Rational Monism and Rational Pluralism

Jack Spencer

October 18, 2019

We can distinguish two sorts of consequentialist theories of rational choice. On one side are expected value theory, conditional expected value theory, minimax, maximin, and other forms of *rational monism*. On the other side are various forms of *rational pluralism*, none of which enjoy much familiarity. Rational monism is more commonly defended. But I believe that consequentialists should favor rational pluralism.

1 Rational Monism v. Rational Pluralism

The dispute between rational monists and rational pluralists is a dispute about an analogical claim in metaethics. To put the analogical claim in its proper context, it will be helpful to remind ourselves about consequentialism and its reductive ambitions.

According to consequentialism, the realization of value is all that fundamentally matters.¹ The deontic, therefore, is reducible: every normatively significant deontic notion somehow reduces to facts about value and its realization.

There are at least two parts to the reductive task facing consequentialists because there are at least two normatively significant deontic notions. There

 $^{^{1}}$ See *e.g.* Bentham (1961[1789]), Mill (1988[1861]), Moore (1903;1912) and Ramsey (1990[1926]).

is *objective* permission and also *rational* (sometimes called 'subjective') permission.² The distinction between the two is made vivid by examples like the following:

Boxes like Miners. There are three opaque boxes, arranged left-to-right. The agent must choose exactly one. The agent knows that the middle box contains \$9. Of the other two boxes, she knows that one contains \$0 and that the other contains \$10. But she is uncertain which box contains \$10, and divides her credence equally between the two possibilities. (In fact, the right box contains \$10.)

An agent facing *Boxes like Miners* is objectively required to take the right box and rationally required to take the middle box.

An adequate consequentialist reduction of objective permission is already at hand. According to consequentialists, whenever an agent faces a decision, each of the agent's options has an actual value, and objective permission reduces to actual value maximization. Actual value maximization is not just necessarily coextensive with objective permission; according to consequentialists, it's prior. Whenever an option is objectively permissible, it is so because it maximizes actual value. For example, in Boxes like Miners, if we equate dollars and units of value, the actual values of choosing the left, middle, and right boxes are, respectively, 0, 9, and 10. And according to consequentialists, that's what makes choosing the right box objectively required.

Reducing *rational* permission is harder. I will assume that the reduction takes the same basic shape: that, whenever an option is rationally permissible, it is so because it maximizes some quantity.³ But the hard-to-answer

²The claim that there are both objective and rational permissions is not entirely uncontroversial; see *e.g.* [redacted], Kolodny and MacFarlane (2010), and Thomson (2008).

³For reasons discussed in §5.3, it is the *stable* maximization, not the *mere* maximization, of a quantity that makes options rationally permissible.

question remains: what quantity, or quantities, feature in the reduction? What *makes* options rationally permissible?

It's here that rational monists and rational pluralists disagree. According to rational monists, some quantity stands to rational permission as actual value stands to objective permission. Actual value is the universal objective-maker: objective permissions always reduce to the maximization of actual value. According to rational monists, some special quantity Q is the universal rational-maker: rational permissions always reduce to the maximization of it.

Of course, rational monists disagree about what the universal rational-maker is. Some think it's expected value.⁴ Others think it's conditional expected value,⁵ or maximin value,⁶ or risk-adjusted expected value.⁷ But they agree that *some* quantity is the universal rational-maker.

According to rational pluralists, no quantity stands to rational permission as actual value stands to objective permission.⁸ Consequentialism is true. Whenever an option is rationally permissible, it's made so by maximizing some quantity. But no quantity is the *universal* rational-maker. On one occasion it might be the maximization of Q_1 that makes options rationally permissible; on another occasion it might be the maximization of Q_2 that makes options rationally permissible, in which case the Q_1 -values of options will have no bearing on what the agent is rationally permitted to choose.

The essay below has a negative part and a positive part. The negative part is contained in §§2–4. By combining some recent work in decision theory with some metaethical considerations, I will argue against rational monism. The positive part is contained in §§5–7. Many forms of rational pluralism are

⁴See *e.g.* Hammond (1988), Joyce (1999; 2012; 2018), Lewis (1981), von Neumann and Morgenstern (1944), Pettigrew (2015), Ramsey (1990[1926]), Savage (1954), Skyrms (1982; 1984; 1990), Sobel (1994), and Stalnaker (1981).

⁵See e.g. Ahmed (2014a), Eells (1982), and Jeffrey (1965; 1983).

⁶See *e.g.* Rawls (1971).

⁷See *e.g.* Buchak (2013).

 $^{^{8}}$ There are other rational pluralists; see *e.g.* Robinson (dissertation) and Weirich (1988; 2004).

messy and unsystematic. But the form of rational pluralism that I develop is principled. I believe that what makes a quantity a rational-maker on a given occasion is being the best quantity that can guide the agent on that occasion. I take quantities to be mathematical objects: functions from decision problems to functions from world-option pairs to real numbers. I will propose a way of scoring these quantities; I will put forward conditions that a quantity must meet in order to be capable of guiding a given agent on a given occasion; and I will argue that the rational-maker on a given occasion is the highest-scoring quantity that can guide the agent on that occasion. The form of rational pluralism that I develop arises naturally from this constrained optimization conception of occasional rational-making.

2 Two Rules for Reducing

Consequentialists ultimately need a reduction of rational permission that holds true for all agents, ideal and nonideal. As I will point out in the final section, it's very unlikely that any form of rational monism holds true both for ideal and nonideal agents. But many philosophers are inclined to think that some form of rational monism holds true for ideal agents. So, for now, let's set nonideal agents aside, and focus exclusively on ideal agents, who have unlimited powers of introspection and deduction.

The dispute between rational monists and rational pluralists (concerning ideal agents) can be made more precise by appealing to some familiar formalism. In the usual way, let's represent an (ideal) agent facing a decision with a decision problem, which we'll take to be an ordered quadruple, $\langle C, u, A, K \rangle$.

The first coordinate, C, is the agent's *credence function*, a probability function that represents the agent's confidence in each proposition. Here and throughout, I assume the realist view that credences are among the agent's fundamental psychological states.⁹

⁹For more on the dispute between instrumental and realist views of credences, see *e.g.* Eriksson and Háyek (2007), List and Dietrich (2016), and Pettigrew (2019).

The second coordinate, u, is the agent's *utility function*, which maps each possible world to a real number and thereby represents the degree to which the agent finds the world desirable. I will assume the same realist view of utilities that I assume of credences.

The third coordinate, $A = \{a_1, a_2, ..., a_n\}$, is the set of *options*, which are propositions that the agent can make true by deciding. I will assume that options are pairwise exclusive.

The fourth and final coordinate, $K = \{k_1, k_2, ..., k_m\}$, is the set of dependency hypotheses, which are propositions that fully specify how things do and do not depend causally on the agent's decision.¹⁰ I will assume that dependency hypotheses are pairwise exclusive and compossible with each option.

Taken as a whole, a decision problem $d = \langle C, u, A, K \rangle$ represents an (ideal) agent of type $\langle C, u \rangle$ facing a decision of type $\langle A, K \rangle$.

If D is the set of decision problems, then consequentialists, on account of their reductive ambitions, are committed to two metaethical functions that have D as their domain.

The first is the rule for reducing objective permission, a function that maps each decision problem to the quantity the maximization of which makes options objectively permissible relative to that decision problem. In principle, a dispute between monists and pluralists could arise here. Objective monists would claim that the rule for reducing objective permission is a constant function, and objective pluralists would disagree. But I will ignore this dispute and assume, here and throughout, that the rule for reducing objective permission maps every decision problem to actual value.

The second is the rule for reducing rational permission, a function that maps each decision problem to the quantity the maximization of which makes options rationally permissible relative to that decision problem. This is the function that gives precision to the dispute between rational monists and rational pluralists (concerning ideal agents). Rational monists believe that this function is constant, and rational pluralists, like me, disagree.

¹⁰Cf. Lewis (1981) and Skyrms (1982).

It is worth pausing here to say more precisely what I take quantities to be. Most of the important examples will be familiar. But let me offer an abstract characterization.¹¹

Let $W = \{w_1, w_2, ..., w_i\}$ be the set of possible worlds, and, for simplicity, assume that W is finite. A mathematical quantity is any function that maps decision problems to functions that map option-world pairs to real numbers. In other words, if Z is a mathematical quantity and $d = \langle C, u, A, K \rangle$, then Z maps d to some function that maps each $\langle a, w \rangle$ to a real number. Mathematical quantities are then partitioned into quantities by their maximizational structure. If Z is mathematical quantity, let Max(Z, w, d) be the set of options that maximize Z at $\langle w, d \rangle$. Two mathematical quantities, Z_1 and Z_2 , are equivalent, then, just if, for any $\langle w, d \rangle$, $Max(Z_1, w, d) = Max(Z_2, w, d)$.

More simply, then, a quantity can be thought of as a function that maps each $\langle w, d \rangle$ to the set of options that maximize the quantity relative to $\langle w, d \rangle$. Given this abundant conception of quantities, it is almost uncontroversial that some quantity is necessarily coextensive with rational permission—see §4.3. But what we seek is a consequentialist reduction of rational permission, and necessary coextensiveness does not suffice for reduction. If rational monism (concerning ideal agents) is true, then some quantity, in the sense above, is such that: whenever an (ideal) agent faces a decision, it is the maximization of that quantity that makes options rationally permissible.

3 Independent Monism

3.1 V-monism and U-monism

The two most commonly defended forms of rational monism are V-monism (sometimes called 'conditional expected value theory' or 'evidential decision theory') and U-monism (sometimes called 'expected value theory' or 'causal decision theory'). Both V and U can be defined in terms of actual value,

¹¹This characterization assumes that we have set nonideal agents aside.

which, itself, can be defined using the formalism above.

Since dependency hypotheses fully specify how things the agent cares about do and do not depend causally on the agent's decision, the actual value of an option depends only on which dependency hypothesis holds. Let ak be the conjunction of option a and dependency hypothesis k. Every ak-world has the same utility, so, if w is an ak-world, u(ak) = u(w).

The V-value of a—i.e., the conditional expected value of a—is the agent's expectation of the actual value of a, conditional on a:

$$V(a) = \sum_{K} C(k|a)u(ak).$$

According to V-monists, the rule for reducing rational permission maps every decision problem to V.

The U-value of a—i.e., the expected value of a—is the agent's (unconditional) expectation of the actual value a:

$$U(a) = \sum_{K} C(k)u(ak).$$

According to U-monists, the rule for reducing rational permission maps every decision problem to U.

