Expected Value Under Normative Uncertainty

September 2019[1]

Franz Dietrich
Paris School of Economics & CNRS

Brian Jabarian
Paris School of Economics & Princeton University

Abstract

Maximising expected value is the classic doctrine in choice theory under empirical uncertainty, and a prominent proposal in the emerging philosophical literature on normative uncertainty, i.e., uncertainty about the standard of evaluation. But how should Expectationalism be stated in general, when we can face both uncertainties simultaneously, as is common in life? Surprisingly, different possibilities arise, ranging from Ex-Ante to Ex-Post Expectationalism, with several hybrid versions. The difference lies in the perspective from which expectations are taken, or equivalently the amount of uncertainty packed into the prospect evaluated. Expectationalism thus faces the classic dilemma between ex-ante and ex-post approaches, familiar elsewhere in ethics and aggregation theory under uncertainty. We analyse the spectrum of expectational theories, showing that they reach diverging evaluations, use different modes of reasoning, take different attitudes to normative risk as well as empirical risk, but converge under an interesting (necessary and sufficient) condition.

1 The problem

When evaluating choice options, we often face two types of uncertainty. Decision theory has focused on *empirical* uncertainty: uncertainty about empirical facts, e.g., facts about weather or election outcomes. Philosophers have recently turned to *normative* uncertainty: uncertainty about the standard of evaluation, due for instance to competing normative intuitions. Important contributions include Oddie (1994), Lockhart (2000), Jackson and Smith (2006), MacAskill (2014, 2016), Bradley and Drechsler (2014), Sepielli (2009), Weatherson (2014), Lazar (2017), Greaves and Cotton-Barratt (2019), Tarsney (2018a, 2019), MacAskill and Ord (forth.), Riedener (2020), and Dietrich and Jabarian (2019a). Normative uncertainty is omnipresent in deliberation and decision-making: parents may wonder how much they should value child autonomy,

[1] Acknowledgements to be added.

even when certain about all relevant empirical facts; agents may wonder which consequences matter most, and whether deontological considerations matter too; reason-based agents (as in Dietrich and List 2013, 2017) may wonder which properties matter, and how they matter; and so on.

Much can be debated about the metaphysical status of normative uncertainty: Is it uncertainty about subjective or objective facts? About real or constructed facts? And so forth. We put these important debates aside. But we stress two things: normative uncertainty is meaningful under many interpretations and metaethical views, and it is not formally reducible to standard choice-theoretic uncertainty.[2]

A prominent proposal in the theory of normative uncertainty is to maximise expected value across standards of evaluation, an approach pioneered in some of the cited works, by Oddie, MacAskill, and Ord, respectively. Riedener (2020) gives a sophisticated axiomatic characterization. Much of today's debate focuses on whether this approach – to be called *Expectationalism* – is justified. Expectationalism has for instance been defended through analogies with empirical uncertainty and through formal arguments (e.g., Riedener's work), but criticized for relying on a precise quantification of normative uncertainty and on certain measurements and comparisons of value (e.g., Tarsney 2018b). Instead of addressing the important question 'Expectationalism or not?', we shall assume Expectationalism, and ask a new question: 'Expected value of what?' A particular answer was so far taken for granted: expected value *of the option itself*. We call this Standard Expectationalism. Standard Expectationalism is hybrid in nature: it reasons 'empirically ex-ante', but 'normatively ex-post', as will be shown. There are two non-hybrid incarnations of Expectationalism, Ex-Ante Expectationalism and Ex-Post Expectationalism, which form the expected value of prospects that respectively contain *all* uncertainty (empirical or normative) or *no* uncertainty. Between these two extremes, there is a spectrum of hybrid (semi-ex-ante) versions of Expectationalism.

The dilemma between ex-ante and ex-post reasoning is prominent in other fields of formal ethics and aggregation theory; see Diamond (1976), McCarthy (2006, 2008, 2015), Fleurbaey (2010), Fleurbaey and Voorhoeve (2016), and Fleurbaey and Zuber (2017). The theory of normative uncertainty cannot escape this dilemma. Just as social egalitarians face a dilemma between ex-ante and ex-post equality, so expectationalists about normative uncertainty face a dilemma between ex-ante, ex-post, and hybrid formulations of Expectationalism. The dilemma always takes the same form: should a given paradigm or ideal – for instance equality or (in our case) expected value – be pursued from an ex-ante or ex-post perspective? The ex-ante/ex-post dilemma reaches also into non-expectational approaches to normative uncertainty, such as maxmin approaches, of which one could again envision ex-ante, ex-post, and hybrid versions.

[2] Some decision-theoretic models can be reinterpreted in terms of normative uncertainty. Examples are multi-utility models and Harsanyi's (1978) impartial-observer model. But the attempt to simply reinterpret choice-theoretic risk or uncertainty (in von-Neumann-Morgenstern's or Savage's framework) normatively runs into formal problems. For instance, writing 'normative information' into Savage's states implies letting states determine utilities, in ways not even compatible with standard state-dependent utility theory (see however Riedener 2020 for a non-standard analysis). This would ultimately undermine the two-attitude make-up of choice theory, which is based on two independent ingredients, beliefs and values (or formally, probabilities and utilities).

But we here exclusively focus on Expectationalism. Our goal will not be to defend some version of Expectationalism against others, but to put the spectrum of expectational theories on the table. Deciding between expectational theories – choosing the right degree of ex-post-ness – might prove as difficult as deciding between ex-ante and ex-post egalitarianism.

Although existing work has assumed a specific formulation of Expectationalism, it may be taken to address Expectationalism in general, because most arguments raised for or against Expectationalism do not hinge on that formulation. Our contribution is thus not to undermine existing arguments, but to raise the problem of the right formulation of Expectationalism.

The paper has a simple structure. The rest of this section introduces the simplest possible formal framework for capturing normative and (where needed) empirical uncertainty. Section 2 introduces four salient expectational theories – four solutions to the problem 'expectation of what?'. Section 3 defines these four theories formally. Section 4 illustrates how they reach diverging evaluations and take different attitudes to risk. Section 5 defines Expectationalism in general, going beyond the four special theories. Section 6 concludes the paper.

We now present the formal ingredients one by one.

The objects of evaluation. We consider a non-empty set A of objects of evaluation, called '*options*'. They could be policy measures, social arrangements, income distributions etc. So far we leave open whether options contain empirical risk.

Competing valuations. Options have uncertain value. Following the expected-value approach, our notion of value is absolute, not comparative (ordinal). We thus represent a possible standard of evaluation by a function v, called a *valuation*, assigning to each option a in A a value $v(a)$ in $\mathbb{R}$. The agent hesitates between certain valuations. Let $\mathcal{V}$ be their set, formally a finite non-empty set of functions from A to $\mathbb{R}$. Let us give moral examples, without suggesting a restriction to moral choice. $\mathcal{V}$ could contain a utilitarian, an egalitarian, and some deontological valuation. Alternatively, $\mathcal{V}$ could consist of 'similar' valuations which differ only in a parameter: prioritarian valuations with different degrees of prioritarianism, or egalitarian valuations with different degrees of inequality-aversion, or valuations of intertemporal well-being with different discounting of future well-being, or valuations (of risky options) with different degrees of risk-aversion, etc. In such parametric examples, normative uncertainty boils down to uncertainty about the correct parameter value: the correct amount of prioritarianism, inequality-aversion, discounting, risk-aversion, etc.[3]

Credences in valuations. You assign to each valuation v in $\mathcal{V}$ a subjective correctness probability $Pr(v) \geq 0$, where $\sum_{v \in \mathcal{V}} Pr(v) = 1$. Probabilities capture degrees of belief about value.

