

Expected comparative utility theory: A new theory of instrumental rationality

David Robert

January 25, 2024

Copyright © 2024 The Author.

Correspondence

Email: jeandavidrobert@hotmail.com

Expected comparative utility theory: A new theory of instrumental rationality

Abstract

This paper aims to address the question of how one ought to choose when one is uncertain about what outcomes will result from one's choices, but when one can nevertheless assign probabilities to the different possible outcomes. These choices are commonly referred to as *choices* (or *decisions*) *under risk*. I assume in this paper that one ought to make *instrumentally rational* choices—more precisely, one ought to adopt suitable means to one's *morally permissible* ends. *Expected utility (EU) theory* is generally accepted as a normative theory of rational choice under risk, or, more specifically, as a theory of instrumental rationality. According to EU theory, when faced with a decision under risk, one ought to rank one's options (from least to most choiceworthy) according to their EU and one ought 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. In this paper, I argue that EU theory is not the correct theory of instrumental rationality. In its place, I argue for a new theory of instrumental rationality, namely *expected comparative utility (ECU) theory*.

I first 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 of a in G —that is, the difference in utility, in G , between a and whichever alternative to a carries the greatest utility in G . On the basis of this principle, I then argue, roughly speaking, that for any agent, S , faced with any decision under risk, S ought to rank his or her options (in terms of how choiceworthy they are) according to their ECU and S ought to choose whichever option carries the greatest ECU (or one of them in the event that several alternatives are tied). For any option, a , a 's ECU is a probability-weighted sum of a 's comparative utilities across the various possible states of the world. In this paper, I show that in some commonplace decisions under risk, ECU theory delivers different verdicts from those of EU theory.

1. Introduction

This paper is a restatement and expansion of my previously published work: Robert (2018) and Robert (2021). I have revised and restructured the work to present its main points more clearly and convincingly and in greater detail, and to develop new arguments for ECU theory. Sections 3.2.3, 4, and 5 of this paper are completely new.

This paper aims to address the question of how one *ought* to choose when one is *uncertain* about what outcomes will result from one's choices, but when one can nevertheless assign probabilities to the different possible outcomes. These choices are commonly referred to as *choices* (or *decisions*) *under risk*. Along the way, this paper will also address the question of how one *ought* to choose when one is *certain* about what outcomes will result from one's choices and when one assigns probability 1 to those outcomes. These choices are commonly referred to as *choices* (or *decisions*) *under certainty*.

Standard decision theory, otherwise known as *expected utility (EU) theory*, requires that when faced with a decision under risk (or a decision under certainty), one ought to rank one's choice options (from least to most choiceworthy) according to their EU and one ought 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 (from the 1920s onwards), has received foundational support from both economists and philosophers (Bolker, 1966; Jeffrey, 1983; Joyce, 1999; Ramsey, 1931; Savage, 1954; von Neumann

& Morgenstern, 1947).¹

I will assume in this paper that for any agent, *S*, and for any choice option, *a*, for *S*, *a*'s utility is a cardinal index of preference (which measures the strength of preference between options, e.g. *a*) 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)* (see Gintis, 2018). According to the 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*. In this paper, *rational preferences* are understood as preferences that obey the axioms of rational preference of standard decision theory.

I will further assume in this paper that one ought to make *instrumentally rational* choices. Someone makes instrumentally rational choices, according to Kolodny and Brunero (2020), "insofar as she adopts suitable means to her ends." Orthodox EU theory is a theory of instrumental rationality (Buchak, 2022), where the agent's *ends* (or rational preferences) are understood as morally permissible or impermissible ends. As such, orthodox EU theory is not a *normative* theory of choice (even though it is meant to be), since it is not the case that an agent ought to choose in accordance with EU theory if some or all of the agent's ends (or rational preferences) are morally impermissible. I rely here on a unified view of practical reason, which has been recently defended by Brown (2023). Brown (2023) argues that there is such a thing as what one ought *simpliciter* to do, and proposes "an account of our normative concepts according to which only ought *simpliciter* judgments commit one to acting in accordance with those judgments."

Consequently, in this paper, I will depart from orthodoxy by stipulating that to make *instrumentally rational* choices—according to EU theory and any normative alternative to EU theory—is to adopt suitable means to one's *morally permissible* ends, or (i.e.) one's morally permissible *rational preferences*, provided that all of one's ends (or rational preferences) are morally permissible.² If not all of one's ends are morally permissible, then I assume that one ought to adopt morally permissible means to one's ends, which means that if the optimal means to one's ends are morally impermissible, then one ought to adopt morally permissible, suboptimal means to one's ends. Thus, for the sake of brevity, in this paper, *ends* will be understood as morally permissible ends, or (i.e.) morally permissible rational preferences, and *preferences (to prefer a to b or to be indifferent between a and b)* will be understood as morally permissible preferences.

In this paper, I will argue for a new normative alternative to EU theory. Arguments against EU theory typically involve identifying decision situations where EU theory gives the intuitively wrong recommendations. For example, the Allais, Ellsberg, St Petersburg paradoxes do just that (see Allais, 1953; Ellsberg, 1961; Bernoulli, 1738). This paper takes a different approach. Instead of identifying counterexamples to EU theory that strongly suggest a new normative alternative, I will develop conceptual arguments in favor of a new normative alternative, alternative which I will show to be inconsistent with EU theory.

Starting from the premise that EU is the appropriate criterion of *rational preference* to apply to decisions under certainty and decisions under risk, and from the premise that we require a graded, quantitative measure of *choiceworthiness* for decisions under certainty and decisions under risk, I will argue that we need a new normative theory of rational choice under risk, or, more specifically, a new theory of instrumental rationality, namely *expected comparative utility (ECU) theory*. 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 is therefore not the correct normative theory of rational choice under

¹ See Briggs, 2023, for an overview of EU theory. See Buchak, 2022, for an overview of normative rivals to EU theory.

² I follow here Horton (2023), according to whom "an option is permissible if and only if it could be rationally chosen by someone with permissible preferences."

risk, or, more specifically, not the correct theory of instrumental rationality.

ECU Theory aims to improve upon regret-based decision theories.³ These theories are founded on the idea that rational agents strive to minimize 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). The basic regret-based decision theory for decisions under risk makes use of a *benchmark* (or zero point of choiceworthiness). According to this theory, 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, or (i.e.) the probability-weighted sum of each of an option's possible degrees of choiceworthiness, always coincides with maximizing EU.

ECU Theory also makes use of a benchmark (one that is importantly different from the one employed in the basic regret-based decision theory): 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* of a . Roughly speaking, ECU Theory requires that when faced with a decision under risk, one ought to rank one's options (in terms of how choiceworthy they are) according to their ECU and one ought to choose whichever option carries the greatest ECU (or one of them in the event that several alternatives are tied). For any choice option, a , the ECU of a is the probability-weighted sum of a 's comparative utilities across the various states of the world.

The idea of calculating differences between the utility of an option under consideration and the utilities of its alternatives in the choice situation—idea that accords with ECU theory—has been explored in the philosophical literature (Bartha, 2007 & 2016; Colyvan, 2008; Colyvan & Hájek, 2016; Yager, 2017) and the economic modeling literature (Zhang, 2015).

In the following section (Section 2), I will compare EU theory and ECU theory, revealing how and why they differ. I will show that in some commonplace decisions under risk, ECU Theory gives different verdicts from those of EU Theory. In Section 3, I will develop a comprehensive step-by-step argument for ECU theory (and against EU theory). In Section 4, I will respond to some possible objections to ECU theory. In Section 5, I will conclude.

2. EU Theory vs. ECU Theory

This section will explicate and compare EU theory and ECU theory. I will begin by explaining how EU theory and ECU theory are to be formulated and how they work. I will then explain how and why their verdicts differ in various decision problems.

2.1 What Are EU Theory and ECU Theory? How Do They Work?

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

³ See Yager, 2017. For an early normative decision theory that takes into account anticipated regret, namely, *regret theory*, see Bell, 1982; Fishburn, 1982; and Loomes & Sugden, 1982.