Both V-monism and U-monism face challenges. Newcomb problems challenge V-monism, and unstable problems—such as Bostrom's (2001) Meta-Newcomb, Egan's (2007) $Psychopath\ Button$, and Ahmed's (2014b) $Dicing\ with\ Death$ —challenge U-monism. And as we will see in §3.4, if both challenges succeed—if Newcomb problems are counterexamples to V-monism and unstable problems are counterexamples to U-monism—then many forms of rational monism that otherwise might have seemed promising can be shown to be extensionally inadequate.

¹²Other discussions of Newcomb problems and/or unstable problems include: [redacted], Ahmed (2012; 2014a; 2014b; 2018), Artnzenius (2008), Bales (2018), Bassett (2015), Briggs (2010), Eells (1982), Eells and Harper (1991), Gallow (MS), Gibbard and Harper (1978), Gustafsson (2011), Hare and Hedden (2016), Horgan (1981), Hunter and Richter (1978), Jeffrey (1983), Joyce (1999; 2012; 2018), Lewis (1981), Nozick (1969), Oddie and Menzies (1992), Rabinowicz (1988; 1999), Skyrms (1982; 1984; 1990), Stalnaker (1981), Wedgwood (2013), Weirich (1985; 1988; 2004), and Wells (forthcoming).

3.2 Newcomb Problems

Given the set of dependency hypotheses, we can define strict K-domination: option a_i strictly K-dominates option a_j , relative to credence function C, just if every k to which C assigns nonzero probability is such that $u(a_ik) > u(a_ik)$.

Two principles that connect strict K-domination and rational choice then suggest themselves. The first is stronger:

K-Elimination: If, relative to an agent's credence function, option a_i strictly K-dominates option a_j , then it is not rationally permissible for the agent to choose option a_j .

The second is weaker:

K-Selection: If, relative to an agent's credence function, option a_i strictly K-dominates every other option, then the agent is rationally required to choose option a_i .

In Newcomb's problems, V-monism violates both:

Newcomb. There is a transparent box and an opaque box. The agent has two options. She can take only the opaque box, or she can take both boxes $(a_1 \text{ or } a_2)$. The transparent box contains \$1,000. The opaque box contains either \$0 or \$1,000,000, depending on a prediction made yesterday by a reliable predictor. If the predictor predicted that the agent would take both boxes, the opaque box contains \$0. If the predictor predicted that the agent would take only the opaque box, the opaque box contains \$1,000,000. The agent knows all of this.

Taking both boxes strictly K-dominates taking only the opaque box, and thus strictly K-dominates every other option. But V-monism recommends taking only the opaque box:

```
V(a_1) \approx (0)(0) + (1)(1,000,000) = 1,000,000 > V(a_2) \approx (1)(1,000) + (0)(1,001,000) = 1,000.
```

I am convinced that an agent facing Newcomb is rationally required to take both boxes, so I reject V-monism.¹³

3.3 Unstable Problems

But U-monism also faces a challenge. We can distinguish stable and unstable decision problems. If C is an agent's credence function, let C^a be the agent's credence function conditional on a. A decision problem $d = \langle C, u, A, K \rangle$ is stable just if some $a \in A$ maximizes U both relative to d and relative to $d^a = \langle C^a, u, A, K \rangle$, and unstable, otherwise. Some unstable problems appear to be counterexamples to U-monism. For example:¹⁴

The Frustrater. There is an envelope and two opaque boxes, A and B. The agent has three options. She can take box A, box B, or the envelope $(a_A, a_B, \text{ or } a_E)$. The envelope contains \$40. The two boxes together contain \$100. How the money is distributed between the boxes depends on a prediction made yesterday by the Frustrater, a reliable predictor who seeks to frustrate. If the Frustrater predicted that the agent would take box A, box B contains \$100. If the Frustrater predicted that the agent would take box B, box A contains \$100. If the Frustrater predicted that the agent would take the envelope, each box contains \$50. The agent knows all of this.

If we equate dollars and units of value, then, no matter what the agent's credences are, $U(a_E) = 40$ and $U(a_A) + U(a_B) = 100$. So, no matter what the agent's credences are, a_A and/or a_B maximize U. But there is a strong

¹³For a defense of two-boxing, see [redacted].

¹⁴This example is from [redacted]. It's assumed that the agent is unable to randomize their choice.

intuition that an agent facing *The Frustrater* is rationally required to take the envelope, and that intuition is undergirded by an argument that I find compelling.¹⁵

Unstable problems do not only challenge U-monism; they also challenge K-domination principles. The weaker principle, K-Selection, is safe and should be accepted. But the stronger principle, K-Elimination, is challenged by examples like the following:¹⁶

The Semi-Frustrater. There are two opaque boxes, one white and one black. The agent has four options. She can point to either box with either hand $(a_{RW}, a_{LW}, a_{RB}, \text{ or } a_{LB})$. One of the boxes contains \$0; the other contains \$100. The agent receives whichever box she points to. Which box contains which sum depends on a prediction made yesterday by the Semi-Frustrater, a predictor who seeks to frustrate. If the Semi-Frustrater predicted that the agent would point to the black box, the white box contains \$100. If the Semi-Frustrater predicted that the agent would point to the white box, the black box contains \$100. There are two left-right asymmetries. First, the agent will receive an extra \$5 if she points to a box with her right hand. Second, because the Semi-Frustrater scans only half of the agent's brain (the half that controls motor movement on the right-hand side of the body), the Semi-Frustrater is a 90%-reliable predictor of right-handed box-pointings and only a 50%-reliable predictor of left-handed box-pointings. The agent knows all of this.

Each right-handed option strictly K-dominates the corresponding left-handed option. The two relevant dependency hypotheses are k_W , the proposition that the white box contains \$100, and k_B , the proposition that the black box contains \$100, and:

¹⁵See [redacted].

¹⁶This example is from [redacted].

```
u(a_{LW}k_W) = 100 < 105 = u(a_{RW}k_W);

u(a_{LW}k_B) = 0 < 5 = u(a_{RW}k_B);

u(a_{LB}k_W) = 0 < 5 = u(a_{RB}k_W); and

u(a_{LB}k_B) = 100 < 105 = u(a_{RB}k_B).
```

But there is a strong intuition that the rationally permissible options are the left-handed options. I am compelled by this intuition, so I reject both U-monism and K-Elimination.

3.4 Against Independent Monism

One can always bite the bullet. An inveterate V-monist might insist that an agent facing Newcomb is rationally required to take only the opaque box, and an inveterate U-monist might insist that an agent facing $The\ Frustrater$ is not rationally permitted to take the envelope. But whether we want to go in for this sort of bullet-biting depends in part on how attractive the intuition-accommodating alternatives are. So I propose that we provisionally take the intuitions at face value, and thus take V-monism and U-monism to stand refuted.

If we do so, we can prove an important limitative result.¹⁸ Say that a quantity is *independent* if its comparative relations holds independently of alternatives. Height is an example. If x and y are two people in a room and x is taller than y, then x continues to be taller than y, no matter who enters or exits the room.

When an object maximizes an independent quantity, it continues to do so upon the elimination of alternatives. If x is the tallest person in the room, she continues to be tallest, no matter who else exits the room.

 $[\]overline{^{17}}$ For an inveterate defense of V-monism, see e.g. Ahmed (2014a). For an inveterate defense of U-monism, see e.g. Harper (1996) and Joyce (2012; 2018).

¹⁸The proof draws on Ahmed (2012).

Most familiar quantities—height, mass, age, wealth, velocity, brightness, actual value, U, and V, just to mention a few—are independent quantities, and almost every form rational monism on offer is centered on some independent quantity or other. But given two principles that connect strict K-domination to rational choice, we can prove that rational permission is not coextensive with the maximization of any independent quantity. The principles are K-Selection, the weaker of the K-domination principles above, and:

K-Permission: It is sometimes rationally permissible for an agent to choose an option that is strictly K-dominated relative to the agent's credence function.

The motivation for K-Permission comes from cases like *The Semi-Frustrater*.

The proof is then straightforward. Take any example in which an agent is rationally permitted to choose an option that is strictly K-dominated relative to the agent's credence function. For instance, take *The Semi-Frustrater*. If Q-monism is true, then, since the left-handed options are rationally permissible, $Q(a_{LW})$ and $Q(a_{LB})$ are equal to one another and exceed both $Q(a_{RW})$ and $Q(a_{RB})$.

Now eliminate all but two of the agent's options, keeping a K-dominated option that is rationally permissible and an option that K-dominates it. For example,

The Demi-Semi-Frustrater. Everything is as in The Semi-Frustrater, except that the agent cannot point to the black box.

If Q is independent, then, relative to the *The Demi-Semi-Frustrater*, $Q(a_{LW})$ exceeds $Q(a_{RW})$. Hence, according to Q-monism, the agent facing *The Demi-Semi-Frustrater* is rationally required to point left-handedly. But this claim contradicts K-Selection, since, in *The Demi-Semi-Frustrater*, pointing right-handedly to the white box strictly K-dominates every other option.¹⁹

¹⁹ The Demi-Semi-Frustrater is just another Newcomb problem.

It follows, then, that rational permission is not coextensive with Qmaximization. And since Q was chosen arbitrarily, the conclusion generalizes. If K-Permission and K-Selection are both true, rational permission
is not coextensive with the maximization of any independent quantity.

This limitative result can be stated more clearly with the help of some terminology. Call any form of rational monism centered on an independent quantity, an *independent monism*. Every independent monism verifies:

Alpha: For any decision problem $d = \langle C, u, A, K \rangle$, if $a \in A$ is rationally permissible relative to d, and $a \in A' \subset A$, then a is rationally permissible relative to $d' = \langle C, u, A', K \rangle$.²⁰

Our limitative result, then, is this: Alpha is false, and hence every form of independent monism is extensionally inadequate. If we take the intuitions about Newcomb problems and unstable problems at face value, then the only hope for rational monism is some form of dependent monism.

4 Dependent Monism

4.1 V-ratificationism

The most familiar theory that falsifies Alpha is V-ratificationism. ²¹

Let's say that option a is ratifiable, relative to decision problem $d = \langle C, u, A, K \rangle$, just if a maximizes U relative to $d^a = \langle C^a, u, A, K \rangle$, and non-ratifiable, otherwise. According to V-ratificationism, options are lexically ordered by ratifiability, and then ranked by V-value. Hence, if any option is ratifiable, the rationally permissible options are the options that maximize V among the ratifiable options, and if no option is ratifiable, the rationally permissible options are the options that maximize V, simpliciter.