[3]Our examples show that normative uncertainty comes in two species: mere 'parameter uncertainty' and fundamental 'model uncertainty'. A similar distinction is made in other fields, especially statistics and macroeconomics (e.g., Hansen and Sargent 2001).

Meta-value. Given these credences, how should you evaluate options *overall*? An answer takes the form of a *meta-valuation* or *theory*, formally a function assigning to each option in A a (meta-)value in $\mathbb{R}$. To distinguish meta-valuations from valuations, we denote them by upper-case letters like 'V'. Two examples suffice for now. Standard Expectationalism evaluates options $a \in A$ by their expected value: $V(a) = \sum_{v \in \mathcal{V}} Pr(v)v(a)$. Another (decidedly non-expectational) theory evaluates options $a \in A$ by their minimal possible value: $V(a) = \min_{v \in \mathcal{V}: Pr(v) \neq 0} v(a)$.

Some would call our theories 'meta-theories', and our valuations 'first-order theories'. Nothing hinges on having defined theories as value *functions* rather than value *orders* on A; this is for convenience. Readers who prefer an ordinal notion of meta-value can replace our meta-valuations by the orders they induce.

Measurability and comparability of value. As usual in the expectational approach, we take first-order value to be measurable and comparable across valuations. Full measurability makes it meaningful to say that an option x has value 7 under a valuation v ($v(x) = 7$), or is twice as valuable as another option y ($v(x) = 2v(y)$), or exceeds z's value by 2 ($v(x) - v(z) = 2$), etc. Full comparability makes it meaningful to say that two valuations v and v' assign same value to option x ($v(x) = v'(x)$), or same value gain to the change from option x to option y ($v(y) - v(x) = v'(y) - v'(x)$), etc.[4] Such assumptions are strong and debatable. They can be relaxed, in ways that differ across versions of Expectationalism.[5] We set aside when and how measurability and comparability can be justified,[6] and how they could be relaxed by different versions of Expectationalism.

Adding empirical uncertainty: options as lotteries. The above framework is complete as a model of purely normative uncertainty; empirical uncertainty in options is allowed, but not modelled. To add empirical uncertainty explicitly, we hereafter assume that options in A are lotteries on a given set X of *outcomes*, i.e., functions a from X to $[0, 1]$ such that $\sum_{x \in X} a(x) = 1$, where $a(x)$ is non-zero for only finitely many x in X. An option is *riskless* if some outcome has probability one, and *risky* otherwise. Outcomes represent empirical states of affairs after resolution of empirical uncertainty. They may be 'consequences' of actions or go beyond 'consequences', depending on what we wish to model. Interpreting outcomes as consequences limits us to consequentialist valuations, hence to normative uncertainty between types of consequentialism. But if outcomes go beyond consequences, e.g., by capturing intentions or the choice context, then the model is open to non-consequentialist valuations, hence to normative uncertainty between possibly non-consequentialist valuations.[7]

We do not require that *all* lotteries on X count as options, i.e., belong to A. But

[4] Comparability and measurability are addressed by Bossert and Weymark (2004), and in the context of normative uncertainty by, e.g., Ross (2006), Sepielli (2009) and Tarsney (2018b).

[5] For instance, all versions need only *affine* measurements of value, and Standard Expectationalism needs only *unit* comparisons, not *level* comparisons.

[6] Justifying cross-valuation comparisons is easier if $\mathcal{V}$ consists of theories of similar type, e.g., egalitarian theories with different degrees of inequality-aversion.

[7] Normative uncertainty between non-consequentialist valuations is addressed by Barry and Tomlin (2016) and Tenenbaum (2017).

we assume that A contains at least the riskless lotteries, which assign probability one to some outcome. We take valuations v in $\mathcal{V}$ to also evaluate outcomes x in X, by defining $v(x)$ as $v(a)$ where a is the riskless option corresponding to x.

Valuations of vNM type and of non-vNM type. A valuation v in $\mathcal{V}$ could have the notorious von-Neumann-Morgenstern property: it could 'be vNM'. Being vNM means evaluating options a in A by the expected value of the outcome: $v(a) = \sum_{x \in X} a(x)v(x)$. Ever since the Harsanyi-Sen debate, it is controversial whether the vNM property is an arbitrary or an inherently necessary feature of 'value'.[8] Our model is ecumenical: the agent could be utterly certain of the 'vNM Hypothesis', by having positive credence only in vNM valuations in $\mathcal{V}$; or utterly certain of the converse, by having positive credence only in non-vNM valuations in $\mathcal{V}$; or genuinely uncertain, by having positive credence in both types of valuations. Being ecumenical is important because we (as modellers) should not project our own view w.r.t. the vNM Hypothesis upon the agent to be modelled. We should for instance not force $\mathcal{V}$ to contain only vNM valuations or only non-vNM valuations, because this would restrict the entire framework to agents who are utterly certain of the vNM Hypothesis or of its converse, thereby excluding an important kind of normative uncertainty, namely uncertainty about the vNM Hypothesis. There is a deeper methodological point: the field or theory of normative uncertainty is engaged in meta-normativity and should thus avoid prejudging first-order normative questions. It should take people's normative beliefs and uncertainties at face value (without 'forbidding' some of them), and tell people how to respond to their own uncertainties. Earlier work about normative uncertainty has often restricted attention to vNM valuations. This can be legitimate as a working assumption, but we wish to overcome this restriction.

2 Expected value of what?

Expectational theories evaluate options by the expected value of *some object*. That object is the prospect offered by the option, but there are different types of prospect: the ex-ante prospect, the ex-post prospect, and hybrid prospects taking a partly ex-ante and partly ex-post perspective.

Think of prospects as probability distributions. More precisely, one can define prospects equivalently as distributions over empirical-normative worlds ('world prospects') or distributions over resulting value levels ('value prospects'). We shall later only work with value prospects. But let us start with world prospects. An *empirical-normative world* – for short, a *world* – is a pair (x, v) of an outcome in X (an 'empirical world') and a valuation in $\mathcal{V}$ (a 'normative world'). In a world, all empirical or normative uncertainty is resolved. A *world prospect* is a probability distribution over worlds, representing how likely worlds are (where for simplicity only finitely many worlds have non-zero probability). Each option a generates an (ex-ante) world prospect, under which the probability of a world (x, v) is the product $a(x)Pr(v)$ of the probabilities of

[8] See Broome (1991), Weymark (1991), Nissan-Rozen (2015) and Greaves (2017) for analyses with deviating conclusions.

outcome x (under option a) and valuation v. This world prospect is ex-ante because no uncertainty is resolved; ex-post and hybrid world prospects will be defined in a moment.

We can now give four possible answers to the question 'Expected value of what?', hence four ways to reason and ultimately to assign meta-value to a given option a. We keep the answers informal; formal definitions follow in Section 3.

