

Is Normative Uncertainty Irrelevant If Your Descriptive Uncertainty Depends on It?

Pamela Robinson

(This is the penultimate version; the published version is [here](#).)

Abstract

According to ‘Excluders’, descriptive uncertainty—but not normative uncertainty—matters to what we ought to do. Recently, several authors have argued that those wishing to treat normative uncertainty differently from descriptive uncertainty face a *dependence problem* because one’s descriptive uncertainty can depend on one’s normative uncertainty. The aim of this paper is to determine whether the phenomenon of dependence poses a decisive problem for Excluders. I argue that existing arguments fail to show this, and that, while stronger ones can be found, Excluders can escape them.

1. Introduction

Which kinds of uncertainty matter to what we ought to do? According to ‘Excluders’, only descriptive uncertainty matters. Normative uncertainty is irrelevant. (See, for example, Gracely 1996, Harman 2015, Nissan-Rozen 2015, Hedden 2016, and Weatherson 2014, 2019.) But if being descriptively uncertain sometimes involves being normatively uncertain, then taking descriptive uncertainty into account might sometimes require taking normative uncertainty into account as well. And this could be a serious problem for Excluders.

Podgorski (2020) has argued that those wishing to treat normative uncertainty differently from descriptive uncertainty face a *dependence problem* because descriptive uncertainty can depend on normative uncertainty. MacAskill and Ord (2020) also object to Excluders using a case involving dependence.

The aim of this paper is to determine whether the phenomenon of dependence poses a decisive problem for Excluders. I argue that existing arguments fail to show this. But stronger ones can be found. I define three varieties of dependence—‘evidential dependence’, ‘identity dependence’, and ‘higher-order dependence’—and use each to construct an argument against Excluders.

I then show that Excluders can escape these arguments by adopting a specific account of relevance and by rejecting the standard way of defining ‘normative uncertainty’ as uncertainty about normative facts and ‘descriptive uncertainty’ as uncertainty about descriptive facts. I conclude that the phenomenon of dependence doesn’t provide good grounds to reject the view that only descriptive uncertainty matters to what we do.

2. Includers and Excluders

The debate over whether normative uncertainty matters to what we ought to do has been described in different ways and in different terms. Those who think that normative uncertainty does matter, is ‘relevant’ (e.g., Harman 2015), or bears an important ‘symmetry’ (e.g., Weatherson 2019) to descriptive uncertainty have been called ‘metanormativists’ (e.g., MacAskill 2014), ‘uncertaintists’ (e.g., Harman 2015, Sepielli 2018), ‘supersubjectivists’ (Hedden 2016), and ‘normative subjectivists’ (Podgorski 2020).

I’ll focus on the question of whether normative uncertainty matters in the way that descriptive uncertainty does to what we ought to do. I’ll identify the two sides of the debate by names that draw attention to this distinction, and I’ll use the ‘subjective ought’ as a stand-in for all important subjective normative notions.

The Include View: Both descriptive uncertainty and normative uncertainty are relevant to what we subjectively ought to do.

The Exclude View: Descriptive uncertainty is relevant to what we subjectively ought to do, but normative uncertainty is irrelevant. (Normative uncertainty is ‘excluded’.)

Both sides agree that we have reasons to posit a subjective ought that makes descriptive uncertainty relevant to what we ought to do. ‘Includers’ think these reasons carry over to normative uncertainty as well. ‘Excluders’ think they don’t. Some of the explanations Excluders give are: (i) that you’re to blame for getting the normative facts wrong in ways that you aren’t to blame for getting the descriptive facts wrong (e.g., Harman 2015; Weatherson 2019)—perhaps because normative facts can be known *a priori* (Hedden 2016), (ii) that the Include View licences an objectionable kind of ‘normative fetishism’ or aiming at what’s right for the wrong reasons (e.g., Harman 2015; Weatherson 2019), and (iii) that we’re unlikely to find a coherent Includer theory of subjective ought (e.g., Hedden 2016; Nissan-Rozen 2015).

I won’t compare the overall plausibility of the Include View and the Exclude View here. This paper is not an argument in favor of the Exclude View, and I do not conclude that it’s true. Instead, my aim is to determine whether the phenomenon of descriptive uncertainty *depending on* normative uncertainty could make the Exclude View untenable—even by Excluders’ own lights.

3. Existing Arguments

3.1. Podgorski’s Argument

Podgorski (2020) argues that we sometimes have descriptive uncertainty because we’re normatively uncertain about something, and that, because of this, the Exclude View is false.