²⁰ Cf. Sen (1970).

²¹ Cf. Jeffrey (1983).

There are different ways to understand V-ratificationism, but we'll understand it as a particular form of dependent monism.²² The math here isn't important, but let me lay it out anyway. We can bound and normalize the V-values of options by taking their arctangent and dividing by π . We can then add a ratifiability score: if a is ratifiable (relative to $d = \langle C, u, A, K \rangle$) then let $r(a) = \frac{1}{2}$, and if a is nonratifiable (relative to d) then let $r(a) = -\frac{1}{2}$. We then can define the J-value of option a (relative to d) as follows:

$$J(a) = \frac{\tan^{-1}(V(a))}{\pi} + r(a).$$

The J-values of nonratifiable options lie on the open interval (-1,0) and are ordered by their V-values, and the J-values of ratifiable options lie on the open interval (0,1) and are ordered by their V-values. Hence, so long as we care only about the ordinal rankings of options in terms of choiceworthiness, we can take V-ratificationism to be J-monism.

There are some familiar virtues of V-ratificationism. Unlike any form of independent monism, V-ratificationism validates both K-Selection and K-Permission.²³ Moreover, it gives the correct verdicts in all of the cases above: two-boxing in Newcomb, pointing right-handedly in $The\ Demi-Semi-Frustrater$, taking the envelope in $The\ Frustrater$, and pointing left-handedly in $The\ Semi-Frustrater$.

But V-ratificationism also has some vices. It admits of counterexamples and—like every form of dependent monism, I think—is metaethically dubious. The metaethical vice is more important, but let's start with a counterexample.

4.2 A Counterexample to V-ratificationism

V-ratificationism predicts that ratifiable options are always more choiceworthy than nonratifiable options. But that prediction is wrong, as the following

 $^{^{22}}$ We can also formulate V-ratificationism as a form of rational pluralism. The monistic and pluralistic formulations of V-ratificationism do not differ extensionally, but they differ metaethically—see $\S4.3$.

²³If an option K-dominates every other option, then it is the only ratifiable option.

example, adapted from Skyrms (1984), makes clear:

Three Shells. There are three shells, A, B, and C. The agent must choose exactly one $(a_A, a_B, \text{ or } a_C)$. How much money is contained in each shell depends on a prediction made yesterday by a reliable predictor. If the predictor predicted that the agent would choose shell A, then A contains \$5, B contains \$0, and C contains \$0. If the predictor predicted that the agent would choose shell B, then A contains \$0, B contains \$9, and C contains \$10. If the predictor predicted that the agent would choose shell C, then A contains \$0, B contains \$10, and C contains \$9. The agent knows all of this.

Choosing shell A is the only ratifiable option. Hence, according to V-ratificationism, an agent facing $Three\ Shells$ is rationally required to choose A, no matter what credences she has. But that's wrong. If the agent is confident that she will choose A, then she is rationally required to choose A, since she's then confident that A contains \$5 and that the other two shells contain \$0. But if the agent is not confident that she will choose shell A, it's not even rationally permissible for her to do so. $Contra\ V$ -ratificationism, ratifiable options are not, merely by virtue of being ratifiable, more choiceworthy than nonratifiable options.

4.3 Other Forms of Dependent Monism

There are other forms of dependent monism on offer. Wedgwood (2013) defends what we might call B-monism, where:

$$B(a) = \sum_{K} (C(k|a)(u(ak) - \frac{\sum_{A} u(ak)}{\#A})).$$

Gallow (MS) defends what we might G-monism, where:

$$G(a_i) = -1 \times (\max_{a_j \in A} (\sum_K C(k|a_i)u(a_jk)) - \sum_K C(k|a_i)u(a_ik)).$$

Neither makes much progress over V-ratificationism: Three Shells is a counterexample to both.²⁴

But examples and counterexamples can only get us so far. V-ratificationism (a.k.a. J-monism), B-monism, and G-monism are just three forms of dependent monism. There are uncountably many others. And given optimism—the claim, which I accept, 25 that there is a determinate fact of the matter about which options are rationally permissible relative to every $\langle w, d \rangle$ —we can construct an extensionally adequate form of dependent monism as follows:

Let g be a function that maps each decision problem to the set of world-option pairs that are rationally permissible. Hence, $\langle a, w \rangle \in g(d)$ if and only if a is rationally permissible at $\langle w, d \rangle$. Let a mathematical quantity, Z_g , be defined as follows. If $\langle a, w \rangle \in$

$$\begin{split} B(a_A) &= \sum_K (C(k|a_A)(u(ak) - \frac{\sum_A u(ak)}{\#A})) \approx (1)(5 - \frac{5}{3}) + (0)(0 - \frac{19}{3}) + (0)(0 - \frac{19}{3})) = \frac{10}{3}; \\ B(a_B) &= \sum_K (C(k|a_B)(u(ak) - \frac{\sum_A u(ak)}{\#A})) \approx (0)(0 - \frac{5}{3}) + (1)(9 - \frac{19}{3}) + (0)(10 - \frac{19}{3}) = \frac{8}{3}. \\ B(a_C) &= \sum_K (C(k|a_C)(u(ak) - \frac{\sum_A u(ak)}{\#A})) \approx (0)(0 - \frac{5}{3}) + (0)(10 - \frac{19}{3}) + (1)(9 - \frac{19}{3}) = \frac{8}{3}; \text{ and} \end{split}$$

Following Gallow (MS), let $E(a_i|a_j) = \sum_K C(k|a_j)u(a_i)$. Then $G(a_i) = -1 \times (\max_{a_j \in A} E(a_j|a_i) - E(a_i|a_i))$. In Three Shells, no matter what the agent's credences are: $E(a_A|a_A) \approx 5$, $E(a_A|a_B) \approx 0$, $E(a_A|a_C) \approx 0$, $E(a_B|a_A) \approx 0$, $E(a_B|a_B) \approx 9$, $E(a_B|a_C) \approx 10$, $E(a_C|a_A) \approx 0$, $E(a_C|a_B) \approx 10$, and $E(a_C|a_C) \approx 9$. Hence

$$G(a_A) \approx -1 \times (5-5) = 0;$$

 $G(a_B) \approx -1 \times (10-9) = -1;$ and
 $G(a_C) \approx -1 \times (10-9) = -1.$

 25 The most sustained argument against optimism is Briggs' (2010). Briggs argues that any adequate decision theory must verify two principles—a Pareto principle and a self-sovereignty principle—and then proves that no decision theory can verify both. I think that an adequate decision theory must falsify both: the Pareto principle is refuted by The Semi-Frustrater, and the self-sovereignty principle is refuted by Three Shells.

²⁴In *Three Shells*, no matter what the agent's credences are:

g(d), then $Z_g(d)(\langle a, w \rangle) = 1$, and if $\langle a, w \rangle \notin g(d)$, then $Z_g(d)(\langle a, w \rangle) = 0$. If Q_g is the quantity that contains Z_g , then Q_g -maximization is coextensive with rational permission.

We can also construct any number of extensionally adequate forms of rational pluralism.

The disagreement between an extensionally adequate form of dependent monism and an extensionally adequate form of rational pluralism is real. Dependent monists claim that some dependent quantity is the universal rational-maker. Rational pluralists deny that any quantity is the universal rational-maker, and many rational pluralists will insist that options are always made rationally permissible by some independent quantity or other. But the disagreement cannot be settled merely by appeal to examples and counterexamples.

4.4 The Problem of Consequentialist Credentials

Instead, the disagreement is to be settled on metaethical grounds.

A consequentialist who claims that it is the maximization of Q that makes options rationally permissible relative to decision problem d must be able to provide a consequentialist explanation for why it is the maximization of Q, specifically, and not some other quantity, that makes options rationally permissible. We are owed some story about how we get from consequentialism—the claim that the realization of value is all that fundamentally matters—to the claim that Q-maximization is what makes options rationally permissible relative to d. We can call this explanatory task, the problem of consequentialist credentials. My reason for favoring rational pluralism over dependent monism is that I think it's better positioned to answer the problem of consequentialist credentials.

The usual way of trying to answer the problem of consequentialist creden-

tials is by proving a representation theorem. ^{26,27} In a representation theorem, some formal conditions are identified and alleged to be requirements on rational choice. It is proved that, if the formal conditions identified really are requirements on rational choice, then rational permission is coextensive with the maximization of some quantity. And those with reductive ambitions then go on to claim that rational permission reduces to the maximization of that quantity. Consequentialists who respond to the problem of consequentialist credentials by proving a representation theorem inevitably are lead to rational monism. The whole point of proving a representation theorem is to arrive at a representing quantity: that is, a quantity the maximization of which is coextensive with rational permission relative to every decision problem. And all of the familiar representation theorems lead to some form of independent monism, since, in each case, the identified conditions entail Alpha.²⁸ (The importance of the limitative result above thus becomes apparent. A proof that Alpha is false is, inter alia, a proof that none of the familiar representation theorems succeed.)

The problem of consequentialist credentials is an acute problem for dependent monism because, as the normative complexity of the quantity that is alleged to be the universal rational-maker increases, so too does the difficulty of exhibiting its consequentialist credentials. It is one thing to try to explain why U—i.e., expected value—should be the universal rational-maker.

The fundamental source for the normative force of expected utility theory lies in what are known as representation theorems... (Bermúdez 2009: 30).

The standard method for justifying any version of expected utility theory involves proving a representation theorem... (Joyce 1999: 4).

For an alternative way of trying to answer the problem of consequentialists credentials, see Hammond (1988).

²⁶Some representative quotations:

²⁷For an orthogonal critique of using representation theorems to answer the problem of consequentialist credentials, see Meacham and Weisberg (2011).

²⁸The familiar representation theorems include: Bolker (1967), Buchak (2013), Joyce (1999), von Neumann and Morgenstern (1944), and Savage (1954).

It is another thing entirely to try to explain why, say, J—i.e., the quotient of the arctangent of conditional expected value and π plus a ratifiability score—should be the universal rational-maker. It is hard to imagine a representation theorem that purports to prove that J is the representing quantity. Indeed, it is hard to imagine any satisfying explanation for why the maximization of J, specifically, should be so normatively important. And I think the same goes for the other dependent quantities, too. I cannot prove that dependent monists will be unable to provide a satisfactory solution to the problem of consequentialist credentials, but I regard their prospects as very dim indeed.