- *Normatively ex-post reasoning*: You place yourself in a normatively ex-post and empirically ex-ante position, by considering a given valuation v and the lottery of empirical outcomes generated by option a. So you face the *normatively ex-post* world prospect, in which v has (marginal) probability one and any outcome x in X has (marginal) probability $a(x)$. It yields the value $v(a)$. Stepping outside this position, you then form the expectation of the value $v(a)$ across valuations v in $\mathcal{V}$. This is Standard Expectationalism.
- *Ex-post reasoning*: You place yourself in a fully ex-post position, by considering a given outcome x and a given valuation v. So you face the *ex-post* world prospect, in which world (x, v) has probability one. It yields the value $v(x)$. Stepping outside this position, you then form the expectation of the value $v(x)$ across worlds (x, v) in $X \times \mathcal{V}$. This is Ex-Post Expectationalism.
- *Ex-ante reasoning*: You place yourself in the fully ex-ante position, in which both parts of the empirical-normative world are unknown. So you face the *ex-ante* world prospect, defined above. You then form the expected value of this ex-ante prospect; how this works is shown in Section 3. This is Ex-Ante Expectationalism.
- *Empirically ex-post reasoning*: You place yourself in an empirically ex-post and normatively ex-ante position, by considering a given outcome x and the probability distribution over valuations Pr reflecting your normative uncertainty. So you face the *empirically ex-post* world prospect, in which x has (marginal) probability of one and any valuation v has (marginal) probability $Pr(v)$. You then form the expected value of this world prospect, in a way shown in Section 3. This is Reverse Expectationalism. It is the reverse or 'dual' of Standard Expectationalism, as it reasons ex-ante where Standard Expectationalism reasons ex-post, and vice versa.

	normatively ex-post	normatively ex-ante
empirically ex-post	Ex-Post Expectationalism	Reverse Expectationalism
empirically ex-ante	Standard Expectationalism	Ex-Ante Expectationalism

Table 1: Four expectational theories and their modes of reasoning

These four answers to the question 'Expected value of what?' were given in the form of world prospects, but they can be redescribed as *value* prospects. Value prospects are prospects of achieving certain value levels (not worlds) with certain probabilities, for instance achieving value 4 with probability 1/2 and value 0 with probability 1/2. A world prospect immediately induces a value prospect (mathematically by taking the image of the world prospect under the mapping $(x, v) \mapsto v(x)$ from worlds to resulting values). For instance, the ex-post world prospect under which world (x, v) is certain induces the riskless value prospect under which the value $v(x)$ is certain.

Formally, a *value prospect* is simply a lottery over real numbers, i.e., a function p

assigning to each value k in $\mathbb{R}$ a probability $p(k)$ in $[0,1]$ such that $\sum_{k\in\mathbb{R}} p(k) = 1$, where (for simplicity) only finitely many values k in $\mathbb{R}$ have non-zero probability $p(k)$. Each option a generates a value prospect, denoted p_a. It reflects empirical and normative uncertainty, as the resulting value $v(x)$ depends on both x and v, hence on the empirical-normative world (x, v). The probability that the resulting value is (say) 4 is the sum-total probability of all worlds (x, v) such that $v(x) = 4$. The just-defined value prospect p_a of an option a is an ex-ante construct: no uncertainty is yet resolved. Indeed, p_a is simply the value prospect induced by the ex-ante world prospect. Partly or fully ex-post value prospects are definable by eliminating one or both sources of uncertainty.

We now formally define the four kinds of value prospect. They correspond exactly to the four kinds of world prospect above, respectively:[9]

- The **(ex-ante) value prospect of option** $a \in A$ is the value prospect 'p_a' such that any value $k \in \mathbb{R}$ has probability

$$p_a(k) = \text{ 'probability that } a \text{ leads to value } k\text{'} = \sum_{(x,v)\in X\times\mathcal{V}:v(x)=k} \underbrace{a(x)Pr(v)}_{\text{prob. of }(x,v)}.$$

- The **(normatively ex-post) value prospect of option** $a \in A$ **given valuation** $v \in \mathcal{V}$ is the value prospect '$p_{a,v}$' such that any value $k \in \mathbb{R}$ has probability

$$p_{a,v}(k) = \text{ 'probability that } a \text{ leads to value } k \text{ given } v\text{'} = \sum_{x\in X:v(x)=k} a(x).$$

- The **(empirically ex-post) value prospect given outcome** $x \in X$ is the value prospect 'p_x' such that any value $k \in \mathbb{R}$ has probability:

$$p_x(k) = \text{ 'probability that } x \text{ leads to value } k\text{'} = \sum_{v\in\mathcal{V}:v(x)=k} Pr(v).$$

- The **(ex-post) value prospect given both** $x \in X$ **and** $v \in \mathcal{V}$ is the riskless value prospect '$p_{x,v}$' under which the value is $v(x)$ with probability one.[10]

3 Four expectational theories

We now formally define the four expectational theories discussed in Section 2. Each takes the expected value of a certain prospect, as is either clear by definition or established in Theorem 1. We begin with the two theories whose definitions do not explicitly refer to prospects.

Standard Expectationalism ('EV_{stan}'): *The meta-value of an option $a \in A$ is the expected value of the option itself:*

$$EV_{stan}(a) = \sum_{v\in\mathcal{V}} Pr(v)v(a) \text{ ('standard expected value').}$$

[9] Compare our value prospects with Rowe and Voorhoeve's (2018) well-being prospects in a context of health ethics under (purely empirical) risk, uncertainty, or ambiguity.

[10] The value prospects p_x and $p_{x,v}$ can be regarded as special cases of the value prospects p_a and $p_{a,v}$, by choosing a to be the riskless option that yields x for sure.

This theory reasons empirically ex-ante, because the object whose average evaluation it forms is *the option*, which captures empirical risk. The second theory reasons fully ex-post: it forms the average evaluation of *the outcome*, which no longer contains empirical risk. This requires averaging across both outcomes and valuations, hence across empirical-normative worlds (x, v). Formally:

Ex-Post Expectationalism ('EV_{post}'): *The meta-value of an option $a \in A$ is the expected value of the outcome:*

$$EV_{post}(a) = \sum_{(x,v) \in X \times \mathcal{V}} \underbrace{a(x)Pr(v)}_{prob.\ of\ (x,v)} v(x) \ \ (\text{'expected final value'}).$$

The third theory reasons fully ex-ante. It operates neither at the ex-post level of outcomes (EV_{post}), nor at the semi-ex-post level of options (EV_{stan}), but at the level of ex-ante value prospects. But how can valuations v in $\mathcal{V}$ evaluate value prospects rather than options, i.e., how should we define $v(p)$ for a value prospect p? We of course identify $v(p)$ with $v(a)$ for any option a in A chosen such as to have the value prospect p given v. If for instance p is the value prospect 'the value is 1 or 0 equiprobably', then we pick an option a which equiprobably has an outcome x of value $v(x) = 1$ or an outcome y of value $v(y) = 0$, and define $v(p)$ as $v(a)$. Formally, the *value of a value prospect p under a valuation v* in $\mathcal{V}$ – denoted $v(p)$ – is the value $v(a)$ of options $a \in A$ such that $p_{a,v} = p$. This definition implicitly rests on an assumption that we shall maintain for the rest of the paper:

Assumption: Hereafter, for each valuation v in $\mathcal{V}$ and value prospect p we assume that (i) A contains an option a whose value prospect given v, $p_{a,v}$, is p; and (ii) any two such options a in A have same value $v(a)$.

Condition (i) is a typical richness assumption: the set of options A should be sufficiently inclusive, i.e., contain options with any given value prospects. Condition (ii) is a consistency assumption on the valuations in $\mathcal{V}$. It is compatible with most or all first-order theories one would naturally want to consider.