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

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 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 .⁴

According to what we can call *EU theory's choice situation principle*, for any agent, S , faced with any decision under *certainty* or any decision under *risk* between any number of alternative choice situations, a, b, c, d , and e , it is rational for S to prefer a to b , and a is more choiceworthy than b for S , if and only if the EU of whatever option(s) in a carry the greatest EU is greater than the EU of whatever option(s) in b carry the greatest EU; it is rational for S to be indifferent between a to b , and a is just as choiceworthy as b for S , if and only if the EU of whatever option(s) in a carry the greatest EU is equal to the EU of whatever option(s) in b carry the greatest EU; 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 the EU of whatever option(s) in a carry the greatest EU and the EU of whatever option(s) in b carry the greatest EU; finally, it is rational for S to weakly prefer a over the alternative choice situations, b, c, d , and e , and a is choiceworthy for S , if and only if the EU of whatever option(s) in a carry the greatest EU is equal to or greater than the EU of whatever option(s) in b, c, d , and e carry the greatest EU.

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:

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

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 the greatest 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 from a 's utility the utility of whichever alternative(s) to a (i.e.,

⁴ 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 .

b, *c*, *d*, or *e*) carry the greatest 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.

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*.⁵ (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.

In Section 3, I will argue for the following principle: 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 will call this principle the *choiceworthiness maximization* (*CM*) *principle*. According to the conjunction of the CU principle and the CM principle (henceforth, according to the CU principle), 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*'s CU is equal to or greater than zero, and according to the conjunction of the CECU principle and the CM principle (henceforth, according to the CECU principle), 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*'s CECU is equal to or greater than zero.

We are now in a position to precisely define ECU theory: ECU theory is the conjunction of the CU principle (for decisions under certainty), the CECU principle (for decisions under risk) and what I will refer to as *ECU theory's choice situation principle*.

According to ECU theory's choice situation principle, for any agent, *S*, faced with any decision under *certainty* or any decision under *risk* between any number of alternative choice situations, *a*, *b*, *c*, *d*, and *e*, (i) *a* is more choiceworthy than *b* for *S* if and only if the EU of whatever option(s) in *a* carry the greatest EU is greater than the EU of whatever option(s) in *b* carry the greatest EU, or if the EU of whatever option(s) in *a* carry the greatest EU is equal to the EU of whatever option(s) in *b* carry the greatest EU, then the CU/CECU of whatever option(s) in *a* carry the greatest CU/CECU is greater than the CU/CECU of whatever option(s) in *b* carry the greatest CU/CECU; (ii) *a* is just as choiceworthy as *b* for *S* if and only if the EU of whatever option(s) in *a* carry the greatest EU is equal to the EU of whatever option(s) in *b* carry the greatest EU, and the CU/CECU of whatever option(s) in *a* carry the greatest CU/CECU is equal to the CU/CECU of whatever option(s) in *b* carry the greatest CU/CECU; (iii) the extent to which *a* is more choiceworthy than *b* for *S* is the difference in EU between the EU of whatever option(s) in *a* carry the greatest EU and the EU of whatever option(s) in *b* carry the greatest EU, or if the EU of whatever option(s) in *a* carry the greatest EU is equal to the EU of whatever option(s) in *b* carry the greatest EU, then the extent to which *a* is more choiceworthy than *b* for *S* is the difference in CU/CECU between the CU/CECU of whatever option(s) in *a* carry the greatest CU/CECU and the CU/CECU of whatever option(s) in *b* carry the greatest CU/CECU; (iv) *a* is choiceworthy for *S* if and only if *a* is more choiceworthy than or just as choiceworthy as each

⁵ *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).

of the alternative choice situations, *b*, *c*, *d*, and *e*.

To demonstrate how to apply EU theory and ECU theory to a concrete decision problem involving *objective probabilities*, 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* ought to 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* ought to 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 four 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: $0.99 \times 0.98 = 0.9702$), *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: $0.99 \times 0.02 = 0.0198$), *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: $0.01 \times 0.98 = 0.0098$), *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: $0.01 \times 0.02 = 0.0002$), $P(A)$ denotes the probability of state *A*, and $U(a, A)$ denotes the utility of option *a* when state *A* is actual. (See Table 1).

TABLE 1 Decision matrix

	A (0.9702)	B (0.0198)	C (0.0098)	D (0.0002)
<i>a</i>	0	0	1500	1500
<i>b</i>	0	700	0	700

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

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

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

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

2.2 Do EU Theory’s and ECU Theory’s Verdicts Differ? How and Why?

Mark Colyvan (2008) has argued for a new normative decision theory, namely, *relative expectation theory*, that gives the right verdicts in *some* decision problems where there are an infinite number of states with only finite utilities attached (such as the St. Petersburg game), decision problems where EU theory gives no verdicts whatsoever (see Colyvan and Hájek, 2016, pp. 838–839; Meacham, 2019, pp. 1010–1012). According to Colyvan’s new theory, for any agent, *S*, and for any two alternative options, *a* and *b*, *S* rationally prefers *a* to *b* (and *a* is more choiceworthy than *b* for *S*) 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 rationally indifferent between *a* and *b* (and *a* is just as choiceworthy as *b* for *S*) 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

problems where there are only a finite number of possible states and where the states are probabilistically independent of all choice options.

Like relative expectation theory, ECU theory requires that one calculate the probability-weighted sum of the differences in utility between the options for each possible state. Therefore, in its handling of decision problems where there are infinitely many states with only finite utilities attached (such as the St. Petersburg game), ECU theory inherits the advantages of relative expectation theory over EU theory. That is to say, in *some* of those *infinite decision* problems, ECU theory delivers the intuitively correct verdicts, whereas EU theory delivers none.

Now, let a *finite decision* problem 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. Let us consider the following example:⁶ 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 (option *a*), bring an umbrella and wear rain boots (option *b*), bring an umbrella and wear running shoes (option *c*), not bring an umbrella, nor a rain poncho, and wear rain boots (option *d*), and not bring an umbrella, nor a rain poncho, and wear running shoes (option *e*). (From the outset, Alice rules out bringing a rain poncho and wearing running shoes because for some reason she believes that she cannot jointly do so.) 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, nor a rain poncho, but wear rain boots?

The above decision problem 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

	<i>A</i> (0.5)	<i>B</i> (0.5)	EU	CECU
<i>a</i>	2	10	6	0
<i>b</i>	4	8	6	-1
<i>c</i>	6	6	6	-1
<i>d</i>	8	4	6	-1
<i>e</i>	10	2	6	0

According to EU theory, all the options (i.e., *a*, *b*, *c*, *d*, and *e*) are choiceworthy since their EU is the same. By contrast, according to ECU theory, only options *a* and *e* are choiceworthy, since their CECU is equal to or greater than zero. ($ECU(a) = -3$, $ECU(b) = -4$, $ECU(c) = -4$, $ECU(d) = -4$, $ECU(e) = -3$) With respect to Alice's walk, ECU theory therefore recommends that Alice either bring a rain poncho and wear rain boots, or not bring an umbrella, nor a rain poncho, and wear running shoes. This verdict may seem counterintuitive, as ECU theory rules out as unchoiceworthy the middle ground (i.e., bring an umbrella, but wear running shoes, or not bring an umbrella, nor a rain poncho,

⁶ This example is inspired from Briggs' (2023) umbrella example.

but wear rain boots). However, I hope to show that this verdict is not counterintuitive in light of the conceptual arguments for ECU theory that I will develop in Section 3.

Here is a slightly revised decision matrix (see Table 3):

TABLE 3 Decision matrix

	A (0.5)	B (0.5)	EU	CECU
<i>a</i>	2	10	6	0
<i>b</i>	5	8	6.5	-0.5
<i>c</i>	6	6	6	-1
<i>d</i>	8	4	6	-1
<i>e</i>	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* (and also *e*) is choiceworthy tout court, since its CECU is equal to or greater than zero. ($ECU(a) = -3$, $ECU(b) = -3.5$, $ECU(c) = -4$, $ECU(d) = -4$, $ECU(e) = -3$) Therefore, with respect to Alice’s walk, EU theory recommends that Alice bring an umbrella and wear rain boots, whereas ECU theory recommends that Alice either bring a rain poncho and wear rain boots, or not bring an umbrella, nor a rain poncho, and wear running shoes.

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 decision situation illustrated in Table 3. In that decision 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 4):

TABLE 4 Decision matrix

	A (0.5)	B (0.5)	EU	CECU
<i>a</i>	2	10	6	-1
<i>b</i>	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 theory⁷: 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 5):

⁷ Thanks to an anonymous reviewer for giving this example.

TABLE 5 Decision matrix

	A (0.001)	B (0.999)	EU	CECU
<i>a</i>	1000	1	1.999	0.002
<i>b</i>	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 6):

TABLE 6 Decision matrix

	A (0.001)	B (0.999)	EU	CECU
<i>a</i>	1000	1	1.999	-0.898
<i>b</i>	0	2	1.998	0.898
<i>c</i>	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. Thus, ECU theory, contrary to EU theory, violates the *irrelevance of statewise dominated alternatives* (Quiggin, 1994).