One might have thought that the prospects for rational pluralism are even dimmer still, since the usual way of trying to answer the problem of consequentialist credentials inevitably leads to rational monism. But there is an alternative and, I think, superior way to try to answer the problem of consequentialist credentials that does not inevitably lead to rational monism.

5 Constrained Optimization

I think consequentialists should answer the problem of consequentialist credentials in the same way they answer every problem: namely, by optimizing.

The picture I have in mind goes roughly as follows. Each quantity is assigned a score, which measures the degree to which its maximization conduces to the realization of value. The rational-making quantity is then the quantity that scores highest, subject to a guidance constraint.

As a bit of terminology, if an agent of type $\langle C, u \rangle$ faces a decision of type $\langle A, K \rangle$, then we'll say that the agent is *involved* in $d = \langle C, u, A, K \rangle$. And if an agent involved in d is capable of being guided by some quantity Q, then we'll say that Q is d-guiding. Rational permissions are, by their very nature, capable of providing guidance, so the rule for reducing rational permission maps every decision problem d to some d-guiding quantity. But subject to this guidance constraint, we optimize, since the realization of value is all that fundamentally matters. Hence,

Rational Optimization: The rule for reducing rational permission maps each decision problem d to the highest-scoring d-guiding quantity.²⁹

Rational Optimization is my proposed solution to the problem of consequentialist credentials. If the rule for reducing rational permission maps d to Q, it does so because Q is the highest-scoring d-guiding quantity: the quantity the maximization of which best conduces to the realization of value, among the quantities that can guide an agent involved in d. Objective permission is a matter of unconstrained optimization—the rule for reducing objective permission maps every decision problem, d, to the best quantity, namely, actual value. Rational permission is a matter of constrained optimization—the rule for reducing rational permission maps every decision problem, d, to the best d-guiding quantity.

If we accept Rational Optimization, then two matters become pressing. We want to know how to score quantities, and we want to know what it is for a quantity to be d-guiding. My goal in the remainder of this section is to make progress on these two matters. I offer a partial account of how to score quantities and a full account of what it is for a quantity to be d-guiding. By combining the two, we can shed light on the puzzling examples above.

5.1 Scoring

As I envisage it, the score of a quantity should be determined by two factors. The first factor is the various d-scores of the quantity, where, for some decision problem d, the d-score of Q, S(Q,d), is a measure of the degree to which the maximization of Q conduces to the realization of value relative to d, specifically. The second factor is some measure, M(d), assigned to each $d \in D$. The score of a quantity is then some average of its d-scores, weighted by the M(d)'s.

 $^{^{29}}$ The idea that we should be scoring quantities (or programs) and optimizing subject to some constraint has been a mainstay of work in bounded rationality, especially in computer science; see *e.g.* Halpern *et al.* (2014), Icard (2018), and Russell and Subramanian (1995).

My account of how to score quantities is partial because I do not know what the M(d)'s should be. If we put enough constraints on the space of decision problems, then certain measures, like an indifference measure, are tempting and plausible. But without imposing constraints on the space of decision problems, it is hard to know what the M(d)'s should be. So I will leave that matter undecided.

I do, however, have a proposal for how to d-score quantities. To a first approximation, I think that the d-score of Q should be the actual value that an agent involved in d expects to realize by choosing a Q-maximizing option.

More formally, let Max(Q, w, d) be the set of options that maximize Q at $\langle w, d \rangle$, and let #Max(Q, w, d) be the number of options contained in Max(Q, w, d).

Let @(a, w, d) be the actual value of option a at $\langle w, d \rangle$. For example, in Boxes like Miners, if a is the option of choosing the right box, and w_1 and w_2 are worlds at which the right box contains \$10 and \$0, respectively, then, equating dollars and units of value, $@(a, w_1, d) = 10$ and $@(a, w_2, d) = 0$.

Let @(Q, w, d) be the average of the actual values of the Q-maximizing options at $\langle w, d \rangle$ —that is, $\sum_{Max(Q, w, d)} \frac{@(a, w, d)}{\#Max(Q, w, d)}$.

I propose, then, that the d-score of Q should be the credence-weighted average of the @(Q, w, d)'s, as determined by the credence function in d. In other words, I propose that:

$$S(Q,d) = \sum_{W} C(w)@(Q,w,d).$$

One can conceive of many alternative ways to d-score quantities, and, in a fuller discussion, it would be instructive to compare this proposal to its rivals. But this proposal is mathematically simple, metaethically simple, and, unlike many of its rivals,³⁰ rightly ensures that the d-score of actual

 $^{^{30}}$ For example, I have not yet been able to find any (remotely plausible) way of d-scoring quantities that (a) entails that V weakly D-dominates U and (b) does not entail that the D-score of V sometimes exceeds the d-score of actual value. A method of d-scoring quantities cannot be adequate unless in ensures that the d-score of actual value is never exceeded, so this amounts to an outstanding challenge to V-enthusiasts.

value is never exceeded.³¹ It therefore seems reasonable to me to explore this proposal and see what work it can do for us.

Having (provisionally) adopted this way of d-scoring quantities, we still cannot determine the score of any quantity, since I have not offered any proposal about what the M(d)'s are. But I am going to assume that we can determine some facts about how the scores of quantities relate, nevertheless, by appealing to relations of weak domination. If, for some d, $S(Q_1, d) > S(Q_2, d)$, and if, for every d, $S(Q_1, d) \geq S(Q_2, d)$, then Q_1 weakly D-dominates Q_2 . In what follows, I assume that a quantity always scores higher than does any quantity it weakly D-dominates.

5.2 Invariant and Supervenient Quantities

I now want to turn to guidance, building up to my preferred conception in stages.

One necessary condition for guidance is invariance. A quantity is d-invariant just if, for any worlds, w_1 and w_2 , $Max(Q, w_1, d) = Max(Q, w_2, d)$. If a quantity is d-invariant, then the options that maximize it relative to d depend only on the credences and utilities of an agent involved in d.

If quantity Q is d-invariant, let Max(Q, d) be the options that maximize Q relative to d: that is, the options that maximize Q at each $\langle w, d \rangle$. My proposed way to d-score quantities entails:

If Q is a d-invariant quantity, then the d-score of Q is the average of the U-values of the options that maximize Q relative to d.

To see this, suppose that Q is d-invariant. If k is the dependency hypothesis that holds at world w, then the actual value of option a at world w equals u(ak). Hence,

³¹Proof: Let α be the quantity maximized by exactly the actual value maximizing options at every $\langle w, d \rangle$. For any quantity Q and for any $\langle w, d \rangle$, $@(\alpha, w, d) \geq (Q, w, d)$, since the average actual value of the options that maximize actual value at $\langle w, d \rangle$ cannot be less than the average actual value of the options that maximize Q relative to $\langle w, d \rangle$. Hence, for any d, $S(\alpha, d) \geq S(Q, d)$. Hence, $S(\alpha) \geq S(Q)$.

$$S(Q,d) = \sum_{W} C(w)@(Q,w,d) = \sum_{K} C(k) \frac{\sum_{Max(Q,d)} u(ak)}{\#Max(Q,d)} = \frac{\sum_{Max(Q,d)} U(a)}{\#Max(Q,d)}.$$

The fact that the d-score of a d-invariant quantity is the average of the U-values of the options that maximize it is important because it entails that the d-score of a d-invariant quantity never exceeds the d-score of U.

Let's say that a quantity is *supervenient* if, for every d, it is d-invariant. It is often taken for granted that options can be made rationally permissible only by supervenient quantities. All of the quantities above—U, V, J, B, and G—are supervenient. It is therefore remarkable that, given the assumptions above, we can prove that U is the highest-scoring supervenient quantity. Just by appealing to relations of weak D-domination, we can prove that the highest-scoring supervenient quantities are only ever maximized by U-maximizing options. And by imposing a very plausible continuity constraint on the space of quantities, we can prove:

Supervenient Optimality: U is the highest-scoring (continuous) supervenient quantity.

Both the formulation of the continuity constraint and the proof of Supervenient Optimality are in the appendix.

To appreciate the metaethical import of Supervenient Optimality, suppose, just for a moment, that being supervenient is both necessary and sufficient for being d-guiding. Then, by combining Rational Optimization and Supervenient Optimality, we can give a metaethical derivation of U-monism. Three metaethical claims—Rational Optimization, my proposed way to d-score quantities, and the claim that being supervenient is both necessary and sufficient for being d-guiding—jointly entail that the rule for reducing rational permission maps every decision problem to U.

5.3 Guidance and Stable Maximization

I do not, myself, accept U-monism. But the attempted metaethical proof above is helpful because it allows me to say exactly where I think U-monism goes wrong.

In my view, being supervenient is necessary, but not sufficient, for being d-guiding. The attempted proof above establishes this much: that the rule for reducing rational permission maps d to U, whenever U is d-guiding. But it fails to establish U-monism because U is not always d-guiding.

According to the conception of guidance I favor:³²

An agent facing a decision is capable of being guided by some supervenient quantity Q just if, for some option a, the fact that a maximizes Q can be the agent's reason for choosing a.

And according to the conception of reasons for actions I favor:

The fact that a maximizes Q can be an agent's reason for choosing a just if (1) the agent is in a position to know that a maximizes Q and (2) conditional on a, the agent (still) is in a position to know that a maximizes Q.³³

Condition (1) is an epistemic constraint ensuring that reasons are within the agent's ken. If we make things simple and take knowledge to be truth plus certainty, then condition (1) says that p can be an agent's reason for choosing a only if, relative to the agent's credence function, p is true and certain. Condition (2) is a non-self-undermining constraint ensuring that agents can choose on the basis of their reasons.³⁴ If we again take knowledge to be truth plus certainty, then condition (2) says that p can be an agent's reason for choosing a only if, relative to the agent's credence function conditional on a, p is true and certain.³⁵

³²A similar conception of guidance is defended in [redacted].

³³Condition (2) is akin to, but not quite equivalent to, a principle that Hare (2011: 196) calls "Reasons are not Self-Undermining."

³⁴See [redacted].

³⁵Note an important distinction here. What condition (2) requires is that it be true and certain relative to C^a that a maximizes the quantity relative to C^a , not that it be true and certain relative to C^a that a maximizes the quantity relative to C. Thanks to [redacted] for discussion here.

Condition (1) is trivial for ideal agents (although, as we will see, non-trivial for nonideal agents). So, for ideal agents, the action lies entirely with condition (2), which is all about stability.