Ex-Ante Expectationalism ('EV_{ante}'): *The meta-value of an option $a \in A$ is the expected value of the ex-ante prospect:*

$$EV_{ante}(a) = \sum_{v \in \mathcal{V}} Pr(v)v(p_a) \ \ (\text{'expected ex-ante value'}).$$

Note an intended peculiarity: $v(p_a)$ uses a *given* valuation (v) to evaluate a prospect (p_a) which carries normative uncertainty about the correct valuation. Precisely this is what ex-ante reasoning should do, as it should ask how attractive each *ex-ante* prospect is on average across possible valuations.

The fourth theory calculates the average evaluation of yet another object: neither the option (EV_{stan}), nor the outcome (EV_{post}), nor the ex-ante prospect (EV_{ante}), but the empirically ex-post value prospect. This requires averaging across outcomes and valuations, hence across empirical-normative worlds (x, v).

Reverse Expectationalism ('EV_{rev}'): *The meta-value of an option $a \in A$ is the expected value of the empirically ex-post prospect:*

$$EV_{rev}(a) = \sum_{(x,v)\in X\times\mathcal{V}} \underbrace{a(x)Pr(v)}_{prob.\ of\ (x,v)} v(p_x) \text{ ('reverse expected value').}$$

This theory reverses the reasoning of Standard Expectationalism: it reasons empirically ex-post rather than normatively ex-post.

The following theorem re-expresses the four theories in a comparable format, showing that they only differ in the 'locus' of expectation-taking, i.e., in the sort of prospect whose expected value they maximise:

Theorem 1 *Each expectational theory $V \in \{EV_{ante}, EV_{post}, EV_{stan}, EV_{rev}\}$ evaluates any option $a \in A$ by the expected value of a specific value prospect, i.e.,*

$$V(a) = \sum_{(x,v)\in X\times\mathcal{V}} \underbrace{a(x)Pr(v)}_{prob.\ of\ (x,v)} v(p),$$

where 'p' stands for the

- *ex-ante value prospect p_a if $V = EV_{ante}$,*
- *ex-post value prospect $p_{x,v}$ if $V = EV_{post}$,*
- *normatively ex-post value prospect $p_{a,v}$ if $V = EV_{stan}$,*
- *empirically ex-post value prospect p_x if $V = EV_{rev}$.*

4 Illustration of these expectational theories and their risk attitudes

Suppose you hesitate between just two valuations, v and v'. You have credence $\frac{1}{2}$ in each of them, and credence 0 in all other valuations in $\mathcal{V}$ (if any). Both valuations v and v' are risk-averse. So you are sure that risk-aversion is correct: your normative uncertainty does not pertain to the risk-attitude (more on risk attitudes in the next section). You now compare two options. Both options lead to the value prospect 'value 4 with probability $\frac{1}{2}$, value 0 with probability $\frac{1}{2}$', denoted $4_{50\%}0_{50\%}$, but for very different reasons. The first option involves only normative risk: it surely has outcome x, whose value is either $v(x) = 4$ or $v'(x) = 0$. The second option involves only empirical risk: it has either outcome y or outcome z (equiprobably), where it is uncontroversial between v and v' that y has value 4 and z has value 0. By risk-aversion, the option is evaluated below the expected resulting value of $\frac{1}{2}4 + \frac{1}{2}0 = 2$; let the value be 1 under both v and v'. The gap from 1 to 2 is a 'risk penalty' or 'risk premium'.

Table 2 displays the (ex-ante and normatively ex-post) value prospects of options and the evaluations of options by both first-order theories and the four expectational theories. The four meta-evaluations are obtained as follows:

- *Standard* Expectationalism forms the average value of the option. This yields $\frac{1}{2}4 + \frac{1}{2}0 = 2$ or $\frac{1}{2}1 + \frac{1}{2}1 = 1$, respectively.

	value prospect			evaluation by					
	given v	given v'	ex-ante	v	v'	EV_{stan}	EV_{post}	EV_{ante}	EV_{rev}
option 1	$4_{100\%}$	$0_{100\%}$	$4_{50\%}0_{50\%}$	4	0	2	2	1	1
option 2	$4_{50\%}0_{50\%}$	$4_{50\%}0_{50\%}$	$4_{50\%}0_{50\%}$	1	1	1	2	1	2

Table 2: Four expectational theories applied to two concrete options

- *Ex-Post* Expectationalism forms the average value of the outcome. In principle, this requires averaging across valuations in $\mathcal{V}$ (normative uncertainty) and outcomes (empirical uncertainty). Yet our options effectively need just one dimension of averaging, as they have just one source of uncertainty. The first option has just normative uncertainty: it surely has outcome x, of value 4 or 0. The second option has just empirical uncertainty: it has outcome y of sure value 4 or outcome z of sure value 0. Each option thus has the same average value of the outcome: $\frac{1}{2}4 + \frac{1}{2}0 = 2$.
- *Ex-Ante* Expectationalism forms the average value of the ex-ante value prospect, i.e., of $4_{50\%}0_{50\%}$ for each option. So we must calculate $\frac{1}{2}v(4_{50\%}0_{50\%}) + \frac{1}{2}v'(4_{50\%}0_{50\%})$. What are $v(4_{50\%}0_{50\%})$ and $v'(4_{50\%}0_{50\%})$? As $4_{50\%}0_{50\%}$ is option 2's value prospect given v, $v(4_{50\%}0_{50\%}) = v(\text{option } 2) = 1$. As $4_{50\%}0_{50\%}$ is also option 2's value prospect given v', $v'(4_{50\%}0_{50\%}) = v'(\text{option } 2) = 1$. So, $\frac{1}{2}v(4_{50\%}0_{50\%}) + \frac{1}{2}v'(4_{50\%}0_{50\%}) = \frac{1}{2}1 + \frac{1}{2}1 = 1$.
- *Reverse* Expectationalism forms the average value of the empirically ex-post value prospect. Like for EV_{post}, this can require averaging across both outcomes and valuations, but for our two options one dimension of averaging drops out, as option 1 is empirically riskless and option 2 is normatively riskless. Option 1 surely has outcome x, whose value prospect $4_{50\%}0_{50\%}$ is evaluated at 1 by both (risk-averse) valuations, as just seen. The average value is thus $\frac{1}{2}1 + \frac{1}{2}1 = 1$. Option 2 either has outcome y, whose value prospect $4_{100\%}$ has value 4 under both v and v'; or has outcome z, whose value prospect $0_{100\%}$ has value 0 under both v and v'. The average value is thus $\frac{1}{2}4 + \frac{1}{2}0 = 2$.