He divides up the terrain using different terminology, distinguishing between ‘N-subjectivists’, who hold that we must sometimes appeal to *normatively subjective* norms to account for all subjective oughts, and ‘D-subjectivists’ who think that subjective oughts are fully explained by

descriptively subjective norms (2020, p. 46). The prescriptions of a ‘normatively subjective’ norm ‘depend on an agent’s uncertainty about distinctively normative facts’, while the prescriptions of a ‘descriptively subjective’ norm ‘depend on an agent’s uncertainty about descriptive facts but not on their uncertainty about normative facts (such as what is fundamentally of value or which moral theory is true)’ (2020, p. 45). Podgorski explains that:

[The norm] ‘*My Favourite Theory*’, which instructs an agent to perform whichever action is recommended by the moral theory in which they are the most confident, is an example of an N-subjective norm. An agent’s descriptive uncertainty may matter, on this view, for what they ought to do (since it affects the recommendations of different moral theories), but so does their uncertainty about what the correct moral theory is. (2020, p. 45)

[The norm] *expectational hedonistic utilitarianism*, which says that an agent ought to perform whichever action maximizes the expected pleasure generated by their action, is an example of a D-subjective norm. Its prescriptions depend on the agent’s uncertainty about descriptive facts about pleasure, but not on the agent’s uncertainty about normative facts, including uncertainty about whether pleasure is valuable. (2020, p. 45)

I frame the debate here as that between the Include View and the Exclude View because they say the same things—or similar enough things—in simpler terms. ‘Includers’ are Podgorski’s ‘N-subjectivists’ and ‘Excluders’ are Podgorski’s ‘D-subjectivists’.¹

Podgorski’s argument against Excluders (or D-subjectivists) uses the following two scenarios.

Angel’s Present: Heloise wakes up on the morning of her birthday to a knock at the door. On her doorstep stands an angel, bearing a gift for her. If she opens its gift, the angel informs her warmly, something very good will happen. Heloise knows, we will postulate, that the angel has a perfect understanding of what is good and what is bad,

¹ I’ll argue that it matters a great deal how the notion of ‘relevance’ is understood in the definitions of the Include and Exclude Views, in which case it also matters a great deal how the notion of a norm’s ‘dependence’ on kinds of uncertainty is understood in Podgorski’s definitions. This invites confusion within his framework because the kind of ‘dependence problem’ he raises involves cases in which one’s normative uncertainty depends on one’s descriptive uncertainty—a different kind of dependence from the way in which a *norm* can depend on a kind of uncertainty. Defining the two sides of the debate in terms of relevance helps keep these two dependence relations separate.

and that it is honest in describing its gift. She takes the box and places it on her coffee table, and considers whether to open it. (2020 p. 50)

Devil's Present: Heloise wakes up on the morning of her birthday to a knock at the door. On her doorstep stands a devil, bearing a gift for her. If she opens its gift, the devil tells her, cackling through a toothy grin, something very bad will happen. Heloise knows, we will postulate, that the devil has a perfect understanding of what is good and what is bad, and that it is honest in describing its gift. She takes the box and places it on her coffee table, and considers whether to open it. (2020 p. 51)

Add to each that: Heloise holds the box to her ear and shakes it a little, to try and guess what is inside. She hears an unmistakable squeal—a sound that is made by only two rare creatures—the azzip and the buzzip. Opening the box will, she knows, release one of these two creatures. Now Heloise, it turns out, is an ethics professor, and has narrowed the correct theory of value down to two candidates, Theory A and Theory B, with her credence split evenly between them. According to Theory A, an azzip being released would be good (+1 units value), but a buzzip being released would be very bad (-2 units value). According to Theory B, the opposite is true. And let us suppose that in fact, Theory A is true. (2020 p. 51)

Heloise's descriptive uncertainty about which creatures the boxes contain *depends on* her normative uncertainty about which theory is true. She's only uncertain about which creatures the boxes contain because she's uncertain about which theory is true.

Podgorski explains that this additional information doesn't change what Heloise ought to do. She ought to open the angel's present but not the devil's. However, he argues that no Excluder theory of subjective ought can deliver these verdicts. According to him, a popular Excluder theory of subjective ought is:

Maximize Expected Value (MEV): One ought to perform an action that maximizes expected value, EV, where $EV(\phi) = \sum_i P(O_i|\phi)V(O_i)$, $P(O_i|\phi)$ is the agent's credence in outcome O_i conditional on ϕ , and $V(O_i)$ is the objective value of that outcome, according to the true theory. (2020 p. 55)

Podgorski explains that it follows from MEV that Heloise ought not to open either box. If we assume that Theory A is true,² then:

Heloise’s expected value for opening the angel’s box is 0.5 (her credence that the box contains an azzip) times 1 (the objective value of the outcome if the box contains an azzip) plus 0.5 (her credence that the box contains a buzzip) times -2 (the objective value of the outcome if the box contains a buzzip), or -0.5. Her expected value for leaving the box unopened is 0. Therefore, Heloise ought to leave the angel’s gift unopened. Since her descriptive credences about what is in the boxes are exactly the same in the case of the devil’s gift, she likewise ought to leave that gift unopened. (2020 p. 55)

If this is right, Excluders adopting MEV must represent Heloise’s decisions in Angel’s/Devil’s Present as follows. (We’re still assuming that Theory A is true.)