Say that option a stably maximizes Q relative to $d = \langle C, u, A, K \rangle$ just if a maximizes Q both relative to d and relative to $d^a = \langle C^a, u, A, K \rangle$. If my conception of guidance is correct, then guidance requires stable maximization. An ideal agent involved in d can be guided by a supervenient quantity Q just if some option stably maximizes Q relative to d. And since the fact that an option maximizes a quantity can be the agent's reason for choosing the option only if the option stably maximizes the quantity, it is the stable maximization, as opposed to the mere maximization, of the rational-making quantity that makes options rationally permissible.

Putting Rational Optimization together with my preferred account of d-guidance, we get:

Expanded Rational Optimization: What makes options rationally permissible relative to decision problem d is the stable maximization of the highest-scoring supervenient quantity that is stably maximized relative to d.

And with Expanded Rational Optimization in hand, we can shed some light on the puzzling examples from above.

5.4 Why *U*-monism Is Nearly True

Expanded Rational Optimization entails that an agent involved in d is capable of being guided by U just if some option stably maximizes U relative to d. This idea should sound familiar. Recall the distinction between stable and unstable decision problems. A decision problem is stable just if, relative to it, some option stably maximizes U. Thus, according to Expanded Rational Optimization, an agent is capable of being guided by U just if the agent is involved in some stable decision problem.

At an intuitive level, this prediction about guidance seems exactly right to me. An ideal agent facing a stable decision problem—Newcomb, say—can be guided by U. The fact that taking both boxes (uniquely) maximizes U can be the agent's reason for taking both boxes. But an ideal agent facing an unstable decision problem— $The\ Frustrater$, say—cannot be guided by U. In an unstable decision problem, U-maximization is too elusive to be a guide. The agent cannot both know which option she will choose and that she will choose a U-maximizing option.

By putting Expanded Rational Optimization and Supervenient Optimality together, we can explain the importance of the division between stable and unstable decision problems. For any decision problem d, it is the stable maximization of the highest-scoring d-guiding quantity that makes options rationally permissible. If d is a stable decision problem, then U is the highest-scoring d-guiding quantity. Hence, U is the rational-maker relative to every stable decision problem. But if d is an unstable decision problem, then U is not d-guiding, so the rational-maker will be some quantity other than U. Thus, U is the rational-maker relative to d if and only if d is a stable decision problem.

So let's turn to the next obvious question: which quantity or quantities are the rational-makers relative to unstable decision problems?

6 *U*-Pluralism

The form of rational pluralism that I develop in this section is speculative, but I think it compares favorably, in terms of theoretical simplicity, to the dependent monisms considered in §4. And, so far as I know, it is the only theory of rational choice on offer that can handle both Newcomb problems and the suite of unstable problems considered herein, a suite that includes Bostrom's *Meta-Newcomb*, Egan's *Psychopath Button*, and Ahmed's *Dicing with Death*. So, even if the theory proves mistaken, it still might help point us in the right direction.

$6.1 \quad UV$ -ism

As a warmup, consider, UV-ism, which says that the stable maximization of U makes options rationally permissible relative to stable decision problems and that the stable maximization of V makes options rationally permissible relative to unstable decision problems.

The success of UV-ism is strange, but striking. It verifies both K-Selection and K-Permission. It recommends two-boxing in Newcomb, the right-handed option in $The\ Demi-Semi-Frustrater$, the envelope in $The\ Frustrater$, and the left-handed options in $The\ Semi-Frustrater$. (It also correctly recommends: one-boxing in Bostrom's Meta-Newcomb, not pressing in Egan's $Psychopath\ Button$, and paying to flip in Ahmed's $Dicing\ with\ Death$.) And it correctly handles $Three\ Shells$, recommending shell A if the agent is highly confident that she will choose shell A, and recommending shells B and C if the agent is not highly confident that she will choose shell A. Nevertheless, I think that we should reject UV-ism, for two reasons.

The first is metaethical. Given Expanded Rational Optimization, UV-ism is tantamount to a bold metaethical prediction: that V is the highest-scoring quantity that can guide an agent whenever U fails to be. The presupposition of this prediction is correct. Whenever an option maximizes V, it also stably maximizes V. Therefore, for any decision problem d, V is d-guiding. But the substantive claim is dubious. In fact, I think it's false. I think that there are cases in which the highest-scoring quantity that can guide an ideal agent is neither U, nor V.

The second reason is related and extensional. There are cases—admittedly, rather complicated case—that challenge UV-ism. The cases are both unstable and Newcomb-like. Here's an example:

The Meta-Frustrater. There are two opaque boxes, one white and one black. The agent has four options. She can point to either box with either hand $(a_{RW}, a_{LW}, a_{RB}, \text{ or } a_{LB})$. One of the boxes contains \$0; the other contains \$100. The agent receives the con-

tents of whichever box she points to. Which box contains which sum depends on a prediction made yesterday by a minion who seeks to frustrate. If the minion predicted that the agent would point to the white box, the black box contains \$100. If the minion predicted that the agent would point to the black box, the white box contains \$100. There are two left-right asymmetries. The first is straightforward. The agent receives an extra \$5 if she points to a box with her right hand. The second is more complicated. There are two minions: one is a 90%-reliable predictor of both left-handed and right-handed box-pointings, and the other is a 50%-reliable predictor of both left-handed and right-handed box-pointings. Which minion is up against the agent depends on a prediction made two days ago by the Meta-Frustrater, who is a very reliable predictor. If the Meta-Frustrater predicted that the agent would point with her right hand, then the Meta-Frustrater put the agent up against the minion who is 90% reliable. If the Meta-Frustrater predicted that the agent would point with her left hand, then the Meta-Frustrater put the agent up against the minion who is 50% reliable. The agent knows all of this.

The similarities between *The Semi-Frustrater* and *The Meta-Frustrater* are obvious. If the Meta-Frustrater is a (nearly) perfect predictor, then, in the two examples, the four options have (nearly) the same V-values and U-values.³⁶ According to UV-ism, the examples are also normatively alike. In

$$V(a_{RW}) = V(a_{RB}) = (105)(0.1) + (5)(0.9) = 15$$
; and $V(a_{LW}) = V(a_{LB}) = (100)(0.5) + (0)(0.5) = 50$.

The *U*-values are sensitive to the agent's credences over *A*. If, for example, $C(a_{RW}) = C(a_{RB}) = C(a_{LW}) = C(a_{RW}) = 0.25$, then:

$$U(a_{RW}) = U(a_{RB}) = (105)(0.5) + (5)(0.5) = 55$$
; and

$$U(a_{LW}) = U(a_{LB}) = (100)(0.5) + (0)(0.5) = 50.$$

³⁶If the Meta-Frustrater is perfect, then in both examples:

both examples, UV-ism recommends the left-handed options.

Intuitions about examples this complicated are not always clear. But it seems to me, and it has seemed to quite a number of other people, too, that UV-ism mishandles The Meta-Frustrater: that, on account of something Newcomb-like, the rationally permissible options in The Meta-Frustrater are the right-handed options. In The Semi-Frustrater, the predictor who seeks to frustrate the agent has a predictive weakness, which the agent can exploit by pressing left-handedly. But in The Meta-Frustrater, the predictor who seeks to frustrate the agent—whichever minion it happens to be—has no predictive weakness to exploit. An agent who points left-handedly in The Meta-Frustrater thus seems to be merely managing the news:³⁷ forgoing a certain benefit in order to produce evidence that she is up against the predictively weaker minion.

6.2 *U*-pluralism

The form of rational pluralism that I favor explains why UV-ism has the success it does, better coheres with Expanded Rational Optimization, and correctly handles $The\ Meta$ -Frustrater. I call it, U-pluralism.

Recall that U is defined in terms of K, the set of dependency hypotheses. To formulate U-pluralism, I will assume that there is some privileged way to gradually coarsen K. The set of these coarsenings, K, is linearly ordered by granularity. The least member of K is the set of dependency hypotheses, to which I will append superscripted zeroes, $K^0 = \{k_1^0, k_2^0, ..., k_n^0\}$. As we gradually coarsen, we might arrive at some intermediate partition, $K^j = \{k_1^j, k_2^j, ..., k_m^j\}$, and then some coarser intermediate partition, $K^l = \{k_1^l, k_2^l, ..., k_k^l\}$. The coarsest and greatest member of K is the trivial partition, $K^{\top} = \{k_1^{\top}, k_2^{\top}, ..., k_k^{\top}\}$, which has $k^{\top} = \top$ as its only member.

Just as U-monists can disagree about how best to conceptualize K, U-pluralists can disagree about how best to conceptualize K.³⁸ In the interest

³⁷ Cf. Lewis (1981).

³⁸See e.g. Ahmed (2014a), Joyce (1999), and Lewis (1981).

of efficiency, I will help myself to one plausible conception.

Let $\{lh_1, lh_2, ..., lh_n\}$ be the set of propositions that specify the laws of nature and the history of the world up to the time of decision, insofar as those matters are beyond the agent's control. If we take $\{lh_1, lh_2, ..., lh_n\}$ to be the set of dependency hypotheses, then the natural way to gradually coarsen is by removing successive slices of history, producing ever shorter initial segments, and then finally removing the laws, themselves. (This is not the only possible conception of K, 39 but it has the virtue of being easy to work with.)

However K is characterized, it gives rise to a spectrum of supervenient quantities, which range from U, at one extreme, to V, at the other:

$$U^{0}(a) =_{def} \sum_{K^{0}} C(k^{0}) V(ak^{0}) = \sum_{K^{0}} C(k^{0}) u(ak^{0}) = U(a);$$
...
$$U^{j}(a) =_{def} \sum_{K^{j}} C(k^{j}) V(ak^{j});$$
...
$$U^{\top}(a) =_{def} \sum_{K^{\top}} C(k^{\top}) V(ak^{\top}) = V(a^{\top}) = V(a).$$

Let U be the set of these quantities, and let U inherit the order on K: $U^i \prec U^j$ if and only if K^i is a refinement of K^j . Then the least member of U is the causally finest member, namely, U. The greatest member is the causally coarsest member, namely, V. And the intermediate members have intermediate degrees of causal granularity and are ordered accordingly.

With U characterized, we can state U-pluralism:

U-pluralism: What makes an option rationally permissible relative to decision problem d is the stable maximization of the least member of U that is stably maximized relative to d.

³⁹One alternative I find attractive is purely modal. Each fact is assigned some counterfactual fixity, à la Kment (2014), and we gradually coarsen by progressively removing the facts with the least counterfactual fixity. This purely modal characterization of K is harder to work with, but probably superior.