Interestingly, the reasoning mode – the version of Expectationalism adopted – determines the risk attitude underlying our (meta-)evaluations. We shall explain this only informally; the relevant definitions and formal results are given in a separate paper dedicated to risk attitudes (Dietrich and Jabarian 2019b). By 'risk' we mean risk about resulting value, such as (in our example) the risk of achieving value 4 or 0. Such risk can have empirical or normative origin: it can stem from an uncertain outcome (empirical uncertainty) or an uncertain value of the outcome (normative uncertainty) or even a combination. A (meta-)theory is risk-averse if its evaluation of options contains a penalty for risk, i.e., falls below the value achieved in expectation. The theory is risk-neutral if risk is not penalized. One might feel uncomfortable with imposing a particular risk attitude on meta-evaluations: where would the justification come from? An interesting alternative to imposition is to adopt whatever risk attitude you believe to be correct for first-order valuations: if you are certain of a particular risk attitude such as risk-aversion, i.e., hold positive credence only in valuations with that risk-attitude, then your meta-evaluations adopt that same risk-attitude. We call this *risk-impartiality* because your meta-level risk attitude defers to your risk-attitudinal judgments. In our example, you are certain that risk-aversion is correct, as you are certain that one of the risk-averse valuations v and v' is correct; so a risk-impartial theory is risk-averse. Risk-impartiality seems to be a natural default, at least in the absence of a convincing argument for any particular meta-level risk attitudes (see Dietrich and Jabarian 2019b for risk-impartiality in cases of not being certain of any given risk-attitude).[11]

Which risk-attitudes do our four theories have? As summarized in Table 3, one is risk-neutral (no penalty for risk), one is risk-impartial (deference to risk-attitudinal judgments), and two have hybrid risk-attitudes, i.e., are risk-neutral or risk-impartial depending on the origin of risk. To explain why, we use our example, in which all

	neutral to normative risk	impartial to normative risk
neutral to empirical risk	Ex-Post Expectationalism	Reverse Expectationalism
impartial to empirical risk	Standard Expectationalism	Ex-Ante Expectationalism

Table 3: The risk attitudes of the four meta-theories in our example with risk-averse first-order theories

valuations are risk-averse, so that risk-impartiality becomes risk-aversion.

- *Standard* Expectationalism applies the valuations v and v' to the option, which captures only empirical risk. This leads (by risk-aversion of v and v') to a penalty or discount for empirical risk only: the theory is averse to empirical risk, but neutral to normative risk. This explains why in Table 2 the normatively risky option 1 gets the undiscounted value of 2, while the empirically risky option 2 gets the discounted value of 1.
- *Ex-Post* Expectationalism applies the two valuations to the outcome, which captures no risk. So no risk is penalized: the theory is globally risk-neutral. This explains why both options in Table 2 get the undiscounted value of 2.
- *Ex-Ante* Expectationalism applies the two valuations to the ex-ante value prospect, which captures risk of both origins. So all risk is penalized: the theory is globally risk-averse. This explains why both options in Table 2 get the discounted value of 1.
- *Reverse* Expectationalism applies the two valuations to the empirically ex-post value prospect, which captures only normative risk. So only normative risk is penalized: the theory is averse only to normative risk. This explains why in Table 2 only the normatively risky option gets the discounted value of 1.

5 The full taxonomy of Expectationalism

We now come to the unification. We introduce a single generic expectational theory, of which our four earlier theories are nothing but special cases. The generic theory depends on a parameter that determines the reasoning mode, i.e., the extent of ex-post-ness. Particular choices of this parameter yield our four special expectational

[11] Risk attitudes have been analysed extensively in the different context of purely empirical uncertainty. For different accounts, see Weirich (1986), Buchak (2013), Bradley and Stefánsson (2017) and Baccelli (2018).

theories, and all other expectational theories. So there are not just four expectational theories, but a large and unified class of expectational theories.

The parameter determining the expectational theory is the type of *information* relative to which reasoning is ex-post: full information yields Ex-Post Expectationalism, no information yields Ex-Ante Expectationalism, purely normative information yields Standard Expectationalism, purely empirical information yields Reverse Expectationalism, and yet other types of information yield other expectational theories. We model an information by an *empirical-normative event* $I \subseteq X \times \mathcal{V}$, containing the empirical-normative worlds (x, v) which are consistent with the information. The information of a full empirical-normative world (x, v) is $I = \{(x, v)\}$, containing just one world; the vacuous or tautological information is $I = X \times \mathcal{V}$, containing all worlds; the information of a valuation v is $I = X \times \{v\}$, containing worlds of type $(*, v)$; and the information of an outcome x is $I = \{x\} \times \mathcal{V}$, containing worlds of type $(x, *)$. Recall that each option a generates a world prospect, i.e., a probability function over worlds. Let us denote it by P_a. The probability of a world (x, v) is $P_a(x, v) = a(x)Pr(v)$, the product of the probabilities of x and v.

To define our general expectational theory, we need a general notion of value prospect, which has an arbitrary degree of ex-post-ness, i.e., conditionalises on an arbitrary information I. We call it the 'ex-I value prospect', as it is the value prospect from the perspective of I. Formally, for any option $a \in X$ and information $I \subseteq X \times \mathcal{V}$ (of non-zero probability $P_a(I)$), the *ex-I value prospect* of a is the value prospect $p_{a,I}$ such that the probability of a value level $k \in \mathbb{R}$ is the probability that a results in value k given I:

$$\begin{aligned} p_{a,I}(k) &= \text{probability of final value } k \text{ given } I = \frac{\text{prob. of } [I \text{ \& final value } k]}{\text{prob. of } I} \\ &= \frac{P_a(\{(x,v) \in I : v(x) = k\})}{P_a(I)}. \end{aligned}$$

This general notion of value prospect encompasses our four earlier notions:

Proposition 1 *The ex-I value prospect $p_{a,I}$ of an option $a \in A$ given an information $I \subseteq X \times \mathcal{V}$ (of non-zero probability $P_a(I)$) coincides with the*

- *ex-ante value prospect p_a if $I = X \times \mathcal{V}$ (no information),*
- *ex-post value prospect $p_{x,v}$ if $I = \{(x, v)\}$ (information of a full world (x, v)),*
- *normatively ex-post value prospect $p_{a,v}$ if $I = X \times \{v\}$ (information of a valuation v),*
- *empirically ex-post value prospect p_x if $I = \{x\} \times \mathcal{V}$ (information of an empirical outcome x).*

Recall that each valuation v in $\mathcal{V}$ can evaluate not just options, but also value prospects. So we can form $v(p_{a,I})$, which tells how valuable v finds the prospect of option a given I. We call $v(p_{a,I})$ a's *ex-I value of* a, according to v.

An expectational theory reasons ex-post w.r.t. some type of information. A type of information is represented by an *information partition*: a partition $\mathcal{I}$ of the set $X \times \mathcal{V}$ of empirical-normative worlds. $\mathcal{I}$ contains those information I on which the reasoner conditionalises when conceptualizing options as prospects. As such, $\mathcal{I}$ defines a degree of ex-post-ness of reasoning. Fully ex-post reasoning is defined by the finest information partition $\mathcal{I} = \{\{(x,v)\} : (x,v) \in X \times \mathcal{V}\}$; fully ex-ante reasoning by the coarsest partition $\mathcal{I} = \{X \times \mathcal{V}\}$; normatively ex-post reasoning by the partition $\mathcal{I} = \{X \times \{v\} : v \in \mathcal{V}\}$ into 'valuation events'; empirically ex-post reasoning by the partition $\mathcal{I} = \{\{x\} \times \mathcal{V} : x \in X\}$ into 'outcome events'; and other hybrid reasoning modes by other partitions.

An information partition $\mathcal{I}$ – a degree of ex-post-ness – determines an expectational theory, which evaluates options by the expected value (across empirical-normative worlds) of the prospect w.r.t. $\mathcal{I}$. Formally:

Ex-$\mathcal{I}$ Expectationalism ('$EV_{\mathcal{I}}$'): *The meta-value of an option $a \in A$ is the expected value of the ex-$\mathcal{I}$ prospect:*[12]

$$EV_{\mathcal{I}}(a) = \sum_{(x,v)\in X\times\mathcal{V}} \underbrace{a(x)Pr(v)}_{prob.\ of\ (x,v)} v(p_{a,\mathcal{I}(x,v)}) \text{ ('expected ex-}\mathcal{I}\text{ value')}$$

where $\mathcal{I}(x,v)$ is the information in empirical-normative world (x,v), i.e., the $I \in \mathcal{I}$ containing (x,v).