Angel’s/Devil’s Present

	Box contains an azzip (Pr = 0.5)	Box contains a buzzip (Pr = 0.5)	Expected value
Open	+1	-2	-0.5
Don’t	0	0	0

MEV thus gives the same verdicts about Angel’s Present and Devil’s Present, when they should be different. MEV is one possible Excluder theory. But Podgorski’s thought is that all Excluder theories of subjective ought will make Heloise’s decision in these cases depend on her descriptive uncertainty—that is, her uncertainty about which creatures the boxes contain. So, none of these theories will get the verdict that Heloise ought to open the angel’s present but not the devil’s.

He claims that MEV and other Excluder theories fail because they’re insensitive to the dependence relationship between Heloise’s uncertainty about which creature is in the box and her uncertainty about which moral theory is true. He argues that Includers—and only Includers (2020, Section 3)—can get different verdicts for the two cases by conditionalizing on the truth

² MEV gives the same verdicts if Theory B is assumed.

of the two theories (2020, pp. 57-9). Includers can represent Heloise's decisions in the following way:

Angel's Present

	Theory A is true (Pr = 0.5)		Theory B is true (Pr = 0.5)		Something like expected value
	It's an azzip (Pr = 0.5)	It's a buzzip (Pr = 0)	It's an azzip (Pr = 0)	It's a buzzip (Pr = 0.5)	
Open	+1	-2	-2	+1	+1
Don't	0	0	0	0	0

Devil's Present

	Theory A is true (Pr = 0.5)		Theory B is true (Pr = 0.5)		Something like expected value
	It's an azzip (Pr = 0)	It's a buzzip (Pr = 0.5)	It's an azzip (Pr = 0.5)	It's a buzzip (Pr = 0)	
Open	+1	-2	-2	+1	-2
Don't	0	0	0	0	0

This strategy isn't available to Excluders, because it makes Heloise's normative uncertainty between theories A and B relevant to her decision.

However, Excluders can respond to this argument as follows. The only differences between the Angel's/Devil's Present cases are: *whether the gift is given by an angel or devil* and *what the angel or devil says*. If Heloise ought to open the angel's present but not the devil's, it's because the angel has said that opening the present will be very good, the devil has said that opening the present will be very bad, and Heloise completely trusts them. Excluders can simply represent her decisions like this:

Angel's Present*

	Angel is trustworthy (Pr = 1)	Angel isn't trustworthy (Pr = 0)	Expected value
Open	+1	?	+1
Don't	0	0	0

Devil's Present*

	Devil is trustworthy (Pr = 1)	Devil isn't trustworthy (Pr = 0)	Expected value
Open	-2	?	-2
Don't	0	0	0

On this way of thinking, both Heloise's uncertainty about which creatures the boxes contain and Heloise's uncertainty about normative theories are irrelevant to her decision. She's able to know 'the objective values of the outcomes according to the true theory of value' without needing to know what that true theory of value is, and without needing to rely on her credences about what the true theory of value is. A theory like MEV can use her credences about other descriptive facts (e.g., whether the angel/devil is trustworthy) to get at these objective values.

Podgorski's argument is supposed to work by showing that Excluders always get the wrong verdict about one of the cases because they represent Heloise's decision as relying on her uncertainty about whether her presents contain azzips or buzzips. Because Heloise's uncertainty about this is the same in both cases, Podgorski thinks Excluder theories fail because they can't include information about the *relationship* between Heloise's uncertainty about what the boxes contain and her uncertainty about which normative theory is true. It's true that the only differences between the Angel's and Devil's Present tables above are Heloise's subjective probabilities about which creatures are in the boxes, conditional on which theories are true. But what determines these subjective probabilities isn't the fact that Heloise's uncertainty about what the boxes contain *depends on* her uncertainty about which theory is true. They're determined by Heloise's credences about what the angel and devil have said to her. Both Excluders and Includers can rely on this to represent Heloise's decisions, so there's no need for either to adopt Podgorski's solution. The Angel's and Devil's Present* tables seem to represent Heloise's decisions for what they are: simple decisions made under certainty that don't require factoring in her uncertainty about creatures or theories.

One might object that MEV isn't a view whose recommendations are relative to how we decide to draw decision tables. It shouldn't be possible to get different recommendations just by partitioning descriptive states in a different way. Since Podgorski hasn't made a calculation error

in drawing his table, how could he be mistaken? In response, Excluders can claim that Podgorski did make an error in drawing the table for proponents of MEV: he didn't just partition things differently; he also left out important descriptive information.