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 7):

TABLE 7 Decision matrix

	1	2	3	4
<i>a</i>	5	5	5	5
<i>b</i>	1	1	1	1
<i>c</i>	1	2	2	2
<i>d</i>	1	3	3	3
<i>e</i>	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 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. In other words, compared to utility, CU is a more discriminating measure of the extent to which *a* is more choiceworthy than *b* in situations 1 to 4. And there are not any contrary cases where CU (or CECU) gives a *less* differentiated picture than does utility (or EU).

ECU theory delivers different verdicts from those of EU theory not only with respect to *choice options*, but also with respect to *choice situations*, both in decisions under certainty and decisions under risk. Let us consider the following four examples:

According to ECU theory, for any agent, S , faced with a decision (under certainty) between choice situations 1 and 2 (see Table 7), choice situation 1 is more choiceworthy for S than choice situation 2 since the utility of whatever option(s) in 1 carry the greatest utility is equal to the utility of whatever option(s) in 2 carry the greatest utility, and the CU of whatever option(s) in 1 carry the greatest CU is greater than the CU of whatever option(s) in 2 carry the greatest CU. By contrast, according to EU theory, for any agent, S , faced with a decision (under certainty) between choice situations 1 and 2, choice situation 1 is just as choiceworthy for S as choice situation 2 since the utility of whatever option(s) in 1 carry the greatest utility is equal to the utility of whatever option(s) in 2 carry the greatest utility.

According to both ECU theory and EU theory, for any agent, S , faced with a decision (under certainty) between choice situations 1 to 4 (see Table 7), choice situation 4 is uniquely choiceworthy for S since the utility of whatever option(s) in 4 carry the greatest utility is greater than the utility of whatever option(s) carry the greatest utility in each of the alternative choice situations (1 to 3).

According to ECU theory, for any agent, S , faced with a decision (under risk) between choice situations 13 and 14 (see Tables 13 and 14 in Section 3.2.2), choice situation 13 is more choiceworthy for S than choice situation 14 since the EU of whatever option(s) in 13 carry the greatest EU is equal to the EU of whatever option(s) in 14 carry the greatest EU, and the CECU of whatever option(s) in 13 carry the greatest CECU is greater than the CECU of whatever option(s) in 14 carry the greatest CECU. By contrast, according to EU theory, for any agent, S , faced with a decision (under risk) between choice situations 13 and 14, choice situation 13 is just as choiceworthy for S as choice situation 14 since the EU of whatever option(s) in 13 carry the greatest EU is equal to the EU of whatever option(s) in 14 carry the greatest EU.

According to both ECU theory and EU theory, for any agent, S , faced with a decision (under risk) between choice situations 13 to 16 (see Tables 13 to 16 in Section 3.2.2), choice situation 16 is uniquely choiceworthy for S since the EU of whatever option(s) in 16 carry the greatest EU is greater than the EU of whatever option(s) carry the greatest EU in each of the alternative choice situations (13 to 15).

3. The Argument for ECU Theory

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

First, note that, in what follows, *what it is to be choiceworthy* and *choiceworthiness* will be given the following conceptual analyses: 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 axioms of standard decision theory, and 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 axioms of standard decision theory.

More explicitly, 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 decision matrix, that is, in light of what utilities S assigns to each option within each of the various possible states of the world and what probabilities S assigns to each of the various possible state of the world, and 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 decision matrix, that is, in light of what utilities S assigns to each option within each of the various possible states of the world and what probabilities S assigns to each of the various possible state of the world.

Following Horton (2023), in this paper, I stipulate that all of S 's preferences are *rational* and *morally permissible*. If some or all of S 's preferences within each of the various possible states of the world were *morally impermissible*, then it would follow that S ought to choose out of what options are morally permissible, even if those options are suboptimal in light of S 's (morally permissible and morally impermissible) rational preferences within each of the various possible states of the world.

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 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 (i.e., the CECU principle).
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*, S ought to measure and rank his or her options (for the purpose of choice) in terms of *how choiceworthy* they are for S .
11. It is *not* the case that for any agent, S , faced with any decision under *risk*, S ought to measure and rank his or her options (for the purpose of choice) in order of rational preference. (11 follows from 3, 9 and 10.)
12. For any agent, S , faced with any decision under *risk*, S ought to measure and rank his or her options (for the purpose of choice) 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, as well as from 3, 4, 6, 9 and 13.)

13. For any agent, S , faced with any decision under *risk*, S ought to choose out of what option(s) are *choiceworthy* for S .
14. It is *not* the case that for any agent, S , faced with any decision under *risk*, S ought to 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*, S ought to 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 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. 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 (i.e., the CECU principle).

3.1 Step 4: 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 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

According to EU theory, 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(s) to *a* (either *b*, *c*, *d*, or *e*) more than *S* (rationally) prefers *b* to the most (rationally) preferred alternative(s) to *b* (either *a*, *c*, *d*, or *e*). It follows that the extent to which *a* is choiceworthy for *S* tout court is the extent to which *S* (rationally) prefers *a* to the most (rationally) preferred alternative(s) to *a* (either *b*, *c*, *d*, or *e*). In other words, 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.

The rationale behind the CU principle is simple: If *S* must forgo option *a*, then *S* ought to choose (any of) the most (rationally) preferred alternative(s) to *a* (either *b*, *c*, *d*, or *e*), and not necessarily the option to which *S* is comparing *a* (i.e., option *b*). If *S* *does* choose option *a*, then *S* necessarily forgoes the most (rationally) preferred alternative(s) to *a* (either *b*, *c*, *d*, or *e*). It follows that the extent to which *a* is choiceworthy for *S* tout court is the extent to which *a* is worthy of being chosen over the most (rationally) preferred alternative(s) to *a* (either *b*, *c*, *d*, or *e*). In other words, the extent to which *a* is choiceworthy for *S* is the extent to which *S* (rationally) prefers *a* to the most (rationally) preferred alternative(s) to *a* (either *b*, *c*, *d*, or *e*)—which is *a*'s CU. To express this point differently, in determining the extent to which *a* is choiceworthy for *S* tout court, *S* ought to take into account the utility gain or cost for *S* of choosing (any of) the most (rationally) preferred alternative(s) to *a* over *a*. The greater the utility *cost* to *S* of choosing (any of) the most (rationally) preferred alternative(s) to *a* over *a*, the more choiceworthy *a* is to *S*, and the greater the utility *gain* for *S* of choosing (any of) the most (rationally) preferred alternative(s) to *a* over *a*, the less choiceworthy *a* is to *S*. This observation

agrees with the CU principle.

An alternative approach is to say 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, the difference in utility between *a* and whichever *option(s)* carry the greatest utility (i.e., *a, b, c, d, or e*). I will refer to this as the *CU* principle*.⁸ The *CU** principle is however untenable, not only for the reasons just stated, but also because it results in a double standard: On the one hand, the degrees of choiceworthiness of all the option(s) that do *not* carry the greatest utility *do* depend (beyond the fact of their *not* carrying the greatest utility) on what other options are available in the choice set—those degrees of choiceworthiness may be different negative numbers, but never 0. On the other hand, the degrees of choiceworthiness of all the option(s) that *do* carry the greatest utility *do not* depend (beyond the fact of their carrying the greatest utility) 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 (beyond the fact of their *not* carrying the greatest utility) 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 8): 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 8 Decision matrix^a

	1	2	3	4
<i>a</i>	5	5	5	5
<i>b</i>	1	1	1	1
<i>c</i>	1	2	2	2
<i>d</i>	1	3	3	3
<i>e</i>	1	3	5	8

^aTable 8 is identical to Table 7.

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.

⁸ 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*). According to the rule of maximizing *expected CU** (or *ECU**), one ought 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 that both rules deliver the same verdicts in all decision cases. See the discussion of the notion of "expected shortfall from optimality" in Wedgwood, 2003.

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). (Step 3 of the argument for ECU theory)
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 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 consider a final argument for a graded, quantitative choiceworthiness measure for decisions under certainty. *Measures* of quantities that have an interval scale, 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 accordance with the measurement principle, 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. 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).

Let us consider two decision situations: 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 9):

TABLE 9 Decision matrix

	1	2
<i>a</i>	100	100
<i>b</i>	-100	99
<i>c</i>	-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. 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 200 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.