To derive U-pluralism from Expanded Rational Optimization, we would need two claims. First, we would need the claim that U is linearly ordered by score: that every member of U scores higher than does every greater member. Second, we would need the claim that U is exhaustive: that the rule for reducing rational permission maps every decision problem to some member of U.

The first claim is supported by a computer simulation. As I said above, with enough constraints on the space of decision problems, it can be plausible to use an indifference measure to average the d-scores of a quantity. In the simulation run, I did precisely that. The simulation involved 16 dependency hypotheses and four options. For each cycle, probabilities and utilities (between 0 and 100) were randomly distributed over the 64 atoms in $A \times K$. Since there are 16 dependency hypotheses, there are five members of U: $U^0 = U$, U^1 , U^2 , U^3 , and $U^{\top} = V$. Each quantity in U is supervenient, and the d-score of a supervenient quantity is the average U-value of the options that maximize the quantity relative to d, so, in each cycle of the simulation, for each $U^i \in U$, I recorded the U-value of the U^i -maximizing option. After running the simulation 15,000 times, the following curve emerged:

Normalizing the average d-score of U to 1, the average d-scores of U^1 , U^2 , U^3 , and U^{\top} were, respectively, .9778, .9685, .9651, and .9642. Now, of course, this does not *prove* that each member of U is higher-scoring than every greater member, but it does lend considerable support to that claim.

The second claim is speculative, but more plausible in light of the first. We know that the rule for reducing rational permission maps every stable decision problem to U. And since V is d-guiding for every decision problem d, we know that the rule for reducing rational permission never maps a decision problem to a quantity that scores lower than V. Therefore, if the first claim is true—if U is linearly ordered by score—then the second claim is plausible, although not by any means obvious or trivial.

One important virtue of U-pluralism is that it delivers the recommen-

dations we seek. It verifies both K-Selection and K-Permission. It recommends two-boxing in Newcomb, the right-handed option in $The\ Demi-Semi-Frustrater$, the envelope in $The\ Frustrater$, and the left-handed options in $The\ Semi-Frustrater$. (It recommends: one-boxing in Bostrom's Meta-Newcomb, not pressing in Egan's $Psychopath\ Button$, and paying to flip in Ahmed's $Dicing\ with\ Death$.) It correctly handles $Three\ Shells$, recommending shell A if the agent is highly confident that she will choose shell A, and recommending shells B and C if the agent is not highly confident that she will choose shell A. And it also correctly handles $The\ Meta-Frustrater$, recommending the right-handed options.

I will not go through all of the relevant calculations. But let me go through two cases, *The Frustrater* and *The Meta-Frustrater*.

To make things simple, suppose that the agent facing *The Frustrater* is certain that the Frustrater's prediction was made instantaneously j units prior to the time of decision. (This assumption merely vivifies the metaethical structure.⁴⁰) Let the k^j 's be propositions that specify the laws of nature and the history of the world up to j units prior to the time of decision, and let U^j be defined in terms of K^j . As we work our way through the members of U, from least to greatest, we encounter a metaethical shift.

⁴⁰Assuming that the agent is certain that the prediction was made instantaneously j units prior to the time of decision makes the metaethical transition sudden. For every $U^i \prec U^j$, $U^i(a_A) + U^i(a_B) = 100$. And, for every $U^k \succeq U^j$, $U^k(a_A) + U^k(a_B) \approx 0$. If we drop the assumption that the agent is certain that the prediction was made j units prior to the time of decision, the metaethical transition might instead be gradual. If the decrease is gradual, then emphasizing stable maximization might be important. In a version of $tallow{The Frustrater}$ in which the agent is uncertain when the prediction was made, it may be the case that the least member of $tallow{U}$ that is stably maximized, say, $tallow{U}$, is maximized both by, say, $tallow{U}$, and $tallow{U}$. This sort of co-maximization would not make $tallow{U}$, and $tallow{U}$, is maximized by maximize any member of $tallow{U}$. In fact, neither $tallow{U}$, nor $tallow{U}$ as stably maximize any member of $tallow{U}$. If $tallow{U}$ maximizes some $tallow{U}$, then $tallow{U}$ can be $tallow{U}$ as evidence in favor of $tallow{U}$ and $tallow{U}$ are the agent then will regard $tallow{U}$ as evidence in favor of $tallow{U}$ stably maximizes the least member of $tallow{U}$ that is stably maximized, whatever that proves to be.

If $U^i \prec U^j$ —in other words, if K^i is any refinement of K^j —then, depending on the agent's credences, a_A and/or a_B maximizes U^i , but no option stably maximizes U^i . Each k^i specifies how much money is in each box. As a result, for any k^i , $V(a_E k^i) = 40$ and $V(a_a k^i) + V(a_b k^i) = 100$. Hence, $U^i(a_E) = 40$ and $U^i(a_A) + U^i(a_B) = 100$. But no option stably maximizes U^i because the agent regards choosing a box as strong evidence that the box is empty:

$$\sum_{K^{i}} C(k^{i}|a_{A}) V(a_{A}k^{i}) < \sum_{K^{i}} C(k^{i}|a_{A}) V(a_{B}k^{i}), \text{ and}$$

$$\sum_{K^{i}} C(k^{i}|a_{B}) V(a_{B}k^{i}) < \sum_{K^{i}} C(k^{i}|a_{B}) V(a_{A}k^{i}).$$

By contrast, if $U^j \preceq U^k$, then taking the envelope stably maximizes U^k . No k^k specifies how much money is in each box. As a result, $V(a_Ak^k)$, $V(a_Bk^k)$, and $V(a_Ek^k)$ are determined by what k^k says about the Frustrater's reliability. If the agent is certain that the Frustrater is almost perfectly reliable, then, for any k^k , $V(a_Ek^k) = 40$ and $V(a_Ak^k) = V(a_Bk^k) \approx 0$. Hence, $U^k(a_Ek^k) = 40$ and $U^k(a_Ak^k) = U^k(a_Bk^k) \approx 0$. And choosing the envelope also stably maximizes U^k . Indeed, if the agent has no uncertainty about the Frustrater's predictive powers, then, for any k^k , $C(k^k|a_A) = C(k^k|a_B) = C(k^k|a_E)$. Thus, according to U-pluralism, an agent facing The Frustrater is rationally required to choose the envelope, and is so because choosing the envelope is the only option that stably maximizes the least member of U that is stably maximized, namely, U^j .

When we see how U-pluralism handles The Frustrater, we can understand why UV-ism has success.

Think about stable decision problems. What makes options rationally permissible relative to a stable decision problem is the stable maximization of U. But in the simplest stable decision problems, V-maximization and U-maximization coincide. Thus, although V-monism is mistaken metaethically, it very often delivers the correct recommendations. The only stable decision problems in which V-monism gives the wrong recommendations are Newcomb problems.

A similar thing holds true of unstable decision problems. What makes options rationally permissible relative to an unstable decision problem is the stable maximization of the least member of U that is stably maximized. But in the simplest unstable decision problems, V-maximization coincides with the maximization of the least member of U that is stably maximized. Thus, although UV-ism is mistaken metaethically, it very often delivers the right recommendations. The only unstable decision problems in which UV-ism delivers the wrong recommendations are unstable Newcomb problems, like $The\ Meta-Frustrater$.

Let's now consider The Meta-Frustrater. To make things simple, suppose that the agent is certain that the Meta-Frustrater's prediction was made instantaneously l units of time prior to the decision, and that the agent is certain that the minion, whichever one the agent is up against, made their prediction instantaneously j units of time prior to the decision, j < l. (Again, these assumptions merely serve to vivify the metaethical structure.⁴¹) Let each k^l specify the laws and the history of the world up to l units prior to the decision; let each k^j specify the laws and history up to j units prior to the decision; and let U^l and U^j be defined in terms of K^l and K^j , respectively. As we work our way through the members of U, from least to greatest, we encounter two metaethical shifts.

If $U^i \prec U^j$, then, depending on the agent's credences, a_{RW} and/or a_{LW} maximize U^i , but no option stably maximizes U^i . Each k^i specifies which box contains \$100. As a result, for any k^i :

$$V(a_{RW}k^i) = 5 + V(a_{LW}k^i);$$

 $[\]overline{}^{41}$ There is one added complication. As [redacted] pointed out to me, according to U-pluralism as formulated, it is essential that the Meta-Frustrater makes his prediction before the minions do. If the minions make their prediction first, then the options that stably maximize the least member of U that is stably maximized will be the left-handed options. I am not sure whether this prediction is wrong. (Flipping the temporal order makes my intuitions less clear.) But, when I am inclined to think that flipping the temporal order makes no normative difference, I am inclined, not to abandon U-pluralism, but to adopt an alternative conception of K. See note 39.

$$V(a_{RB}k^{i}) = 5 + V(a_{LB}k^{i});$$
 and $V(a_{RW}k^{i}) + V(a_{RB}k^{i}) = 110.$

This ensures that choosing a_{RW} and/or a_{RB} maximize U^i . But no option stably maximizes U^i because the agent regards pointing right-handedly to a box as strong evidence that the box is empty:

$$\begin{split} & \sum_{K^i} C(k^i | a_{RW}) V(a_{RW} k^i) < \sum_{K^i} C(k^i | a_{RW}) V(a_{RB} k^i), \text{ and} \\ & \sum_{K^i} C(k^i | a_{RB}) V(a_{RB} k^i) < \sum_{K^i} C(k^i | a_{RB}) V(a_{RW} k^i). \end{split}$$

By contrast, if $U^j \leq U^k$ and $U^k \prec U^l$, then the two right-handed options both stably maximize U^k . The k^k 's do not specify which box contains \$100, but they do specify which minion the agent is up against. If k^k says that the agent is up against the 50% reliable minion, then:

$$V(a_{RW}k^k) = V(a_{RB}k^k) = (0.5)(105) + (0.5)(5) = 55$$
, and $V(a_{LW}k^k) = V(a_{LB}k^k) = (0.5)(100) + (0.5)(0) = 50$.

If k^k says that the agent is up against the 90% reliable minion, then:

$$V(a_{RW}k^k) = V(a_{RB}k^k) = (0.1)(105) + (0.8)(5) = 15$$
, and
$$V(a_{LW}k^k) = V(a_{LB}k^k) = (0.1)(100) + (0.9)(0) = 10.$$

Hence, no matter how the agent distributes her credence over K^k , the two right-handed options maximize U^k .