To see this, it helps to compare Podgorski's cases to the following parallel versions that don't involve any normative uncertainty:

Friend's Pie: Your friend offers you a pie. She tells you it's from the new bakery in town. You know your friend would usually only offer you a pie that you'd enjoy—at least, 90% of the time. She'd rarely play a trick on you.

Enemy's Pie: Your enemy offers you a pie. She tells you it's from the new bakery in town. You know your enemy would usually only offer you a pie that you'd hate—at least, 90% of the time. She'd rarely do something nice.

Add to each that: An avid pie-lover, you've been to the new (and rather unusual) bakery and recognise that the pie being offered is either an azzi pie or a buzzi pie. You know that one is the prank version of the other: one is filled with delicious stew (+1 units of value); the other is filled with exploding ink (-2 units of value). You know that the names of the pies are the key to which is which, since the bakery follows one of two rules. According to Rule A, the pie with the name first in alphabet is the regular one and the other is the prank version. According to Rule B, the pie with the name second in the alphabet is the regular one and the other is the prank version.

In these Pie cases, the additional information doesn't involve any normative uncertainty. In both cases, you're descriptively uncertain about whether the pie is an azzi pie or a buzzi pie. But this isn't due to any normative uncertainty—it's just because you don't know which rule the bakery follows. The other difference from the Present cases is that you can't be *certain* that your friend(enemy) would give you a pie you'd enjoy(hate). Since Heloise can be certain that the angel(devil) would give her a good(bad) present, it might seem that it's only this certainty that gives the Excluders a way to respond.

Your decisions can be represented like this:

Friend's Pie

	Bakery follows Rule A (Pr = 0.5)		Bakery follows Rule B (Pr = 0.5)		Something like expected value
	azzi pie (Pr = 0.45)	buzzi pie (Pr = 0.05)	azzi pie (Pr = 0.05)	buzzi pie (Pr = 0.45)	
Eat	+1	-2	-2	+1	+0.7
Don't	0	0	0	0	0

Enemy's Pie

	Bakery follows Rule A (Pr = 0.5)		Bakery follows Rule B (Pr = 0.5)		Something like expected value
	azzi pie (Pr = 0.05)	buzzi pie (Pr = 0.45)	azzi pie (Pr = 0.45)	buzzi pie (Pr = 0.05)	
Eat	+1	-2	-2	+1	-1.7
Don't	0	0	0	0	0

But again, this is an unnecessarily complex representation. Your decisions can also be described as follows (with or without the precise values):

Friend's Pie*

	Friend is being nice (Pr = 0.9)	Friend is being mean (Pr = 0.1)	Expected value
Open	+1	-2	+0.7
Don't	0	0	0

Enemy's Pie*

	Enemy is being nice (Pr = 0.1)	Enemy is being mean (Pr = 0.9)	Expected value
Open	+1	-2	-1.7
Don't	0	0	0

In the Pie cases, the additional information about kinds of pies and rules is ultimately irrelevant to the decision. Likewise, what Heloise subjectively ought to do doesn't rely on her uncertainty about creatures or theories.

Still, why is it wrong to draw the table as follows?

Friend's/Enemy's Pie

	It's an azzi pie (Pr = 0.5)	It's a buzzi pie (Pr = 0.5)	Expected value
Open	+1	-2	-0.5
Don't	0	0	0

This is analogous to how Podgorski thinks a proponent of MEV would have to draw the decision tables for the Present cases. It's wrong because it leaves out the important descriptive information about what kind of people Friend and Enemy are or what their intentions are.³ Similarly, the Excluder can say, Podgorski's account of how a proponent of MEV would draw the decision table is wrong because it leaves out the important descriptive information about both what the angel and devil each say and Heloise's certainty that they speak the truth.

A final objection is that Podgorski left this out for a good reason: talking about whether the angel and devil are trustworthy is just a roundabout way of talking about whether the presents are good or bad. The fact that the angel and devil always speak the truth seems to be a straightforwardly descriptive fact. But this objection is right about one thing: the fact that they are trustworthy combined with what they say does at least *imply* some normative facts—for example, that something very good(bad) will happen if Heloise opens her present. Still, it's reasonable to wonder why this isn't an available way out for Excluders.⁴ If appealing to the Angel's and Devil's Present* tables isn't a way of ensuring that normative uncertainty is irrelevant to what one subjectively ought to do, then more must be said about relevance.⁵

³ And if this seems like normative information—for example, Friend is 'nice' and Enemy is 'mean'—it should be possible to redescribe it in purely descriptive language. A similar move is also available for Podgorski's angel and devil.