We can now define 'Expectationalism' as a general approach:

Expectationalism: *Meta-value is given by <u>some</u> expectational theory, i.e., by Ex-$\mathcal{I}$ Expectationalism for <u>some</u> information type $\mathcal{I}$ (some partition of $X \times \mathcal{V}$).*

Our four earlier theories are special cases, obtained by plugging in certain information types, i.e., certain degrees of ex-post reasoning:

Theorem 2 *Ex-$\mathcal{I}$ Expectationalism coincides with*

- *Ex-Ante Expectationalism if $\mathcal{I} = \{X \times \mathcal{V}\}$ (no information),*
- *Ex-Post Expectationalism if $\mathcal{I} = \{\{(x,v)\} : (x,v) \in X \times \mathcal{V}\}$ (full information),*
- *Standard Expectationalism if $\mathcal{I} = \{X \times \{v\} : v \in \mathcal{V}\}$ (normative information),*
- *Reverse Expectationalism if $\mathcal{I} = \{\{x\} \times \mathcal{V} : x \in X\}$ (empirical information).*

Are there any circumstances under which it becomes irrelevant how you reason? That is, can it happen that all degrees of ex-post reasoning extensionally yield the same expectational theory, hence the same evaluative judgments, albeit through different procedures? This question obviously matters. If all reasoning modes were extensionally equivalent, then you could reason as you wish or find easiest. The question has a sharp answer: the reasoning mode is irrelevant *if and only if* you have full credence in the vNM Hypothesis, i.e., assign zero probability to all valuations in $\mathcal{V}$ that are not vNM ('von-Neumann-Morgenstern'). Recall that vNM valuations v evaluate options a in A by the expected value of the outcome: $v(a) = \sum_{x\in X} a(x)v(x)$. See Section 1 for discussion.

[12] Although $p_{a,\mathcal{I}(x,v)}$ becomes undefined in the zero-probability case $P_a(\mathcal{I}(x,v)) = 0$, no ambiguity arises. Whenever $p_{a,\mathcal{I}(x,v)}$ is undefined, the value $v(p_{a,\mathcal{I}(x,v)})$ can be interpreted arbitrarily, as it is multiplied by 0 ($= P_a(x,v) = a(x)Pr(v)$) and so has no effect.

Theorem 3 *All expectational theories $EV_{\mathcal{I}}$ coincide (i.e., the reasoning mode has no effect) if and only if you are certain of the vNM Hypothesis, i.e., $Pr(v) = 0$ for all valuations v in $\mathcal{V}$ that are not vNM.*

Some scholars have defended the vNM Hypothesis (and many have assumed it to simplify models). But few of them would go so far as to be utterly certain of that hypothesis. These few people can safely reason as they wish: their reasoning mode has no effect by Theorem 3. All others, who have at least some doubt in the vNM Hypothesis, face the hard choice between reasoning modes, i.e., between expectational theories.

6 Conclusion

There is more than one expected-value theory for evaluating options. The various expectational theories differ *ethically*, by reaching different evaluations; *procedurally*, by using different reasoning; and *risk-attitudinally*, by taking different attitudes towards empirical as well as normative risk. But all theories take the expected value of some type of prospect. At the two ends of the spectrum, *Ex-Ante* and *Ex-Post* Expectationalism respectively take the expected value of the ex-ante or ex-post prospect, hence reason from the perspective before or after resolution of uncertainty of any type (empirical or normative). *Standard* Expectationalism lies in between: it takes the expected value of the option itself, thereby effectively reasoning from an empirically ex-ante, but normatively ex-post perspective. *Reverse* Expectationalism does the opposite: it reasons empirically ex-post, but normatively ex-ante. The four mentioned theories stand out as salient, but they are just examples. In general, to any type of information (technically, any 'information partition' of the set of empirical-normative worlds) corresponds an expectational theory, which reasons ex-post w.r.t. *this* information.

The classical question 'Expectationalism or not?' should therefore be complemented by another pressing question: 'Expected value of what?' The problem of deciding between versions of Expectationalism might prove to be as difficult as the classic problem of deciding between ex-ante and ex-post versions of egalitarianism – or perhaps even more difficult, as we face normative uncertainty besides empirical uncertainty.

References

Baccelli, J. (2018) Risk Attitudes in Axiomatic Decision Theory—A Conceptual Perspective, *Theory and Decision* 84: 61-82

Barry, C., Tomlin, P. (2016) Moral uncertainty and permissibility: evaluating option sets, *Canadian Journal of Philosophy* 46: 898-923

Bossert, W., Weymark, J. (2004) Utility in social choice. In: Barberà, S., Hammond P., Seidl, C. (eds.) *Handbook of utility theory, vol. 2: Extensions*, Kluwer, Dordrecht, pp. 1099-1177

Bradley, R., Stefánsson, O. (2017) What Is Risk Aversion? *The British Journal for the Philosophy of Science*

Bradley, R., Drechsler, M. (2014) Types of Uncertainty, *Erkenntnis* 79(6):1225-1248

Broome, J. (1991) *Weighing Goods*, Oxford: Blackwell

Buchak, L. (2013) *Risk Aversion and Rationality*, Oxford University Press

Diamond, P. A. (1967) Cardinal welfare, individual ethics, and interpersonal comparison of utility: comment, *Journal of Political Economy* 75: 765-6

Dietrich, F., Jabarian, B. (2019a) Axiomatics of normative uncertainty, working paper

Dietrich, F., Jabarian, B. (2019b) The risk attitude under normative uncertainty, working paper

Dietrich, F., List, C. (2013) A reason-based theory of rational choice, *Noûs* 47: 104-134

Dietrich, F., List, C. (2017) What matters and how it matters: a choice-theoretic representation of moral theories, *Philosophical Review* 126: 421-479

Fleurbaey, M. (2010) Assessing risky social situations, *Journal of Political Economy* 118: 649-680

Fleurbaey, M., Voorhoeve, A. (2016) Priority or equality for possible people? *Ethics* 126: 929-954

Fleurbaey, M., Zuber, S. (2017) Fair management of social risk, *Journal of Economic Theory* 169: 666-706

Greaves, H., Cotton-Barratt, O. (2019) A bargaining-theoretic approach to moral uncertainty', working paper

Greaves, H. (2017) A reconsideration of the Harsanyi-Sen-Weymark debate on utilitarianism, *Utilitas* 29(2): 175-213

Hansen, L. P., Sargent, T. J. (2001) Robust Control and Model Uncertainty, *American Economic Review* 91: 60-66

Harsanyi, J. (1978) Bayesian decision theory and utilitarian ethics, *American Economic Review* 68: 223-228.