Here is a different example (see Table 10):

TABLE 10 Decision matrix

	1	2
<i>a</i>	100	100
<i>b</i>	-100	-100
<i>c</i>	-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. 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.

Let me now introduce Ralph Wedgwood’s (2013) *benchmark theory (BT)*.⁹ 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. Wedgwood explicitly rejects

⁹ For critiques of BT, see Bassett, 2015, and Briggs, 2010.

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.

Let us consider the following example (see Table 11):

TABLE 11 Decision matrix

	1	2
<i>a</i>	100	100
<i>b</i>	-100	-500
<i>c</i>	-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. 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 counterexample to *BT**.

These three examples serve to illustrate that if there are *not* 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 defined in relation 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.¹⁰ As previously indicated, I will refer to this principle as the *CU* principle. The *CU* principle entails the *UM* principle.

3.2 Step 4: The CECU Principle

As I argued in discussing Step 2 of the argument for *ECU* theory, we require a graded, quantitative measure of how choiceworthy options are. It is sometimes assumed that if expected utility theory is the correct normative theory of choice under *certainty* and choice under *risk*, then for any agent, *S*,

¹⁰ Choiceworthiness is thus a relative concept. 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”.

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 is a 's *expected utility*. For example, according to Briggs (2023), "expected utility theory provides a way of ranking the acts according to how *choiceworthy* they are: the higher the expected utility, the better it is to choose the act. (It is therefore best to choose the act with the highest expected utility—or one of them, in the event that several acts are tied.)" Similarly, according to Steele & Stefánsson (2020), "preference between options is a judgment of comparative desirability or choice-worthiness" and von Neumann and Morgenstern's (1947) account of expected utility presumes that "lotteries are evaluated in terms of their expected choice-worthiness or desirability. [...] That is, the desirability of [a] lottery is a probability weighted sum of the utilities of its prizes, where the weight on each prize is determined by the probability that the lottery results in that prize."

Let us suppose 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 is a 's *expected utility*. Since a 's expected utility is a 's *utility* in the state of the world to which S assigns probability 1, it then follows 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 a 's *utility* in the state of the world to which S assigns probability 1. I will refer to the latter as the *utility principle*. The UM principle is true if (*but not only* if) the utility principle is true. However, in light of the CU principle, the utility principle can be falsified. It therefore 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 a 's *expected utility*.

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.

In Section 3.1.1, I argued 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 the most (rationally) preferred alternative(s) to a (either b , c , d , or e) *more* than S (rationally) prefers b to the most (rationally) preferred alternative(s) to b (either a , c , d , or e). It could be proposed that the same holds

for any agent, S , faced with any decision under *risk*. Therefore, in decisions under *risk*, the extent to which a is more choiceworthy than b , for S , would be the extent to which {the difference in *expected utility* between a and whichever alternative(s) to a carry the greatest *expected utility* (i.e., b, c, d , or e)} is greater than {the difference in *expected utility* between b and whichever alternative(s) to b carry the greatest *expected utility* (i.e., a, c, d , or e)}. It would follow that 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 *expected utility* between a and whichever alternative(s) to a carry the greatest *expected utility* (i.e., b, c, d , or e). The same would be the case for option b . However appealing this theory might seem, it runs afoul of Ralph Wedgwood's (2013) *Gandalf's principle* (see Section 3.2.1).¹¹

When we move from decision-making under certainty to decision-making under risk, we can, in light of the CU principle, plausibly 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*. Indeed, according to Martin Peterson (2017, p. 66), "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."

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 2.2, 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 situations *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) (as well as ECU theory's choice situation principle). 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.

In what follows, I will develop three lines of argument in support of the CECU principle:

3.2.1 Argument 1

One line of argument in support of the CECU principle is that, contrary to EU theory, the CECU principle entails Wedgwood's *Gandalf's principle*: the choiceworthiness of an option *within 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 determine how one should choose when one is certain about what the actual state of nature is, one simply *does not need to know* whether one is in a nice state of nature or a nasty state of nature. [...] I suggest the following way of generalizing Gandalf's point to cases involving uncertainty [(i.e., risk)]: to make a rational choice in such cases, one *does not need to consider* whether one is in a nice state of nature or a nasty one. All that one

¹¹ The idea of calculating the difference in expected utility between some option and whichever alternative option(s) carry the greatest expected utility is discussed in Zhu, 2018.

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 (see Wedgwood, 2013, pp. 2652–2655). 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 world, 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 problem: 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):

TABLE 12 Decision matrix

	A (0.51)	B (0.49)
<i>a</i>	110	−1000
<i>b</i>	80	110
<i>c</i>	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 *disutility* of option *a* in state *B* is very high (−1110 c-utiles). I take this to be a counterexample to MLCU theory.

3.2.2 Argument 2

I can offer a second line of argument in support of the CECU principle: 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.

3.2.3 Argument 3

One standard argument for maximizing *expected utility*—a *long-run argument*—is based on the assumption that what an agent ought to care about maximizing in the long run is *utility*. Feller (1968) offers a version of this argument. In this section, I will argue for maximizing *expected comparative utility* (and, by implication, CECU) under the premise that what an agent ought to care about maximizing in the long run is not utility, but *choiceworthiness*. My argument comprises nine steps, which are numbered as follows:

(1) Let a sequence of choices (or trials) be *independent* if and only if (i) the range of options available in each choice situation is independent from the range of options available in each of the other choice situations in the sequence, and (ii) the probability and utility of the outcomes of each choice are independent from the probability and utility of the outcomes of each of the other choices in the sequence. What's more, let a sequence of choices be *identically distributed* if and only if the outcomes of each choice have the same probability distribution. For any agent, S , and for any long sequence of independent and identically distributed (IID) choices under risk, φ , for S , φ is maximally choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to come the closest to cumulatively maximizing the quantity, q , such that for any agent, S , faced with a decision under certainty and for any option, a , for S , the measure of the choiceworthiness of a for S is q , and φ is *not* maximally choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to *not* come the closest to cumulatively maximizing the quantity, q . (See Step 2 of the argument for ECU theory and the discussion thereof.) (2) For any agent, S , and for any choice or sequence of choices, φ , for S , φ is choiceworthy for S if and only if φ is maximally choiceworthy for S —that is to say, φ is at least as choiceworthy as the most choiceworthy alternative to φ .

(3) Given (1) and (2), for any agent, S , and for any long sequence of IID choices under risk, φ , for S , φ is choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to come the closest to cumulatively maximizing the quantity, q , such that for any agent, S , faced with a decision under certainty and for any option, a , for S , the measure of the choiceworthiness of a for S is q , and φ is *not* choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to *not* come the closest to cumulatively maximizing the quantity, q . That is to say, over any long sequence of IID trials, an agent ought to always choose, if she can, according to a decision rule, r , such that if she always chooses according to r over that sequence of trials, she will almost certainly accumulate a greater amount of whatever quantity, q , she ought to care about maximizing in the long run (or lose a smaller amount of q) than if she always chooses according to a different rule which delivers different verdicts for some or all of those trials. And the quantity, q , that an agent ought to care about maximizing in the long run is choiceworthiness (and not utility).

(4) 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*—i.e., the difference in utility between any option, a , and whichever alternative to a carries the greatest utility (or one of them in the event that several alternatives are tied). (Note that the concept of comparative utility applies only to individual choices. Thus, the (cumulative) comparative utility of a sequence of choices should be understood not as the difference between the utility of that sequence and the utility of the best alternative sequence of choices, but rather as the sum of the comparative utilities of each individual choice in that sequence.)

(5) Given (3) and (4), for any agent, S , and for any long sequence of IID choices under risk, φ , for S , φ is choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to come the closest to cumulatively maximizing comparative utility, and φ is *not* choiceworthy for S if φ is the sequence of choices which, by S 's lights, is almost certain to *not* come the closest to cumulatively maximizing comparative utility, where the *comparative utility of a choice* is the difference in utility

between whichever option is chosen in a given decision situation and whichever alternative would carry the greatest utility if it were chosen in that situation. In other words, over any long sequence of IID trials, an agent ought to always choose, if she can, according to a decision rule, r , such that if she always chooses according to r over that sequence of trials, she will almost certainly accumulate a greater amount of comparative utility (or lose a smaller amount of comparative utility) than if she always chooses according to a different rule which delivers different verdicts for some or all of those trials.