⁴ One might object that the values that go into the Angel's and Devil's Present* tables are sensitive to Heloise's beliefs about moral theories. For example, suppose that Heloise changes her normative beliefs so that, instead of Theory A, she has a 0.5 credence in Theory C, on which azzips are twice as good as they are on Theory A. This will change the expected value of opening the present. However, the same applies to the values that go into the decision table that Podgorski thinks MEV theorists must adopt. So, on his view this type of sensitivity must be available to Excluders. (Thanks to an anonymous reviewer for this objection.)

⁵ Or, to put things in Podgorski's terms: if appealing to the Angel's and Devil's Present* tables isn't a way of ensuring that none of the norms we must appeal to (to account for subjective oughts) depend on this normative uncertainty, then more must be said about this kind of dependence. (Note that this kind of dependence involves norms and states of uncertainty as opposed to states of uncertainty and other states of uncertainty, which is why I've opted to call it 'relevance'.)

3.2. MacAskill and Ord's Argument

MacAskill and Ord (2020) describe a scenario in which someone is uncertain about what kind of uncertainty she has. Because facts about what kind of uncertainty one has are descriptive facts, this is another case in which one's uncertainty about a descriptive fact depends on one's normative uncertainty—or potential normative uncertainty—about something.

Chicken: Sophie is deciding whether to eat chicken. She's certain that she ought not to eat an animal if that animal is a person, but she's uncertain whether chickens are persons or not. And she has no idea whether her uncertainty stems from descriptive uncertainty about chickens' capacity for certain experiences, or from normative uncertainty about what attributes qualify one as a person in the morally relevant sense. (2020, p. 339; paraphrased slightly)⁶

MacAskill and Ord use this case to make the following objection to the Exclude View:

Suppose we think that normative and descriptive uncertainty should be treated in different ways. Then, if Sophie's uncertainty stems from descriptive uncertainty, that uncertainty should be taken into account, and she (subjectively) ought not to eat the chicken. And if her uncertainty stems from normative uncertainty and normative and descriptive uncertainty should be treated differently, then it might be that she ought to eat the chicken. But then, because finding out the nature of her uncertainty could potentially completely change her decision, she should potentially invest significant resources into finding out what the nature of her uncertainty is. This seems bizarre. (2020, pp. 339-40; paraphrased slightly)

Their complaint can be put as follows. According to the Exclude View, whether Sophie's uncertainty is normative or descriptive can make a difference to what she subjectively ought to do. Because of this, Sophie ought to try to determine what kind of uncertainty she really has. But this doesn't seem right.

Here are three ways Excluders might respond. First, they can point out that many Includers will think that what you subjectively ought to do depends on the precise *degrees* of uncertainty

⁶ I've named their case. The same example and objection appears in MacAskill, Bykvist, and Ord 2020, pp. 49-50.

you have, and that, because of this, it's sometimes worth investing significant resources into find out what these precise degrees are. If it's true in general that it's sometimes worth investing resources into discovering things about one's uncertainty, then it's less troubling that it might sometimes be worth investing resources into discovering what kind of uncertainty one has.

Second, Excluders can claim that, while the kind of uncertainty you have matters to what you ought to do, it doesn't follow from this that you ought to try to determine which kind you have. Because Excluder theories of subjective ought are sensitive to the normative truth (about what you subjectively ought to do) as opposed to your state of normative uncertainty, finding out whether you're normatively or descriptively uncertain won't help *you* figure out what you subjectively ought to do. The kind of uncertainty you have will only matter to *what you actually subjectively ought to do*—something you may be unable to know if you're normatively uncertain. From the perspective of the Include View, this will seem odd. Includers may think that there's a sense of subjective ought that you can have access to, and so, if what you subjectively ought to do depends on which kind of uncertainty you have, this is a reason for you to investigate. But on the Exclude View, this fact may not give you a reason to investigate.

Finally, Excluders can claim that, even if it follows from their view that those who are uncertain about what kind of uncertainty they have ought sometimes to try to determine what kind of uncertainty they have—and even if this is a bit odd—this isn't a decisive argument against the Exclude View.

4. The Dependence Problem

I've argued that neither Podgorski nor MacAskill and Ord succeed in showing that the Exclude View must be rejected on grounds of dependence. Nonetheless, I think the phenomenon of dependence does pose a serious challenge. Excluders want to separate states of uncertainty into two groups and to maintain that only those in one group matter. But the phenomenon of dependence seems similar to the phenomenon of entangled particles. Just as the properties of two entangled particles are connected, the relevance of two states of uncertainty related by

dependence is also connected. If a state of descriptive uncertainty depends on a state of normative uncertainty, then it seems that the state of descriptive uncertainty is relevant to what one subjectively ought to do if and only if the state of normative uncertainty is relevant. The serious problem posed by dependence isn't about whether the Exclude View is implausible. It's about whether the Exclude View is even coherent.

4.1. Relevance

The Include View and the Exclude View are views about which kinds of uncertainty are *relevant* to what we subjectively ought to do. In the next set of arguments, everything rests on whether particular states of uncertainty are relevant or not. So, we need a definition of 'relevance'.