Jackson, F., Smith, M. (2006) Absolutist Moral Theories and Uncertainty, *Journal of Philosophy* 103: 267-283

Lazar, S. (2017) Deontological Decision Theory and Agent-Centred Options, *Ethics* 127: 579-609

Lockhart, T. (2000) *Moral Uncertainty and its Consequences*, Oxford University Press

MacAskill, W. (2014) *Normative Uncertainty*, Doctoral thesis, University of Oxford. URL: http://commonsenseatheism.com/wp-content/uploads/2014/03/MacAskill-Normative-Uncertainty.pdf

MacAskill, W. (2016) Normative Uncertainty as a Voting Problem, *Mind* 125: 967-1004

MacAskill, W., Ord, T. (forth.) Why maximize expected choice-worthiness, *Noûs*. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12264

McCarthy, D. (2006) Utilitarianism and prioritarianism I, *Economics and Philosophy* 22:335-363

McCarthy, D. (2008) Utilitarianism and prioritarianism II, *Economics and Philosophy* 24: 1-33

McCarthy, D. (2015) Distributive equality, *Mind* 124: 1045–1109

Nissan-Rozen, I. (2015) Against Moral Hedging. *Economics and Philosophy* 31: 1-21

Oddie, G. (1994) Moral uncertainty and human embryo experimentation. In: Fulford, K. W. M., Gillett, G., Soskice, J. M. (eds.) *Medicine and Moral Reasoning*,

Cambridge University Press, pp. 3-144
Riedener, S. (2020) An axiomatic approach to axiological uncertainty, *Philosophical Studies* 177: 483-504
Rowe, T., Voorhoeve, A. (2018) Egalitarianism under Severe Uncertainty, *Philos Public Aff* 46: 239-268
Savage, L. J. (1954) *The Foundations of Statistics*, New York: Wiley
Sepielli, A. (2009) What to Do When You Don't Know What To Do. In: Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*, Oxford University Press, p. 35
Tarsney, C. (2018a) Moral Uncertainty for Deontologists, *Ethical Theory and Moral Practice* 21: 505–520
Tarsney, C. (2018b) Intertheoretic Value Comparison: A Modest Proposal, *Journal of Moral Philosophy* 15: 324-344
Tarsney, C. (2019) Normative uncertainty and social choice, *Mind*
Tenenbaum, S. (2017) Action, Deontology, and Risk: Against the Multiplicative Model, *Ethics* 127: 674-707
von Neumann, J., Morgenstern, O. (1944) *Theory of Games and Economic Behavior*, Princeton University Press
Weatherson, B. (2014) Running Risks Morally, *Philosophical Studies* 167(1): 141-63
Weirich, P. (1986) Expected utility and risk, *The British Journal for the Philosophy of Science* 37(4): 419-442
Weymark, J. (1991) A reconsideration of the Harsanyi–Sen debate on utilitarianism. In Elster, J., Roemer, J. E. (eds.) *Interpersonal Comparisons of Well-Being*, Cambridge University Press, pp. 255

A Proofs

Proof of Theorem 2. Let $a \in X$. Firstly,

$$EV_{ante}(a) = \sum_{v\in\mathcal{V}} Pr(v)v(p_a) = \sum_{v\in\mathcal{V}} Pr(v)v(p_a) \sum_{x\in X} a(x) = \sum_{(x,v)\in X\times\mathcal{V}} a(x)Pr(v)v(p_a),$$

where the second equality holds as $\sum_{x\in X} a(x) = 1$. Secondly,

$$\begin{aligned} EV_{stan}(a) &= \sum_{v\in\mathcal{V}} Pr(v)v(a) = \sum_{v\in\mathcal{V}} Pr(v)v(p_{a,v}) \\ &= \sum_{v\in\mathcal{V}} Pr(v)v(p_{a,v}) \sum_{x\in X} a(x) = \sum_{(x,v)\in X\times\mathcal{V}} a(x)Pr(v)v(p_{a,v}), \end{aligned}$$

where the second equality holds because $v(p_{a,x}) = v(a)$, and the third because $\sum_{x\in X} a(x) = 1$. Thirdly, the expression for $EV_{rev}(a)$ holds by definition. Finally,

$$EV_{post}(a) = \sum_{(x,v)\in X\times\mathcal{V}} a(x)Pr(v) \underbrace{v(x)}_{=v(p_{x,v})} = \sum_{(x,v)\in X\times\mathcal{V}} a(x)Pr(v)v(p_{x,v}). \blacksquare$$

Proof of Proposition 1. Consider an option $a \in A$ and an information $I \subseteq X \times \mathcal{V}$ such that $P_a(I) \neq 0$. As our definitions easily imply, if $I = X \times \mathcal{V}$ then $p_{a,I} = p_a$, while if

$I = \{(x,v)\}$ where $(x,v) \in X \times \mathcal{V}$ then $p_{a,I} = p_{x,v}$. If $I = X \times \{v\}$ where $v \in \mathcal{V}$, then $p_{a,I} = p_{a,v}$ because for all $k \in \mathbb{R}$

$$\begin{aligned} p_{a,I}(k) &= \frac{P_a(\{x \in X : v(x) = k\} \times \{v\})}{Pr(v)} = \frac{a(\{x \in X : v(x) = k\})Pr(v)}{Pr(v)} \\ &= a(\{x \in X : v(x) = k\}) = \sum_{x \in X : v(x) = k} a(x) = p_{a,v}(k). \end{aligned}$$

Finally, if $I = \{x\} \times \mathcal{V}$ where $x \in X$, then $p_{a,I} = p_x$ because for all $k \in \mathbb{R}$

$$\begin{aligned} p_{a,I}(k) &= \frac{P_a(\{x\} \times \{v \in \mathcal{V} : v(x) = k\})}{a(x)} = \frac{a(x)Pr(\{v \in \mathcal{V} : v(x) = k\})}{a(x)} \\ &= Pr(\{v \in \mathcal{V} : v(x) = k\}) = \sum_{v \in \mathcal{V} : v(x) = k} Pr(v) = p_x(k). \ \blacksquare \end{aligned}$$

Proof of Theorem 2. Regarding EV_{ante}, for each option $a \in A$

$$EV_{ante}(a) = \sum_{(x,v) \in X \times \mathcal{V}} a(x)Pr(v)v(p_a) = EV_{\mathcal{I}}(a) \text{ for } \mathcal{I} = \{X \times \mathcal{V}\}.$$

where the first identity holds by Theorem 1 and the second identity holds because by Proposition 1 we can replace p_a by $p_{a,X \times \mathcal{V}} = p_{a,\mathcal{I}(x,v)}$. Analogously, for each $a \in A$

$$\begin{aligned} EV_{stan}(a) &= \sum_{(x,v) \in X \times \mathcal{V}} a(x)Pr(v)v(\underbrace{p_a}_{p_{a,X \times \{v\}}}) = EV_{\mathcal{I}}(a) \text{ for } \mathcal{I} = \{X \times \{v\} : v \in \mathcal{V}\} \\ EV_{rev}(a) &= \sum_{(x,v) \in X \times \mathcal{V}} a(x)Pr(v)v(\underbrace{p_x}_{p_{a,\{x\} \times \mathcal{V}}}) = EV_{\mathcal{I}}(a) \text{ for } \mathcal{I} = \{\{x\} \times \mathcal{V} : x \in X\} \\ EV_{post}(a) &= \sum_{(x,v) \in X \times \mathcal{V}} a(x)Pr(v)v(\underbrace{p_{x,v}}_{p_{a,\{(x,v)\}}}) = EV_{\mathcal{I}}(a) \text{ for } \mathcal{I} = \{\{(x,v)\} : (x,v) \in X \times \mathcal{V}\}, \end{aligned}$$

where on each line the two identities use Theorem 1 and Proposition 1, respectively. $\blacksquare$

The proof of Theorem 3 begins with a lemma.