(6) Let a *random variable* be a rule or function that assigns a value to each possible outcome of a random trial or experiment. Moreover, let the *expected value* of a random variable (or decision option) be a probability-weighted average of each of its possible values. The *strong and weak laws of large numbers* state that as the number of IID random variables in a sequence approaches infinity, their sample average converges with overwhelming probability to their expected value. Now, it is straightforward to come up with a rule or function such that the values of a random variable can be expressed as comparative utilities (see variable x below). For any choice option, a , the *expected comparative utility* of a is the expected value of a 's comparative utility, or (i.e.) a *probability-weighted average* of a 's comparative utilities across the various states of the world. It follows from the strong and weak laws of large numbers that for any agent, S , and for any number of alternative options, a, b, c, d , and e , the expected comparative utility of a , for S , is overwhelmingly likely to be close to the *long-run average value* of a 's comparative utility. The long-run average value of a 's comparative utility is given by the following notation:

$$\lim_{n \rightarrow \infty} \frac{1}{n} (x_1 + \dots + x_n)$$

where n is the ordinal number of a trial in a sequence of IID trials where five agents ($S_1 \dots S_5$) who are in all relevant respects identical to S respectively and simultaneously perform options a, b, c, d , and e , and where x (the random variable) is the difference between the utility of a in a given trial and the utility of b, c, d , or e in that trial, whichever has the greatest utility in that trial.

What this means is that over any long sequence of IID trials (or choices under risk), an agent will almost certainly accumulate a greater amount of comparative utility (or lose a smaller amount of comparative utility) if she always chooses whichever option has the greatest expected comparative utility (or one of them in case of a tie) than if she always chooses according to a different rule which delivers different verdicts for some or all of those trials.

(7) Given (5) and (6), for any agent, S , and for any long sequence of IID choices under risk, φ , for S , φ is choiceworthy for S if and only if φ is such that each choice in the sequence maximizes expected comparative utility over the space of all alternatives. In other words, over any long sequence of IID trials, an agent ought to always choose whichever option has the greatest expected comparative utility (or one of them in case of a tie).

(8) In long sequences of IID trials, the rule of maximizing expected utility always delivers the same verdicts (for all those trials) as the rule of maximizing expected comparative utility.¹² Let us consider the following case: An agent, S , is faced with a choice between three independent options or gambles: one option, a , offering a 0.01 probability of winning a prize worth 1500 utiles (and nothing otherwise), one option, b , offering a 0.02 probability of winning a prize worth 700 utiles (and nothing otherwise), and one option, c , offering a 0.03 probability of winning a prize worth 400 utiles (and nothing otherwise). According to the rule of maximizing *expected utility*, S ought to choose option a , since its expected utility is greater than that of every other option in the decision situation. According to the rule of maximizing *expected comparative utility* (or CECU), S ought to also choose

¹² As far as I can tell, this holds for all cases. I leave to others the task of proving this conjecture.

option *a*, since its expected comparative utility (or CECU) is greater than that of every other option in the decision situation. See Table 17.

TABLE 17 Decision matrix

	A (0.9410)	B (0.0095)	C (0.0001)	D (0.0000)
<i>a</i>	0	1500	1500	1500
<i>b</i>	0	0	700	700
<i>c</i>	0	0	0	400

	E (0.0192)	F (0.0005)	G (0.0291)	H (0.0002)
<i>a</i>	0	0	0	1500
<i>b</i>	700	700	0	0
<i>c</i>	0	400	400	400

	EU	CECU
<i>a</i>	14.7	1.8
<i>b</i>	13.86	-1.8
<i>c</i>	11.92	-5.69

Note: The following notation is used: *A* denotes the state “If *S* chooses *a*, then *S* will not win the prize (0 utiles), if *S* chooses *b*, then *S* will not win the prize (0 utiles), and if *S* chooses *c*, then *S* will not win the prize (0 utiles)” (probability: $0.99 \times 0.98 \times 0.97 = 0.9410$), *B* denotes the state “If *S* chooses *a*, then *S* will win the prize (1500 utiles), if *S* chooses *b*, then *S* will not win the prize (0 utiles), and if *S* chooses *c*, then *S* will not win the prize (0 utiles)” (probability: $0.01 \times 0.98 \times 0.97 = 0.0095$), *C* denotes the state “If *S* chooses *a*, then *S* will win the prize (1500 utiles), if *S* chooses *b*, then *S* will win the prize (700 utiles), and if *S* chooses *c*, then *S* will not win the prize (0 utiles)” (probability: $0.01 \times 0.02 \times 0.97 = 0.0001$), *D* denotes the state “If *S* chooses *a*, then *S* will win the prize (1500 utiles), if *S* chooses *b*, then *S* will win the prize (700 utiles), and if *S* chooses *c*, then *S* will win the prize (400 utiles)” (probability: $0.01 \times 0.02 \times 0.03 = 0.0000$), *E* denotes the state “If *S* chooses *a*, then *S* will not win the prize (0 utiles), if *S* chooses *b*, then *S* will win the prize (700 utiles), and if *S* chooses *c*, then *S* will not win the prize (0 utiles)” (probability: $0.99 \times 0.02 \times 0.97 = 0.0192$), *F* denotes the state “If *S* chooses *a*, then *S* will not win the prize (0 utiles), if *S* chooses *b*, then *S* will win the prize (700 utiles), and if *S* chooses *c*, then *S* will win the prize (400 utiles)” (probability: $0.99 \times 0.02 \times 0.03 = 0.0005$), *G* denotes the state “If *S* chooses *a*, then *S* will not win the prize (0 utiles), if *S* chooses *b*, then *S* will not win the prize (0 utiles), and if *S* chooses *c*, then *S* will win the prize (400 utiles)” (probability: $0.99 \times 0.98 \times 0.03 = 0.0291$), *H* denotes the state “If *S* chooses *a*, then *S* will win the prize (1500 utiles), if *S* chooses *b*, then *S* will not win the prize (0 utiles), and if *S* chooses *c*, then *S* will win the prize (400 utiles)” (probability: $0.01 \times 0.98 \times 0.03 = 0.0002$),

(9) Given (7) and (8), maximizing *expected comparative utility* (or CECU) makes for an optimal decision rule when an agent is faced with any long sequence of IID trials (or choices under risk). Since in long sequences of IID choices, the rule of maximizing expected comparative utility always delivers the same verdicts (for all those choices) as the rule of maximizing expected utility, and since agents typically face long sequences of choices of this sort in games of chance, this vindicates the standard way of gambling when the probabilities involved are known—that is, by maximizing *expected utility*.

3.3 The Argument for ECU Theory: Steps 5 to 15

In this section, I will discuss the remaining steps (5 to 15) of the argument for ECU theory:

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, 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 , if S fails to prefer a to b when a 's EU is greater than b 's, or if S fails to be indifferent between a and b when a 's EU is equal to b 's, then S violates at least one of a series of axioms of *rational preference*,¹³ one of which is the IIA. Besides the intuitive plausibility of the axioms of rational preference, further justifications for Step 6 come from money-pump arguments for those axioms (see Gustafsson, 2022) as well as arguments for those axioms from consequentialist foundations (see Hammond, 1987, 1988). These justifications for Step 6 represent challenges for several recent normative alternatives to Step 6—for example, Buchak's (2013) *risk-weighted expected utility theory* and Bradley & Stefansson's (2017) *counterfactual desirability theory*.

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 2.2.

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 criterion of rational preference (i.e., EU) satisfies the IIA, the criterion of choiceworthiness (i.e., CU/CECU) violates that principle (see Section 2.2 for examples). In other words, whereas the criterion of rational preference (i.e., EU) only takes into account whichever two choice options are being compared within the set of alternatives, the criterion of choiceworthiness (i.e., CU/CECU) takes into account the entire set of alternatives when evaluating each option, as should any plausible criterion of choice. This gives us a reason to opt for the proposed criterion of choice (i.e., choiceworthiness) over the standard choice criterion (i.e., rational preference). At this point, 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 conditions, 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*, S ought to measure and rank his or her options (for the purpose of choice) in terms of *how choiceworthy* they are for S .

¹³ Paraphrasing Briggs (2023).

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 my discussion of Step 2 of the argument for ECU theory.)

