-0566-3

Bernoulli, D. (1738 [1954]). Exposition of a new theory on the measurement of risk. *Econometrica*, *22*, 23–36. https://doi.org/10.2307/1909829

Bolker, E. (1966). Functions resembling quotients of measures. *Transactions of the American Mathematical Society*, *2*, 292–312. https://doi.org/10.2307/1994401

Briggs, R. (2010). Decision-theoretic paradoxes as voting paradoxes. *Philosophical Review*, *119*, 1–30. https://doi.org/10.1215/00318108-2009-024

Briggs, R. A. (2019). Normative theories of rational choice: Expected utility. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. https://plato.stanford.edu/archives/fall2019/entries/rationality-normative-utility/

Buchak, L. (2013). *Risk and rationality*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199672165.001.0001

Chalmers, A. (2017). *An offer you can't (rationally) refuse: Systematically exploiting utility-maximisers with malicious gambles* (Unpublished honours thesis). The University of Sydney. https://hdl.handle.net/2123/16131

Colyvan, M. (2008). Relative expectation theory. *Journal of Philosophy*, *105*, 37–44. https://doi.org/10.5840/jphil200810519

Colyvan, M., & Hájek, A. (2016). Making ado without expectations. *Mind*, *125*, 829–857. https://doi.org/10.1093/mind/fzv160

Gustafsson, J. E. (2022). *Money-Pump Arguments*. Cambridge University Press.

Gustafsson, J. E. (2021). The sequential dominance argument for the independence axiom of expected utility theory. *Philosophy and Phenomenological Research*, *103*, 21–39. https://dx.doi.org/10.1111/phpr.12669

Henderson, D. R. (2008). Opportunity cost. In *The concise encyclopedia of economics*. Library of Economics and Liberty. https://www.econlib.org/library/Enc/OpportunityCost.html

Hurka, T. (n.d.). *More seriously wrong.* Unpublished manuscript. https://www.brown.edu/academics/philosophy/ethical-inquiry/sites/brown.edu.academics.philosophy.ethical-inquiry/files/uploads/seriously%20wrong%202016.pdf

Hurka, T. (2019). More seriously wrong, more importantly right. *Journal of the American Philosophical Association*, *5*, 41–58. https://doi.org/10.1017/apa.2018.41

Jeffrey, R. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.

Joyce, J. M. (1999). *The foundations of causal decision theory*. Cambridge University Press.

Kolodny, N., & Brunero, J. (2020). Instrumental rationality. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. https://plato.stanford.edu/archives/spr2020/entries/rationality-instrumental/

Monton, B. (2019). How to avoid maximizing expected utility. *Philosophers' Imprint*, *19*, 1–25. https://hdl.handle.net/2027/spo.3521354.0019.018

Olsen, K. (2018). Subjective rightness and minimizing expected objective wrongness. *Pacific Philosophical Quarterly*, *99*, 417–441. https://doi.org/10.1111/papq.12188

Peterson, M. (2017). *An introduction to decision theory*. Cambridge University Press. https://doi.org/10.1017/9781316585061

Peterson, M. (2022). Moral rightness comes in degrees. *Journal of the American Philosophical Association*, 1–20. https://doi.org/10.1017/apa.2021.36

Quiggin, J. (1994). Regret theory with general choice sets. *Journal of Risk and Uncertainty*, *8*, 153–165. https://doi.org/10.1007/BF01065370

Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–184). Routledge and Kegan Paul.

Savage, L. (1954). *The foundations of statistics*. Wiley.

Sen, A. (1993). Internal consistency of choice. *Econometrica*, *61*, 495–521. https://doi.org/10.2307/2951715

Sinhababu, N. (2018). Scalar consequentialism the right way. *Philosophical Studies*, *175*, 3131–3144. https://doi.org/10.1007/s11098-017-0998-y

Smith, N. J. J. (2014). Is evaluative compositionality a requirement of rationality? *Mind*, *123*, 457–502. https://doi.org/10.1093/mind/fzu072

Smith, N. J. J. (2016). Infinite decisions and rationally negligible probabilities. *Mind*, *125*, 1199–1212. https://doi.org/10.1093/mind/fzv209

Steele, K., & Stefánsson, H. O. (2020). Decision theory. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. https://plato.stanford.edu/archives/win2020/entries/decision-theory/

Temkin, L. S. (2012). *Rethinking the good: Moral ideals and the nature of practical reasoning*. Oxford University Press.

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton University Press.

Wedgwood, R. (2013). Gandalf's solution to the newcomb problem. *Synthese*, *190*, 2643–2675. https://doi.org/10.1007/s11229-011-9900-1

Wedgwood, R. (2017). Must rational intentions maximize utility? *Philosophical Explorations*, *20*, 1–20. https://doi.org/10.1080/13869795.2017.1356352

Zhang, J. (2015). Relative utility modelling. In S. Rasouli & H. Timmermans (Eds.), *Bounded rational choice behaviour: Applications in transport* (pp. 49–72). Emerald Group Publishing Limited. https://doi.org/10.1108/978-1-78441-072-820151014

Zhang, J., Timmermans, H., Borgers, A., & Wang, D. (2004). Modeling traveler choice behavior using the concepts of relative utility and relative interest. *Transportation Research Part B*, *38*, 215–234. https://doi.org/10.1016/S0191-2615(03)00009-2