Lemma 1 *A valuation $v \in \mathcal{V}$ is vNM if and only if it evaluates value prospects by their expectation, i.e., $v(p) = Exp(p)$ $(= \sum_{k \in \mathbb{R}} p(k)k)$ for all value prospects p.*

Proof. 1. First, let $v \in \mathcal{V}$ be vNM. We fix a value prospect p and prove that $v(p) = Exp(p)$. Pick an option $a \in A$ such that $p_{a,v} = p$. We have

$$Exp(p) = \sum_{k \in \mathbb{R}} kp(k) = \sum_{k \in \mathbb{R}} k \sum_{x \in X : v(x) = k} a(x) = \sum_{k \in \mathbb{R}} \sum_{x \in X : v(x) = k} a(x)k = \sum_{x \in X} a(x)v(x),$$

where the second equality uses that $p(k) = p_{a,v}(k) = \sum_{x \in X : v(x) = k} a(x)$, and the third and fourth equalities follow by reordering terms. The last expression equals $v(a)$ as v is vNM, which equals $v(p)$ by choice of a.

2. Conversely, assume $v(p) = Exp(p)$ for all value prospects p. We let $a \in A$ and show $v(a) = \sum_{x \in X} a(x)v(x)$. Defining p as $p_{a,v}$, we have $Exp(p) = \sum_{x \in X} a(x)v(x)$,

as in part 1 of the proof. So it remains to show $v(a) = Exp(p)$. This holds because $v(a) = v(p)$ (as $p = p_{a,v}$) and $v(p) = Exp(p)$ (by hypothesis). ■

Proof Theorem 3. We shall use standard measure-theoretic arguments.

1. Assume $Pr(v) = 0$ for all non-vNM valuations $v \in \mathcal{V}$. Fix an option $a \in A$. We show that $EV_{\mathcal{I}}(a)$ is independent of the information partition $\mathcal{I}$. On the set of worlds $X \times \mathcal{V}$, consider the probability distribution P_a (the world prospect of a) and the random variables $\mathbf{x} : X \times \mathcal{V} \to X, (x, v) \mapsto x$ and $\mathbf{v} : X \times \mathcal{V} \to \mathcal{V}, (x, v) \mapsto v$. Combining these variables yields a third variable, $\mathbf{v}(\mathbf{x})$, given by $X \times \mathcal{V} \to \mathbb{R}, (x, v) \mapsto v(x)$ and representing resulting value. The value prospect p_a equals the distribution of the variable $\mathbf{v}(\mathbf{x})$, and so its expectation is $Exp(p_a) = Exp_{P_a}(\mathbf{v}(\mathbf{x}))$. More generally, for any information $I \subseteq X \times \mathcal{V}$ (such that $P_a(I) \neq 0$), the value prospect $p_{a,I}$ equals the distribution of $\mathbf{v}(\mathbf{x})$ conditional on I, and so $Exp(p_{a,I}) = Exp_{P_a}(\mathbf{v}(\mathbf{x})|I)$. Now for any information partition $\mathcal{I}$ (identifiable with the variable mapping (x, v) to $\mathcal{I}(x, v)$),

$$\begin{aligned}
EV_{\mathcal{I}}(a) &= Exp_{P_a}(\mathbf{v}(p_{a,\mathcal{I}})) && \text{by definition} \\
&= Exp_{P_a}(Exp(p_{a,\mathcal{I}})) && \text{by Lemma 1} \\
&= Exp_{P_a}(Exp_{P_a}(\mathbf{v}(\mathbf{x})|\mathcal{I})) && \text{as } Exp(p_{a,\mathcal{I}}) = Exp_{P_a}(\mathbf{v}(\mathbf{x})|\mathcal{I}) \\
&= Exp_{P_a}(\mathbf{v}(\mathbf{x})) && \text{by the law of iterated expectations,}
\end{aligned}$$

where Lemma 1 is applicable as valuations generated by $\mathbf{v}$ (with non-zero probability) are vNM. The last expression for $EV_{\mathcal{I}}(a)$ shows that $EV_{\mathcal{I}}(a)$ is independent of $\mathcal{I}$.

2. Conversely, let $\mathcal{V}$ contain a non-vNM valuation $\tilde{v}$ of probability $Pr(\tilde{v}) \neq 0$. As $\tilde{v}$ is non-vNM, we may pick an option $a \in A$ such that $\tilde{v}(a) \neq \sum_{x \in X} a(x)\tilde{v}(x)$. Denote the information of valuation $\tilde{v}$ by $I = X \times \{\tilde{v}\}$. We construct two information partitions $\mathcal{I}_1$ and $\mathcal{I}_2$ for which $EV_{\mathcal{I}_1}(a) \neq EV_{\mathcal{I}_2}(a)$. Let $\mathcal{I}_1$ and $\mathcal{I}_2$ coincide outside I and be, respectively, maximally coarse or maximally fine within I. So $\mathcal{I}_1 = \mathcal{I}_0 \cup \{I\}$ and $\mathcal{I}_2 = \mathcal{I}_0 \cup \{\{(x, v)\} : (x, v) \in I\}$, for some partition $\mathcal{I}_0$ of $(X \times \mathcal{V})\backslash I$. Thus $EV_{\mathcal{I}_1}(a) = S + S_1$ and $EV_{\mathcal{I}_2}(a) = S + S_2$ where

$$\begin{aligned}
S &= \sum_{(x,v) \in (X \times \mathcal{V})\backslash I} a(x)Pr(v)v(p_{a,\mathcal{I}_0(x,v)})) \\
S_1 &= \sum_{(x,v) \in I} a(x)Pr(v)v(p_{a,\mathcal{I}_1(x,v)})) = \sum_{x \in X} a(x)Pr(\tilde{v})\tilde{v}(p_{a,I}) \\
S_2 &= \sum_{(x,v) \in I} a(x)Pr(v)v(p_{a,\mathcal{I}_2(x,v)})) = \sum_{x \in X} a(x)Pr(\tilde{v})\tilde{v}(p_{a,\{(x,\tilde{v})\}}).
\end{aligned}$$

By Proposition 1, $p_{a,I} = p_{a,\tilde{v}}$ and $p_{a,\{(x,\tilde{v})\}} = p_{x,\tilde{v}}$. So $\tilde{v}(p_{a,I}) = \tilde{v}(p_{a,\tilde{v}}) = \tilde{v}(a)$ and $\tilde{v}(p_{a,\{(x,\tilde{v})\}}) = \tilde{v}(p_{x,\tilde{v}}) = \tilde{v}(x)$. Thus

$$\begin{aligned}
S_1 &= \sum_{x \in X} a(x)Pr(\tilde{v})\tilde{v}(a) = Pr(\tilde{v})\tilde{v}(a) \sum_{x \in X} a(x) = Pr(\tilde{v})\tilde{v}(a) \\
S_2 &= \sum_{x \in X} a(x)Pr(\tilde{v})\tilde{v}(x) = Pr(\tilde{v}) \sum_{x \in X} a(x)\tilde{v}(x).
\end{aligned}$$

So

$$EV_{\mathcal{I}_1}(a) - EV_{\mathcal{I}_2}(a) = S_1 - S_2 = Pr(\tilde{v}) \left(\tilde{v}(a) - \sum_{x \in X} a(x)\tilde{v}(x) \right).$$

As $Pr(\tilde{v}) \neq 0$ and $\tilde{v}(a) \neq \sum_{x \in X} a(x)\tilde{v}(x)$, we deduce $EV_{\mathcal{I}_1}(a) \neq EV_{\mathcal{I}_2}(a)$. ■