As noted upfront, in this paper, *choiceworthiness*, or (i.e.) *how choiceworthy* an option is, is given the following conceptual analysis: 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 axioms of standard decision theory. Let us now consider the *instrumental rationality argument* (v1):

1. If {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* is the degree to which *a* is a suitable means to *S*'s ends}, then {the *degree* to which *a* is *choiceworthy* 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 axioms of standard decision theory}. (With respect to the consequent, the idea is that the *degree* to which *a* is *choiceworthy* for *S* is the degree to which *a* is worthy of being chosen by *S* *not necessarily* in light of *S*'s rational preferences *tout court* (i.e., *S*'s EU function), but rather in light of *S*'s rational preferences *within each of the various possible states of the world* (i.e., *S*'s alternative possible utility functions). This idea follows from Gandalf's principle, which is discussed in Section 3.2.1.)
2. The *degree* to which *a* is *choiceworthy* for *S* is the degree to which *a* is a suitable means to *S*'s ends. (Assumption)
3. The *degree* to which *a* is *choiceworthy* 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. (3 follows from 1 and 2.)
4. The *degree* to which *a* is a *suitable means* to *S*'s ends 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. (4 follows from 2 and 3.)
5. The *degree* to which *a* is a *suitable means* to *S*'s ends is the degree to which *a* is instrumentally rational for *S*.
6. The degree to which *a* is *instrumentally rational* for *S* is the degree to which *a* is choiceworthy for *S*, or (i.e.) 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. (6 follows from 3 to 5.)

For any agent, *S*, faced with any decision under certainty or any decision under risk, I assume that *S* ought to measure and rank his or her options (for the purpose of choice) in terms of the *degrees* to which they are *instrumentally rational* for *S*, or (i.e.) the degrees to which they are suitable means to *S*'s ends. Therefore, in accordance with the instrumental rationality argument (v1), *S* ought to measure and rank his or her options (for the purpose of choice) in terms of the *degrees* to which they are *choiceworthy* for *S*, or (i.e.) the degrees to which they are worthy of being chosen by *S* 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*, *S* ought to measure and rank his or her options (for the purpose of choice) in order of rational preference. (11 follows from 3, 9 and 10.)

In accordance with my discussion of Step 10, for any agent, *S*, faced with any decision under *risk*, *S* ought to measure and rank his or her options (for the purpose of choice) in terms of the degrees to which they are worthy of being chosen by *S* in light of *S*'s rational preferences *within each of the various possible states of the world* (i.e., *S*'s alternative possible utility functions). It follows that it is a *theoretical possibility* that for any agent, *S*, faced with any decision under *risk*, *S* ought to measure and rank his or her options (for the purpose of choice) in order of rational preference, or (i.e.) in order of EU. However, it does *not* follow that it is true *by definition* that for any agent, *S*, faced with any decision under *risk*, *S* ought to measure and rank his or her options (for the purpose of choice) in order of rational preference, or (i.e.) in order of EU.

12. For any agent, *S*, faced with any decision under *risk*, *S* ought to measure and rank his or her options (for the purpose of choice) 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, as well as from 3, 4, 6, 9 and 13.)

I can offer two further arguments for Step 12: First, there is Wedgwood's Gandalf's principle (see Section 3.2.1). Second, 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 extent to which any given option is more choiceworthy than any other in any ranking of more than two choice options, both in decision situations involving certainty and decision situations involving risk. For decisions under certainty, see Section 3.1.1. For decisions under risk, consider the extent to which option *a* is more choiceworthy than option *b* in decision matrices 13–16 (see Section 3.2.2).

13. For any agent, *S*, faced with any decision under *risk*, *S* ought to choose out of what option(s) are *choiceworthy* for *S*.

Step 13 is a tautology. Therefore, it is necessarily true. Here is another way of arguing for Step 13: As noted upfront, in this paper, *what it is to be choiceworthy* is given the following conceptual analysis: 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 axioms of standard decision theory. Let us now consider the *instrumental rationality argument* (v2):

1. If {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 a suitable means to *S*'s ends}, then {*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 axioms of standard decision theory}. (With respect to the consequent, the idea is that *a* is *choiceworthy* for *S* if and only if *a* is worthy of being chosen by *S* *not necessarily* in light of *S*'s rational preferences *tout court* (i.e., *S*'s EU function), but rather in light of *S*'s rational preferences *within each of the various possible states of the world* (i.e., *S*'s alternative possible utility functions). This idea follows from Gandalf's principle, which is discussed in Section 3.2.1, and the CM principle.)
2. *a* is *choiceworthy* for *S* if and only if *a* is a suitable means to *S*'s ends. (Assumption)
3. *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. (3 follows from 1 and 2.)
4. *a* is a *suitable means* to *S*'s ends 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. (4 follows from 2 and 3.)

5. *a* is a *suitable means* to *S*'s ends if and only if *a* is instrumentally rational for *S*.
6. *a* is *instrumentally rational* for *S* if and only if *a* is choiceworthy for *S*, or (i.e.) *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. (6 follows from 3 to 5.)

For any agent, *S*, faced with any decision under certainty or any decision under risk, I assume that *S* ought to choose out of what option(s) are *instrumentally rational* for *S*, or (i.e.) what option(s) are suitable means to *S*'s ends. Therefore, in accordance with the instrumental rationality argument (v2), *S* ought to choose out of what option(s) are *choiceworthy* for *S*, or (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*, *S* ought to 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.)

In accordance with my discussion of Step 13, for any agent, *S*, faced with any decision under *risk*, *S* ought to choose out of 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* (i.e., *S*'s alternative possible utility functions). It follows that it is a *theoretical possibility* that for any agent, *S*, faced with any decision under *risk*, *S* ought to choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set, or (i.e.) out of what option(s) maximize EU for *S*. However (and contrary to Step 13), it does *not* follow that it is true *by definition* that for any agent, *S*, faced with any decision under *risk*, *S* ought to choose out of what option(s) it is rational for *S* to weakly prefer over the alternative options in the choice set, or (i.e.) out of what option(s) maximize EU for *S*.

15. For any agent, *S*, faced with any decision under *risk*, *S* ought to 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* 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. Potential Objections to ECU Theory

In this section, I will respond to some possible objections to ECU theory. Firstly, I will lay out the *problem of act alterations* and the *problem of act versions* as they apply to ECU theory, and I will provide a solution to these problems taking inspiration from Gustafsson (2014). Secondly, I will respond to objections to the fact that ECU theory violates the IIA condition on choiceworthiness evaluations.

4.1 The Problem of Act Alterations and the Problem of Act Versions

One possible objection to ECU theory is that its ranking of options according to how choiceworthy they

are highly sensitive to the introduction of slight act alterations.¹⁴ We can call this the *problem of act alterations*. Let us consider the following decision problem (see Table 18), where option *a* involves ‘pressing a green button’, and option *b* involves ‘pressing a red button’ (I assume here that not pressing any button is not a feasible option).

TABLE 18 Decision matrix

	A (0.333)	B (0.333)	C (0.333)	EU	CECU
<i>a</i>	12	6	6	8	−2
<i>b</i>	3	12	12	9	2

Let us now consider the following variation on decision matrix 18 (see Table 19):

TABLE 19 Decision matrix

	A (0.333)	B (0.333)	C (0.333)	EU	CECU
<i>a</i>	12	6	6	8	1.667
<i>b</i> ₁	3	12	12	9	−1.667
<i>b</i> ₂	3	11	12	8.667	−2.333

The only difference between decision matrices 18 and 19 is that decision matrix 19 includes options *b*₁ and *b*₂, more specific *versions* of option *b* with option *b*₂ having a slightly different amount of utility in only one state—i.e., 11 utiles instead of 12 utiles in state *B*. Option *b*₁ involves pressing the red button with one’s right index finger, and option *b*₂ involves pressing the red button with one’s left index finger.¹⁵ Yet, introducing the more specific options *b*₁ and *b*₂ in decision matrix 19 reverses ECU theory’s choiceworthiness ranking of options *a* and *b* compared to decision matrix 18. On the other hand, introducing the more specific options *b*₁ and *b*₂ in decision matrix 19 does not reverse EU theory’s ranking of options *a* and *b* compared to decision matrix 18. Thus, the problem of act alterations would appear to be a problem for ECU theory (but not for EU theory).

How can we solve the problem of act alterations? I believe that we can do so by finding a principled way of individuating the various choice options that are available to an agent in any decision problem. One way of individuating options has been proposed by Gustafsson (2014). He has argued that choice options in a decision problem should be construed as sets of acts such that, for each set of acts, one could jointly intentionally perform, at any time *t*, all the acts in the set, but no additional acts (Gustafsson, 2014). One of the reasons given by Gustafsson is that if one construes choice options as individual acts, then one runs into the *problem of act versions* (Bergström, 1966; Castaneda, 1968). Consider the following example:

It is raining outside, but Ann will feel invigorated if she takes a brisk walk around the block (10 utiles), more so than if she stays inside (2 utiles). However, Ann has an injured toenail which causes her a great deal of pain when she tries to walk with her rain boots on. She will therefore experience a great deal of pain if she goes out for a walk wearing her rain boots (−30 utiles), more so than if she stays inside wearing her rain boots (−2 utiles). Luckily, Ann has a very comfortable pair of shoes which do not cause her any pain. However, there is a problem:

¹⁴ Thanks to Paul Bartha for raising this objection and for giving the following examples: Tables 18 and 19, and a variation on the ‘button pressing’ example.