The Newcomb-like phenomenon in *The Meta-Frustrater* is thus made apparent. In *Newcomb*, taking both boxes stably maximizes U, even though the agent regards taking both boxes as bad news, and this is because, no matter how the agent distributes her credence over K, taking both boxes maximizes U. In *The Meta-Frustrater*, for any U^k , $U^j \leq U^k \prec U^l$, the two right-handed options stably maximize U^k , even though the agent regards the right-handed options as bad news, and this is because, no matter how the agent distributes her credence over K^k , the two right-handed options maximize U^k .

Now, if we *continue* working through the members of U, from least to greatest, we will encounter another metaethical shift. If $U^l \prec U^m$, then the two left-handed options stably maximize U^m . This is because the k^m 's specify neither which box contains \$100, nor which minion the agent is up against. Thus, as we know, the two left-handed options maximize V.

But the quantities stably maximized by the left-handed options are metaethically irrelevant. What makes options rationally permissible in *The Meta-Frustrater* is the stable maximization of U^j , the least member of U that is stably maximized. So the rationally permissible options are the right-handed options

Now, as I said, U-pluralism is speculative. But it has at least three things going for it. One: it coheres nicely with Expanded Rational Optimization. Two: it is, so far as I know, the only theory on offer that can handle both Newcomb problems and the full suite of unstable problems considered herein. And three: it is metaethically conservative. According to U-pluralism, expected value theory is nearly true! The stable maximization of U, i.e., expected value, is almost always what makes options rationally permissible. The other members of U, and the additional complications that they bring in tow, are relevant only when we turn our attention to highly unusual cases.

7 Conclusion

An adequate consequentialist reduction of objective permission must involve both an identifying element and an explanatory element. For each decision problem d, we need to identify the quantity that is the objective-maker relative to d, and we need to explain why that quantity is the objective-maker relative to d. As it turns out, both elements are easy. Actual value is the universal objective-maker, and it is so because it is the best quantity to maximize.

An adequate consequentialist reduction of rational permission likewise

must involve both an identifying element and an explanatory element. For each decision problem d, we need to identify the quantity that is the rational-maker relative to d, and we need to explain why that quantity is the rational-maker relative to d. But both elements here are more difficult.

The identifying element is more difficult because there's not just one answer. Many who work on rational choice are searching for the universal rational-maker. But, if I'm right, then rational monism is false: there is no universal rational-maker. And that fact makes the explanatory element more difficult, too; for the usual ways of trying to explain why a quantity is a rational-maker presuppose and require that the quantity be a universal rational-maker. There is no received theory of occasional rational-making.

The positive part of this essay began with me offering a theory of occasional rational-making. I suggested that what makes a quantity a rational-maker relative to decision problem d is being the best quantity to maximize, among the quantities that can guide an agent involved in d. Stated in this way, the proposal is skeletal. But I tried to put some meat on the bones by offering a partial account of how to score quantities and an account of what it is for a quantity to be d-guiding.

When the theory of occasional rational-making is fleshed out in the ways I prefer, it predicts that U is the rational-maker relative to all and only the stable decision problems. To my mind, this brings some needed clarity to the theory of rational choice for ideal agents; for it allows us to reconcile the pro-U intuitions in Newcomb problems with the anti-U intuitions in unstable problems.

But the real value of rational pluralism and the theory of occasional rational-making that gives rise to it become apparent only when we zoom out and consider nonideal agents.

Rational monism, defended in full generality, is very doubtful because we can *minerize* virtually any quantity, constructing a case that stands to it as *Boxes like Miners* stands to actual value. For example, I am convinced that what makes choosing the middle box rationally permissible in *Boxes like*

Miners is the (stable) maximization of U, and the following case, in which there is nonideal agent with unlimited of powers of introspection but limited powers of deduction, minerizes U:⁴²

The Fire. The fire alarm rings and the agent, a firefighter, hurries onto the truck. On the ride over she deliberates. There are three doors into the building, arranged left-to-right. The agent, who cares only about saving lives, must enter the building via one of the three doors. Since she does not know the exact distribution of residents in the building, she does not know which option will result in the most rescues. Based on her credences about the distribution of residents, she calculates the U-value of each option and writes the value on a note card. After exiting the truck and attaching the water hose, she races toward the burning building. She reaches into her pocket, but the note card is gone. Time is of the essence! She knows that all of the residents will die in the time it would take her to recalculate the U-values of the options. She knows that the current U-values of the options are what they were when she calculated them, since she knows that her credences about the distribution of residents are unchanged. But she cannot fully remember the results of her calculations. She remembers that the U-value of entering through the middle door is 9. Of the other two options, she remembers that one has an U-value of 0 and that the other has an U-value of 10. But she cannot remember which option has which U-value, and divides her credence equally between the two possibilities. (In fact, entering through the right door has an U-value of 10, as the lost note card attests.)

The agent facing *The Fire* is rationally required to enter via the middle door, even though she knows for certain that doing so does not maximize

⁴²This example, from [redacted], adapts an example from Kagan (2018). For related discussion, see *e.g.* Feldman (2006) and Weirich (2004).

U. Therefore, if the maximization of U is what makes options rationally permissible in *Boxes like Miners*, rational monism, as applied to both ideal and nonideal agents, stands refuted. And notice that we can replace U with any minerizable quantity. Rational monism, defended in full generality, is hopeless if the quantity at its center is minerizable.

Like others who have discussed structurally similar cases,⁴³ I think that, in *The Fire*, options are made rationally permissible by maximizing U_2 , where the U_2 -value of option a is the agent's expectation of its U-value—that is, $\sum_v C([U(a) = v])v$. The U_2 -value of entering through the middle door is 9, and the U_2 -values of entering through the left door and entering through the right door both are 5.

But if the maximization of U_2 is what makes options rationally permissible in *The Fire*, then we need to explain why it is the maximization of U_2 , specifically, and not some other quantity, that makes options rationally permissible. And the constrained optimization approach is appealing. The firefighter is *not* capable of being guided by U. Condition (1) fails: the firefighter, on account of her limited powers of deduction, is not in a position to know that entering through the right door (uniquely) maximizes U. But the firefighter is capable of being guided by U_2 . And it's tempting to think that what sets U_2 apart from the other quantities that can guide the firefighter, what makes U_2 the occasional rational-maker, is its status as being the best quantity that can guide the firefighter, i.e., the quantity that best conduces to the realization of value, among the quantities that can guide the firefighter.⁴⁴

In this essay I have focused almost exclusively on ideal agents because

⁴³See e.g. [redacted] and Weirich (2004).

⁴⁴Much work on bounded rationality is similarly animated by a constrained optimization conception of rationality. See *e.g.* Bossaerts and Murawksi (2017), Gigerenzer (2008), Gigerenzer and Selten (2001), Griffiths *et al.* (2015), Griffiths and Tenebaum (2006), Halpern *et al.* (2014), Icard (2018), Lokowski and Kreinovich (2018), Paul and Quiggin (2018), Russell and Subramanian (1995), Simon (1956; 1957; 1983), Vul *et al.* (2014), and Weirich (1988; 2004).

nonideal agents are hard to model with mathematical precision. But consequentialists ultimately need a reduction of rational permission that holds true for all agents, ideal and nonideal. And I think that rational pluralism and the theory of occasional rational-making that gives rise to it are worthy of further exploration because I think that they may help unify the theory of rational choice for ideal agents and the theory of rational choice for nonideal agents, thus helping pave the way for a fully general consequentialist reduction of rational permission.⁴⁵

A Proof of Supervenient Optimality

The proof of **Supervenient Optimality** has two parts. First, we show that U weakly D-dominates any supervenient quantity that diverges from U. Then we show that any quantity that is distinct from U, but does not diverge from U, violates a plausible continuity constraint.

If $Max(Q, w, d) \not\subset Max(U, w, d)$ for any point $\langle w, d \rangle$, I will say that there is a point of divergence between Q and U. If Q is supervenient and there is a point of divergence between Q and U, then U weakly D-dominates Q. After all, suppose that $\langle w, d \rangle$ is a point of divergence between Q and U. Since Q and U are both supervenient, the d-score of Q is the average of the U-values of the options in Max(Q, w, d), and the d-score of U is the average of the U-values of the options in Max(U, w, d). Hence, since at least one member of Max(Q, w, d) fails to maximize U relative to d, the average of the U-values of the options in Max(Q, w, d) is strictly less than the average of the U-values of the options in Max(U, w, d). Hence, S(Q, d) < S(U, d). Moreover, if Q is supervenient, then, for any d, $S(Q, d) \leq S(U, d)$. So it follows that U weakly D-dominates Q. And given my assumption that the ordinal rankings of quantities respect relations of weak D-domination, it follows that U scores higher than does Q.

The supervenient quantities that are distinct from U, but score as highly

 $^{^{45}}$ [Acknowledgements].

as U, are subset quantities: quantities that are always maximized by U-maximizing options, but not always maximized by every U-maximizing option. (Think, for example, about the quantity that corresponds to being the leftmost U-maximizing option.) But subset quantities violate an intuitively plausible continuity constraint. If u is a utility function, and u(w) = x, then let $u^{w,\epsilon}$ and $u^{w,-\epsilon}$ be utility functions that are exactly like u, except that $u^{w,\epsilon}(w) = x + \epsilon$ and $u^{w,-\epsilon}(w) = x - \epsilon$. If $d = \langle C, u, A, K \rangle$, then let $d^{w,\epsilon} = \langle C, u^{w,\epsilon}, A, K \rangle$ and let $d^{w,-\epsilon} = \langle C, u^{w,-\epsilon}, A, K \rangle$. The relevant continuity constraint then can be stated as follows:

Utility Continuity. If $a \notin Max(Q, w, d)$, then, for any world w_i there is some ϵ such that $a \notin Max(Q, w, d^{w_i, \epsilon})$ and $a \notin Max(Q, w, d^{w_i, -\epsilon})$.

In effect, Utility Continuity says that small changes to utilities assigned to any particular world should precipitate only small changes in the values that a quantity assigns to options.

To see that every subset quantity violates Utility Continuity, suppose that Q is a subset quantity, and suppose that a is among the options that maximize U at $\langle w, d \rangle$, but not among the options that maximize Q at $\langle w, d \rangle$. There will then be some a-world, w_i , to which the credence function in d assigns nonzero probability, which is such that, increasing its utility, while keeping the utility of every other world the same, increases the U-value of a, but does not increase the U-value of any other option in A. So, for any ϵ , a uniquely maximizes U at $\langle w, d^{w_i, \epsilon} \rangle$. Since Q is a subset quantity, a also uniquely maximizes Q at $\langle w, d^{w_i, \epsilon} \rangle$. But that shows that Q violates Utility Continuity.