It might seem that it's obvious what it means. Perhaps a kind of uncertainty, *K*, is relevant to a theory of ought, *T*, just in case whether or not a person is *K*-uncertain about something can change what she ought to do according to *T*. Call this '**weak relevance**'.

But this isn't the kind of relevance at issue here. For if it were, then most theories of *objective* ought would count as sensitive to uncertainty. For example:

Earning to Donate: I want to do the most good for the world. Becoming an investment banker, making lots of money, and donating most of it to charity is the most effective way of doing the most good. I objectively ought to plan to become an investment banker (since, e.g., it would maximize moral value).

Earning to Donate (with Uncertainty): The same case, except that I face a bunch of uncertainty. I'm uncertain about what I really want and whether earning to donate is the most effective way to do the most good. Given my uncertainty, if I were to become an investment banker, I'd become depressed and would stop caring about others. I wouldn't end up doing much good. I objectively ought not to become an investment banker (since, e.g., it wouldn't maximize moral value).

By hypothesis, the only difference between these two scenarios is the fact that I face uncertainty in the second but not in the first. Since the theory of objective ought *you ought to do what maximizes*

value gives different verdicts about them, this uncertainty counts as being weakly relevant to what I objectively ought to do (according to this theory of objective ought).

I've described a scenario in which one's descriptive uncertainty can change the objective value of one's actions. Likewise, one's normative uncertainty can change the expected values of one's actions. For example, if you're not normatively uncertain, then the expected value of reading a book on normativity might be zero. You won't learn anything new. But if you are normatively uncertain, the expected value could be greater than zero, since you may learn enough to make more rational or moral choices in the future.

Normative uncertainty can be weakly relevant to Excluder theories of subjective ought like MEV. Therefore, this kind of relevance can't be what distinguishes Excluder theories from Includer ones.⁷

To locate the kind of relevance at issue, we can look to how common theories of subjective ought treat descriptive uncertainty. Consider a version of Parfit's Miners Case:⁸

Extra Info Miners Case: Floodwater is rising in mineshafts A and B. Ten miners are trapped in one of the shafts—it isn't known which. You can block shaft A, diverting water to shaft B and saving the miners if they're in shaft A, but killing them if they're in shaft B. You can block B, diverting water to shaft A and saving the miners if they're in shaft B, but killing them if they're in shaft A. You can also do nothing, allowing the water to partially fill each shaft, ensuring that only the lowest miner is killed. You also know that one of the mineshafts is old and that the other is new. You don't think this changes anything about the floodwaters or how many people will die, but you're aware that you can't remember which is old and which is new.

Assume that the miners are in shaft A, but that you have a credence of 0.5 that they're in shaft A and a credence of 0.5 that they're in shaft B. The objective value of blocking shaft A is 10

⁷ Weatherson, a paradigm Excluder, makes a related point (2019, p. 8), noting that "sometimes normative beliefs change the normative significance of other actions". He refers to this as the denial of a 'strong independence claim', and it could easily be extended to cover cases of normative uncertainty. See also Sepielli (2012, p. 47) and Smith (2010).

⁸ Refer to Parfit 1988, pp. 2-3, for the original.

(10 lives saved), the objective value of blocking shaft B is 0 (0 lives saved), and the objective value of blocking neither shaft is 9 (9 lives saved). But, according to a theory of subjective ought like MEV, you subjectively ought to block neither shaft. That's because the expected value of blocking neither shaft is 9, whereas the expected value of blocking either shaft is 5.

We can extract two things from this case about the kind of relevance at issue. First, your uncertainty about which mineshaft is old and which is new is irrelevant to what you ought to do. It doesn't determine, or even partly determine, what you objectively or subjectively ought to do. Second, your uncertainty about which shaft the miners are in is irrelevant to what you objectively ought to do, but is relevant to what you subjectively ought to do. It prevents you from knowing what you objectively ought to do, but doesn't prevent you from knowing what you ought to do according to a theory of subjective ought like MEV.

Given this, here's an account of relevance:

Relevance: your state of uncertainty, *S*, is **relevant** to a theory of subjective ought, *T*, just in case the following conditions are met:

Positive Condition: *S* determines, or partly determines, what you ought to do according to *T*.

Negative Condition: *S* can't prevent you from knowing what you ought to do according to *T*.

A gloss on the negative condition is that, in *S*, you're still in the position to know what you ought to do according to *T* (so long as you don't have other epistemic limitations). For example, your uncertainty about which mineshaft is old and which is new in the Extra Info Miners Case meets the negative condition for relevance to a theory of subjective ought like MEV—and likely to theories of objective ought as well—because nothing morally relevant turns on the age of the mineshafts. Your uncertainty about where the miners are fails the negative condition for relevance to a theory like *you ought to maximize lives saved*, because it does prevent you from knowing what you ought to do (which shaft to block) according to this theory. But your uncertainty about where the miners are does meet the negative condition for

relevance to a theory of subjective ought like MEV, because it doesn't prevent you from knowing what you ought to do according to it (block neither shaft).