¹⁵ I assume here for the sake of argument that pressing the button with one’s left index finger instead of one’s right index finger can make a difference in the decision situation.

it is raining very hard and her feet will get soaked if she wears her shoes. Ann will experience considerable discomfort if she goes out for a walk not wearing her rain boots (-15 utiles), more so than if she stays inside not wearing her rain boots (0 utiles).

Let us suppose that Ann assigns probability 1 to the state of the world as described above. Although the utility of the act ‘Ann stays inside’ is lower than that of the act ‘Ann goes out for a walk’, the utility of at least one version of the act ‘Ann stays inside’—that is, ‘Ann stays inside and does not wear her rain boots’ ($2 + 0 = 2$ utiles)—is greater than the utility of all versions of the act ‘Ann goes out for a walk’—that is, ‘Ann goes out for a walk and wears her rain boots’ ($10 + -30 = -20$ utiles) and ‘Ann goes out for a walk and does not wear her rain boots’ ($10 + -15 = -5$ utiles). Thus, intuitively, Ann ought to stay inside. However, if choice options are construed as individual acts, then EU theory counsels Ann *not* to stay inside, but instead to go out for a walk.

Therefore, to be intuitively plausible, ECU theory should be minimally cashed out as follows:¹⁶

For any agent, S , faced with any decision under certainty or any decision under risk and for any number of mutually exclusive and jointly exhaustive options, or sets of acts, a, b, c, d , and e , such that, for each set, S could jointly intentionally perform, at any time, t , all the acts in the set, but no additional acts,

- a is more choiceworthy than b , for S , at t , if and only if the CU/CECU of S jointly intentionally performing a at t is greater than the CU/CECU of S jointly intentionally performing b at t , and
- a is just as choiceworthy as b , for S , at t , if and only if the CU/CECU of S jointly intentionally performing a at t is equal to the CU/CECU of S jointly intentionally performing b at t .

This implies the following derivative decision rule for individual acts:¹⁷

For any agent, S , faced with any decision under certainty or any decision under risk and for any two mutually exclusive acts, a and b ,

- a is more choiceworthy than b , for S , at any time, t , if and only if a is logically entailed by every set of acts such that, for each set, S could jointly intentionally perform, at t , all the acts in the set, but no additional acts and such that, in accordance with ECU theory, the set of acts would be more choiceworthy for S , at t than each set of acts such that S could jointly intentionally perform, at t , all the acts in the set, but no additional acts and such that the set of acts logically entails b , and
- a is just as choiceworthy as b , for S , at any time, t , if and only if a is not more choiceworthy than b , and a is logically entailed by every set of acts such that, for each set, S could jointly intentionally perform, at t , all the acts in the set, but no additional acts and such that, in accordance with ECU theory, the set of acts would not be less choiceworthy for S , at t than each set of acts such that S could jointly intentionally perform, at t , all the acts in the set, but no additional acts and such that the set of acts logically entails b .

If we follow Gustafsson (2014) in construing *choice options* as sets of acts such that, for each set of acts, one could jointly intentionally perform, at any time t , all the acts in the set, but no additional acts, then we have a principled way of resolving the problem of act alterations. That is to say, it is decision matrix 20 (Table 20) that properly formalizes the ‘button pressing’ decision problem, and not

¹⁶ Inspired by Gustafsson (pp. 593–594).

¹⁷ Inspired by Gustafsson (p. 595).

decision matrix 18 or 19. In decision matrix 20, option a_1 involves ‘pressing the green button with one’s right index finger’, option a_2 involves ‘pressing the green button with one’s left index finger’, option b_1 involves ‘pressing the red button with one’s right index finger’, and option b_2 involves ‘pressing the red button with one’s left index finger’. In other words, option a_1 is ‘the act of pressing the green button and the act of pressing the button with one’s right index finger’, option a_2 is ‘the act of pressing the green button and the act of pressing the button with one’s left index finger’, option b_1 is ‘the act of pressing the red button and the act of pressing the button with one’s right index finger’, and option b_2 is ‘the act of pressing the red button and the act of pressing the button with one’s left index finger’.

TABLE 20 Decision matrix

	A (0.333)	B (0.333)	C (0.333)	EU	CECU
a_1	12	6	6	8	-1.336
a_2	12	6	6	8	-1.336
b_1	3	12	12	9	0.666
b_2	3	11	12	8.667	-0.666

The only difference between decision matrices 18 and 19, on the one hand, and decision matrix 20, on the other hand, is that decision matrix 20 includes options a_1 , a_2 , b_1 and b_2 , versions of options a and b with option b_2 having a slightly different amount of utility in only one state—i.e., 11 utiles instead of 12 utiles in state B . Introducing the more specific options a_1 , a_2 , b_1 and b_2 in decision matrix 20 does *not* reverse ECU theory’s choiceworthiness ranking of options a and b compared to decision matrix 18. Thus, the problem of act alterations is *not* a problem for ECU theory.

4.2 ECU Theory Violates the Independence of Irrelevant Alternatives

In personal communication, Paul Bartha objects to ECU theory on the grounds that it violates the IIA condition on choiceworthiness evaluations:

Let me turn to the problem of act alterations. Your proposed solution is to reject the decision matrix in Table 19, in favour of Table 20. This is a fair response to my example, since I suggested that the difference between two outcomes might correspond to pressing a button with different fingers. Your response addresses the particular example, but it doesn’t seem to me that it solves the fundamental problem. We can change the story so that there really are just three alternatives. With reference to Table 18: the assassin has to choose between poison (a) and a dagger (b). The three possible states are A, B, C with probability 1/3 each—for example, indicating three possible venues where the target individual might be found. Table 18 gives the utilities: the dagger is significantly better in states B and C and EU Max reflects the assassin’s overall preference for the dagger. ECU gives the same verdict. But just as the assassin is reaching for the dagger, he realizes that there is a slightly bigger (and harder to conceal) dagger in his kit—no other tools are available other than the poison (a) and the two daggers, now labelled choices b_1 and b_2 , with utilities as represented in Table 19 because the large dagger is a bit awkward to use if the target is in venue B. The assassin still prefers the dagger b_1 to the poison a and he certainly prefers it to the awkward dagger b_2 . But if he uses ECU reasoning, he now has to switch to the poison because that’s the only choiceworthy option. This seems absurd.

[This is] an objection to the fact that ECU gives up Independence of Irrelevant Alternatives.

The intuition that ECU theory's verdict in the above decision problem is absurd represents a *reductio ad absurdum* of ECU theory and, in particular, of its violation of the IIA condition on choiceworthiness evaluations. But that intuition is, I believe, nullified by my argument for ECU theory. Let me recapitulate that argument: For any agent, S , faced with any decision under *certainty* and for any option, a , for S , the choiceworthiness of a for S depends on both a 's utility and the utility of the best (i.e., most rationally preferred) option(s) to be forgone if S chooses a . The greater the difference in utility between a and the best option(s) to be forgone if S chooses a , the more choiceworthy a is for S . Therefore, the choiceworthiness of a for S is the difference in utility between a and the best option(s) to be forgone if S chooses a (i.e., a 's CU). And a is choiceworthy for S if and only if a maximizes choiceworthiness for S (i.e., CU) over the space of all alternatives in the choice set. This is the CU principle. Gandalf's principle says that for any agent, S , faced with any decision under *risk* and for any option, a , for S , how choiceworthy a is for S *within* each of the various states of the world in which S might find herself is to be calculated as if S were faced with a decision under *certainty* (that is to say, it is to be calculated in accordance with the CU principle). And how choiceworthy a is for S *tout court* is to be calculated by employing the nearly unanimously endorsed decision-theoretic aggregation rule, the *expected value* rule, which involves computing the expected value of options (see Peterson, 2017, p. 66). Therefore, how choiceworthy a is for S *tout court* is roughly the *probability-weighted sum* of a 's degrees of choiceworthiness *within* each of the possible states of the world (i.e., the probability-weighted sum of a 's CUs, or a 's ECU). More precisely, how choiceworthy a is for S *tout court* is a 's CECU, that is to say, the difference in ECU between a and whatever alternative(s) to a carry the greatest ECU. Again, a is choiceworthy for S *tout court* if and only if a maximizes choiceworthiness for S (i.e., CECU) over the space of all alternatives in the choice set. Since S *ought* to choose a if and only if a is *choiceworthy* for S , it follows that for any agent, S , faced with any decision under *risk*, S *ought* to choose out of whatever option(s) are choiceworthy for S (i.e., whatever option(s) maximize CECU), even in cases where whatever option(s) are choiceworthy for S differ from whatever option(s) it is rational for S to weakly prefer over the alternative options in the choice set (i.e., whatever option(s) maximize EU). Thus, ECU theory's verdict in the above decision problem does *not* represent a *reductio ad absurdum* of ECU theory and, in particular, of its violation of the IIA condition on choiceworthiness evaluations.