Thus, Supervenient Optimality holds: U is the highest-scoring supervenient quantity that satisfies Utility Continuity.

References

[1] [redacted]

- [2] Arif Ahmed. 2012. "Press the Button." Philosophy of Science 79: 386–95.
- [3] —. 2014a. Evidence, Decision and Causality. Cambridge University Press.
- [4] —. 2014b. "Dicing with Death." Analysis 74: 587–94.
- [5] —. 2018. "Why Ain'cha Rich?" In A. Ahmed (ed.) *Newcomb's Problem*. Cambridge University Press.
- [6] Frank Arntzenius. 2008. "No Regrets, or: Edith Piaf Revamps Decision Theory." *Erkenntnis* 68: 277–97.
- [7] Adam Bales. 2018. "Decision-Theoretic Pluralism." *Philosophical Quarterly* 68: 801–18.
- [8] Robert Bassett. 2015. "A Critique of Benchmark Theory." Synthese 192: 241–67.
- [9] Jeremy Bentham. 1961[1789]. An Introduction to the Principles of Morals and Legislation. Garden City: Doubleday.
- [10] José Luis Bermúdez. 2009. Decision Theory and Rationality. Oxford University Press.
- [11] Ethan D. Bolker. 1967. "A Simultaneous Axiomatisation of Utility and Subjective Probability." *Philosophy of Science* 34: 333–40.
- [12] Peter Bossaerts and Carsten Murawski. 2017. "Computational Compextity and Human Decision-Making." *Trends in Cognitive Sciences* 21: 917–29.
- [13] Nick Bostrom. 2001. "The Meta-Newcomb Problem." Analysis 61: 309–10.
- [14] R. A. Briggs. 2010. "Decision-Theoretic Paradoxes as Voting Paradoxes." *Philosophical Review* 119: 1–30.

- [15] Lara Buchak. 2013. Risk and Rationality. Oxford University Press.
- [16] Ellery Eells. 1982. Rational Decision and Causality. Cambridge University Press.
- [17] Ellery Eells and William Harper. 1991. "Ratifiability, Game Theory, and the Principle of Independence of Irrelevant Alternatives." *Australasian Journal of Philosophy* 69: 1–19.
- [18] Andy Egan. 2007. "Some Counterexamples to Causal Decision Theory." *Philosophical Review* 116: 94–114.
- [19] Lina Eriksson and Alan Háyek. 2007. "What Are Degrees of Belief?" Studia Logica 86: 183–213.
- [20] Fred Feldman. 2006. "Actual Utility, The Objection from Impracticality, and the Move to Expected Utility." *Philosophical Studies* 129: 49–79.
- [21] J. Dmitri Gallow. MS. "The Causal Decision Theorist's Guide to Managing the News."
- [22] Allan Gibbard and William Harper. 1978. "Counterfactuals and Two Kinds of Expected Utility." In A. Hooker, J. J. Leach, and E. F. McClennen (eds.), Foundations and Applications of Decision Theory, 125–62. Reidel.
- [23] Gerd Gigerenzer. 2008. Rationality for Mortals: How People Cope with Uncertainty. Oxford University Press.
- [24] Gerd Gigerenzer and Reinhard Selten (eds). 2001. Bounded Rationality: The Adaptive Toolkit. Cambridge, MA: MIT Press.
- [25] Thomas L. Griffiths, Falk Lieder, and Noah D. Goodman. 2015. "Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and Algorithmic." *Topics in Cognitive Science* 7: 217–29.
- [26] Thomas L. Griffiths and Joshua B. Tenenbaum. 2006. "Optimal Predictions in Everyday Cognition." *Psychological Science* 17: 767–73.

- [27] Johan Gustafsson. 2011. "A Note in Defense of Ratificationism." *Erkenntnis* 75: 147–50.
- [28] Joseph Y. Halpern, Rafael Pass, and Lior Seeman. 2014. "Decision Theory with Resource-Bounded Agents," *Topics in Cognitive Science* 6: 245–57.
- [29] Peter J. Hammond. 1988. "Consequentialist Foundations for Expected Utility Theory." *Theory and Decision* 25: 25–78.
- [30] Caspar Hare. 2011. "Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had not Done What You Did." Noûs 45: 190–206.
- [31] Caspar Hare and Brian Hedden. 2016. "Self-Reinforcing and Self-Frustrating Decisions." *Noûs* 50: 604–28.
- [32] William Harper. 1986. "Mixed Strategies and Ratifiability in Causal Decision Theory." *Erkenntnis* 24: 25–36.
- [33] Terrence Horgan. 1981. "Counterfactuals and Newcomb's Problem." Journal of Philosophy 78: 331–56.
- [34] Daniel Hunter and Reed Richter. 1978. "Counterfactuals and Newcomb's Paradox." Synthese 39: 249–61.
- [35] Thomas Icard. 2018. "Bayes, Bounds, and Rational Analysis." *Philoso-phy of Science* 85: 79–101.
- [36] Richard Jeffrey. 1965. The Logic of Decision. University of Chicago Press.
- [37] —. 1983. The Logic of Decision, 2nd ed. University of Chicago Press.
- [38] James Joyce. 1999. The Foundations of Causal Decision Theory. Cambridge University Press.

- [39] —. 2012. "Regret and Stability in Causal Decision Theory." Synthese 187: 123–45.
- [40] —. 2018. "Deliberation and Stability in Newcomb Problems and Psuedo-Newcomb Problems." In A. Ahmed (ed.), *Newcomb's Problem*. Oxford University Press.
- [41] Shelly Kagan. 2018 "The Paradox of Methods." *Politics, Philosophy, and Economics* 17: 148–68.
- [42] Boris Kment. 2014. *Modality and Explanatory Reasoning*. Oxford University Press.
- [43] Niko Kolodny and John MacFarlane. 2010. "Ifs and Oughts." *Journal of Philosophy* 107: 115–43.
- [44] David Lewis. 1981. "Causal Decision Theory." Australasian Journal of Philosophy 59: 5–30.
- [45] Christian List and Franz Dietrich. 2016. "Mentalism Versus Behaviorism in Economics: a Philosophy-of-Science Perspective." Economics and Philosophy 32: 249–81.
- [46] Joe Lorkowski and Vladik Kreinovich. 2018. Bounded Rationality in Decision Uncertainty: Towards Optimal Granularity. Springer.
- [47] Christopher J. G. Meacham and Jonathan Weisberg. 2011. "Representation Theorems and the Foundations of Decision Theory." *Australasian Journal of Philosophy* 89: 641–63.
- [48] John Stuart Mill. 1988[1861]. *Utilitarianism*, Roger Crisp (ed.). Oxford University Press.
- [49] G. E. Moore. 1903. *Principia Ethica*. Cambridge University Press.
- [50] —. 1912. Ethics. Oxford University Press.

- [51] John von Neumann and Oskar Morgenstern. 1944. Theory of Games and Economic Behavior. Princeton University Press.
- [52] Robert Nozick. 1969. "Newcomb's Problem and Two Principles of Choice." In N. Rescher (ed.), Essays in Honor of Carl. G. Hempel, 114–46. Reidel.
- [53] Graham Oddie and Peter Menzies. 1992. "An Objectivist's Guide to Subjective Value." *Ethics* 102: 512–33.
- [54] L. A. Paul and John Quiggin. 2018. "Real World Problems." *Episteme* 15: 363–82.
- [55] Richard Pettigrew. 2015. "Risk, Rationality, and Expected Utility Theory." Canadian Journal of Philosophy 47: 798–826.
- [56] —. 2019. Choosing for Changing Selves. Oxford University Press.
- [57] Wlodek Rabinowicz. 1988. "Ratifiability and Stability." In P. Gärdenfors and N. Sahlin (eds.), *Decision, Probability, and Utility*, 406–25. Cambridge University Press.
- [58] —. 1989. "Stable and Retrievable Options." *Philosophy of Science* 56: 624–41.
- [59] Frank P. Ramsey. 1990[1926]. "Truth and Probability." In D. H. Mellor (ed.), *Philosophical Papers*. Cambridge University Press.
- [60] John Rawls. 1971. A Theory of Justice. Harvard University Press.
- [61] Pamela Robinson. Dissertation. "Toward an Ultimate Normative Theory."
- [62] Stuart J. Russell and Devika Subramanian. 1995. "Provably Bounded-Optimal Agents." *Journal of Artificial Intelligence Research* 2: 575-609.
- [63] Leonard J. Savage. 1954. The Foundations of Statistics. Wiley.

- [64] Amartya Sen. 1970. Collective Choice and Social Welfare. Holden-Day.
- [65] Herbert A. Simon. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63: 129–38.
- [66] —. 1957. *Models of Man.* Wiley.
- [67] —. 1983. Reason in Human Affairs. Stanford University Press.
- [68] Brian Skyrms. 1982. "Causal Decision Theory." *Journal of Philosophy* 79: 695–711.
- [69] —. 1984. Pragmatics and Empiricism. Yale University Press.
- [70] —. 1990. The Dynamics of Rational Deliberation. Harvard University Press.
- [71] J. Howard Sobel. 1994. Taking Chances: Essays on Rational Choice. Cambridge University Press.
- [72] Robert Stalnaker. 1981. "Letter to David Lewis of 21 May 1972." In Stalnaker, Harper, and Pearce (eds.), *Ifs: Conditionals, Belief, Decision, Chance and Time*, 151–53. Reidel.
- [73] Judith Jarvis Thomson. 2008. Normativity. Open Court.
- [74] Edward Vul, Noah D. Goodman, Thomas L. Griffiths, and Joshua B. Tenenbaum. 2014. "One and Done? Optimal Decisions from Very Few Samples." Cognitive Science. 38: 599-637.
- [75] Ralph Wedgwood. 2013. "Gandalf's Solution to the Newcomb Problem." Synthese 190: 2643–75.
- [76] Paul Weirich. 1985. "Decision Instability." Australasian Journal of Philosophy 63: 465–72.
- [77] —. 1988. "Hierarchical Maximization of Two Kinds of Expected Utility. *Philosophy of Science* 55: 560–82.

- [78] —. 2004. Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances. Oxford University Press.
- [79] Ian Wells. Forthcoming. "Equal Opportunity and Newcomb's Problem." Mind.