The negative condition, or something like it, is usually what's at issue in debates over the subjective ought, since it's supposed to be *subjective* as opposed to objective. But the positive condition is more important here. The exact formulation of it matters while the exact formulation of the negative condition does not. One gloss on the positive condition is that it's met if (and only if) there's a way to alter S by redistributing your credences in such a way that this would change what one ought to do according to T . Call this the 'counterfactual version of the positive condition', and the resulting account of relevance 'counterfactual relevance'. This is how I'll characterise relevance for now (until section 6.3).

Your uncertainty about which mineshaft is old and which is new doesn't meet the counterfactual version of the positive condition for MEV or theories of objective ought like *you ought to maximize lives saved* because no way of redistributing your credences about which shaft is old and which is new would change what you ought to do according to according to these theories. Your uncertainty about where the miners are does meet the counterfactual version of the positive condition for MEV, because, for example, if you were to have credence 1 that the miners are in shaft A, that would change what you ought to do according to MEV (you ought to block shaft B instead of blocking neither shaft). But your uncertainty about where the miners are doesn't meet the counterfactual version of the positive condition for a theory of objective ought like *you ought to maximize lives saved*, because, for example, your having credence 1 that the miners are in shaft B wouldn't change the fact that you ought to block shaft A according to this objective theory—if we continue to assume that the miners are really in shaft A.

One way to understand Podgorski's argument is as one presupposing a notion of weak relevance. (Or, in his terminology, as assuming that a norm depends on one's state of normative uncertainty if it could give different verdicts to two people who only differ by having different states of normative uncertainty.) If this is how the Exclude View is defined, then Podgorski's

argument succeeds and Excluders cannot adopt the strategy I proposed in section 3.1. However, if it is most charitable to understand Podgorski this way, then we should treat his version of the Exclude View—‘D-subjectivism’—as just *one* version of the view. Most of those on the Excluder side of the debate will want to adopt a more sophisticated position. The goal of this paper could then be seen as an attempt to find a more sophisticated version of the view that’s attractive to most of those on the Exclude side of the debate. And the first step is for Excluders to adopt a more flexible notion of ‘relevance’, like the one(s) I propose here.

4.2. Dependence

What is it for a state of descriptive uncertainty to depend on a state of normative uncertainty?

In general, what’s meant is:

Evidential dependence: One’s uncertain about a descriptive fact d and a normative fact n , and d and n aren’t probabilistically independent.

In addition to this ‘standard’ kind of dependence, I’ll also consider a special kind of evidential dependence:

Identity dependence: One’s uncertainty about a descriptive fact is just another way of expressing or representing one’s uncertainty about a normative fact.

As well as a related, but distinct, kind of dependence:

Higher-order dependence: One’s uncertainty about a descriptive fact is about whether one is normatively uncertain about something.

How might dependence pose a problem for Excluders? If my descriptive uncertainty is relevant to what I subjectively ought to do, and this descriptive uncertainty depends on normative uncertainty, then it’s plausible that my normative uncertainty is relevant if and only if the descriptive uncertainty is relevant. And if my normative uncertainty ends up being relevant to

what I subjectively ought to do according to Excluder theories themselves, the Exclude View can't be true.

5. Three New Dependence Arguments

5.1. The Identity Dependence Argument

Consider the following futuristic scenario:

Superintelligent AI: Tia must decide whether to allow doctors to inject an intelligence serum into her child's brain. It's guaranteed to work and will double her child's IQ, but it must be done now, before her child has learned to talk. Tia's unsure whether it's morally permissible to make this kind of enhancement to her child, and her credences are split between two moral theories. According to Theory A, she should allow the injection. According to Theory B, she shouldn't. Also, a superintelligent AI has just been created. Tia believes that it knows everything and will answer any question truthfully. She thinks: *If only I could ask it which moral theory is true! If it answers, 'Theory A', I should allow the injection. If it answers, 'Theory B', I shouldn't.*

It's possible to describe Tia as being uncertain about whether to allow the injection because she's descriptively uncertain about how the AI would answer. And she's descriptively uncertain of this because she's normatively uncertain between moral theories—because which theory is true bears evidentially on what the AI will say. So, Superintelligent AI is a case of evidential dependence. But which theory is true isn't just evidentially relevant to what the AI would say. Tia's normative uncertainty about which theory is true is *disguised as* descriptive uncertainty about what the AI would say.

This is achieved by finding a clever way to say that n is true, where n is some normative fact. Normally, ' n is true' expresses a normative fact if n does. But, by saying something like ' n is written in book B ' (where it's known that book B contains only true statements), it looks like we've found a way to express a normative fact by expressing a descriptive fact.