Nevertheless, if we give up the IIA (and the irrelevance of statewise dominated alternatives), then 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. This opens the door to strategic manipulation in the decision process.¹⁸ The worry can be overcome, however, if we accept some form of *probability discounting*. According to this view, for any given decision situation, any outcome with probability $\leq p$, where p is very close to 0, should be excluded from consideration in the decision process. As such, if we accept probability discounting, the *very improbable* outcomes of manufactured alternatives cannot alter the degrees of choiceworthiness of the other available options in the choice set. Endorsing some form of probability discounting in our case is not ad hoc since there are independent reasons for doing so: probability discounting, although not uncontentious, provides a unified solution to several thorny decision problems, such as the St. Petersburg paradox and Pascal's Mugging,¹⁹ that involve very small chances of enormous utility payoffs.²⁰

¹⁸ Thanks to Douglas Lackey for raising this objection and for wording suggestions.

¹⁹ The St. Petersburg paradox was analysed by Bernoulli (1738). Pascal's mugging was introduced by Bostrom (2009).

²⁰ Contemporary proponents of probability discounting include Chalmers (2017), Jordan (1994), Monton (2019), and Smith (2014, 2016). Kosonen (n.d.) helpfully formulates and assesses various versions of probability discounting.

5. Conclusion

When arguing for a new normative alternative to EU theory, decision theorists typically come up with decision situations in which their new theory, contrary to EU theory, delivers what seems to be the intuitively correct verdicts. In this paper, I have taken a different approach. I have shown that ECU theory gives verdicts that are different from those of EU theory, and I have presented several lines of reasoning to show that the verdicts of ECU theory are more plausible than those of EU theory, even in decision situations where EU theory, contrary to ECU theory, delivers what may seem *prima facie* to be the intuitively correct verdicts.

In this paper, I have assumed that agents ought to measure and rank their options (for the purpose of choice) in terms of how instrumentally rational they are and that they ought to choose out of what option(s) are instrumentally rational. Furthermore, I have assumed that EU is the appropriate criterion of rational preference. If both of these assumptions are correct, then, as I have argued in this paper, we can derive a notion of *choiceworthiness* according to which (i) agents ought to measure and rank their options (for the purpose of choice) in terms of how choiceworthy they are—that is, according to their CU (in decisions under certainty) or their CECU (in decisions under risk)—and (ii) agents ought to choose out of what option(s) are choiceworthy—that is, what option(s) maximize CU (in decisions under certainty) or what option(s) maximize CECU (in decisions under risk).

Acknowledgments

I wish to thank Paul Bartha, Kenny Easwaran, Douglas Lackey, John Quiggin, and several anonymous reviewers and commenters for very helpful feedback on earlier drafts of this paper.

References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21, 503–546. <https://doi.org/10.2307/1907921>
- 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>
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30, 961–81. <https://doi.org/10.1287/opre.30.5.961>
- Bergström, L. (1966). *The alternatives and consequences of actions*. Stockholm: Almqvist & Wiksell.
- 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>
- Bostrom, N. (2009). Pascal's mugging. *Analysis*, 69, 443–445. <https://doi.org/10.1093/analysis/anp062>
- Bradley, R., & Stefansson, H. O. (2017). Counterfactual desirability. *British Journal for the Philosophy of Science*, 68, 485–533. <https://doi.org/10.1093/bjps/axv023>
- 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. (2023). Normative theories of rational choice: Expected utility. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2023/entries/rationality-normative-utility/>
- Brown, J. L. D. (2023). On scepticism about ought simpliciter. *Australasian Journal of Philosophy*.

- <https://doi.org/10.1080/00048402.2023.2225527>
- Buchak, L. (2013). *Risk and rationality*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199672165.001.0001>
- Buchak, L. (2022). Normative theories of rational choice: Rivals to expected utility. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2022/entries/rationality-normative-nonutility/>
- Castaneda, H.-N. (1968). A problem for utilitarianism. *Analysis*, 28(4), 141–142. <https://doi.org/10.1093/analys/28.4.141>
- 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>
- Ellsberg, D. (1961). Risk, ambiguity and Savage axioms. *The Quarterly Journal of Economics*, 75, 643–669. <https://doi.org/10.2307/1884324>
- Feller, W. (1968). *An introduction to probability theory and its applications, volume I*. Wiley.
- Fishburn, P. C. (1982). Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26, 31–67. [https://doi.org/10.1016/0022-2496\(82\)90034-7](https://doi.org/10.1016/0022-2496(82)90034-7)
- Gintis, H. (2018). Rational choice explained and defended. In G. Bronner & F. Di Iorio (Eds.), *The Mystery of Rationality: Mind, Beliefs and the Social Sciences* (pp. 95–114). Springer. https://doi.org/10.1007/978-3-319-94028-1_8
- Gustafsson, J. E. (2014). Combinative consequentialism and the problem of act versions. *Philosophical Studies*, 167(3), 585–596. <https://doi.org/10.1007/s11098-013-0114-x>
- Gustafsson, J. E. (2022). *Money-pump arguments*. Cambridge University Press.
- Hammond, P. (1987). Consequentialism and the independence axiom. In B. Munier (Ed.), *Risk, Decision and Rationality*. Springer. https://doi.org/10.1007/978-94-009-4019-2_27
- Hammond, P. (1988). Consequentialist foundations for expected utility theory. *Theory and Decision*, 25, 25–78. <https://doi.org/10.1007/BF00129168>
- 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>
- Horton, J. (March 16, 2023). Practical reason and permissible preference. Dianoia Institute of Philosophy Seminar Series in Philosophy, Australia. <https://philevents.org/event/show/108913>
- Jeffrey, R. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.
- Jordan, J. (1994). The St. Petersburg Paradox and Pascal's Wager. *Philosophia*, 23, 207–222. <https://doi.org/10.1007/BF02379856>
- 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/>
- Kosonen, P. (n.d.). How to discount small probabilities. Unpublished manuscript. <https://www.petrakosonen.com/research/how-to-discount-small-probabilities.pdf>
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92, 805–24. <https://doi.org/10.2307/2232669>
- Meacham, C. J. G. (2019). Difference Minimizing Theory. *Ergo: An Open Access Journal of Philosophy*, 6, 999–1034. <https://doi.org/10.3998/ergo.12405314.0006.035>
- Monton, B. (2019). How to avoid maximizing expected utility. *Philosophers' Imprint*, 19, 1–25. <https://hdl.handle.net/2027/spo.3521354.0019.018>
- Peterson, M. (2017). *An introduction to decision theory*. Cambridge University Press. <https://doi.org/10.1017/9781316585061>
- 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.
- Robert, D. (2018). Expected comparative utility theory: A new theory of rational choice. *The Philosophical Forum*, 49, 19–37. <https://doi.org/10.1111/phil.12178>
- 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>
- Savage, L. (1954). *The foundations of statistics*. Wiley.

- 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. <https://doi.org/10.1093/acprof:oso/9780199759446.001.0001>
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton University Press.
- Wedgwood, R. (2003). Choosing rationally and choosing correctly. In S. Stroud & C. Tappolet (Eds.), *Weakness of Will and Practical Irrationality* (pp. 201–229). Oxford University Press.
<https://doi.org/10.1093/0199257361.003.0009>
- 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>
- Yager, R. R. (2017). Generalized regret based decision making. *Engineering Applications of Artificial Intelligence*, 65, 400–405. <https://doi.org/10.1016/j.engappai.2017.08.001>
- 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>
- Zhu, P. (2018). A quantitative approach to choose among multiple mutually exclusive decisions: comparative expected utility theory. arXiv:1801.02422 [q-fin.EC]. <https://doi.org/10.48550/arXiv.1801.02422>