I call the kind of (evidential) dependence in Superintelligent AI ‘identity dependence’. The simplest method for constructing cases of identity dependence is:

The Truth-Substitute Method: An agent is normatively uncertain about whether n is true. Find a way to say that n is true by saying that n is F , where ‘ n is F ’ expresses a descriptive fact. Now the agent is descriptively uncertain about whether n is F , and it’s just another way of expressing or representing her normative uncertainty about whether n is true.

Here’s the argument from identity dependence for the claim that there’s a dependence problem for Excluders. A case like Superintelligent AI can be described as a decision under normative uncertainty, and as a decision under descriptive uncertainty:

Normative Uncertainty Representation

	Theory A is true	Theory B is true
Allow injection	Good	Bad
Don’t allow	Bad	Good

Descriptive Uncertainty Representation

	AI would answer ‘Theory A’	AI would answer ‘Theory B’
Allow injection	Good	Bad
Don’t allow	Bad	Good

Because both representations appear to accurately describe the same case, no theory of subjective ought can give different verdicts about each. Facts about what one ought to do don’t depend on how one draws decision tables—at least, not if the representations are equally accurate.

And it looks like Excluders must apply something like MEV to the descriptive uncertainty representation. If, for example, we suppose that Tia is confident enough that the AI would answer ‘Theory A’, then she subjectively ought to allow the injection. But then Excluders make Tia’s normative uncertainty about which theory is true relevant to what she subjectively ought to do. It meets the negative condition on relevance because Tia’s normative uncertainty won’t

prevent her from knowing what she subjectively ought to do. And it meets the positive condition on relevance because we can change what Tia subjectively ought to do by changing her credences in theories A and B.

So, the conclusion of the Identity Dependence Argument is that the Exclude View is inconsistent because it says that Tia's normative uncertainty is irrelevant to what she subjectively ought to do, but Excluder theories make it relevant. The Exclude View is also incoherent if it gives conflicting verdicts about the same case depending on how the case is represented.

5.2. The Evidential Dependence Argument

Podgorski's Present cases involve straight evidential dependence. Heloise is descriptively uncertain about whether the box contains an azzip or a buzzip because she's normatively uncertain between theories A and B—which theory is true bears evidentially on which creature the box contains. But Heloise's descriptive uncertainty about what's in the boxes isn't relevant to what she subjectively ought to do because it doesn't meet the positive condition for relevance. We can change Heloise's credences about theories A and B (or her credences about which creatures are in the boxes) without changing what she subjectively ought to do.

To find an effective case of evidential dependence, the agent's uncertainty about the descriptive fact must be relevant to what that agent subjectively ought to do. And, though he doesn't use them for this, some of Podgorski's other examples of dependence are better-suited for the job. Here's a case based on his example of a virtuous friend (2020, p. 50):

Saintly Friend: An urgent message arrives for my friend Zoe and I must deliver it as soon as possible. It's Saturday, and Zoe spends her Saturdays volunteering for what she judges to be the worthiest cause. I know she's a wiser and better person than I, and I'm uncertain between two moral theories. According to Theory A, it's most valuable to relieve the suffering of the poor. If it's true, Zoe's probably at the homeless shelter. According to Theory B, it's most valuable to build and maintain relationships. If it's true, Zoe's probably helping her new neighbours move in. Zoe's neighbours' house and the

shelter are in opposite directions. I can only get to one on time. (I can't call Zoe or gather more information.) I have more credence in Theory B. But Theory A is actually true.

In this case, I'm uncertain about where to go because I'm uncertain of a descriptive fact: where Zoe is. And I'm uncertain about this because I'm uncertain about which normative theory is true. Which theory is true bears evidentially on where Zoe is likely to be. We can represent the case in two ways:

Normative Uncertainty Representation

	Theory A is true	Theory B is true
Go to shelter	Good	Bad
Go to Zoe's neighbours'	Bad	Good

Descriptive Uncertainty Representation

	Zoe's at the shelter	Zoe's at her neighbours'
Go to shelter	Good	Bad
Go to Zoe's neighbours'	Bad	Good

According to an Excluder theory like MEV, I subjectively ought to go to the shelter if I'm more confident that Zoe's there. But I'm uncertain about where Zoe is because I'm normatively uncertain between theories A and B. This normative uncertainty meets the negative condition for relevance. It doesn't prevent me from being in a position to know what I ought to do according to a theory like MEV. It also meets the positive condition for relevance. If I were more confident in Theory B, then it would be that I subjectively ought to go to the shelter. So, the Exclude View is incoherent because it says that my normative uncertainty is irrelevant to what I subjectively ought to do, yet paradigm Excluder theories make it relevant.

5.3. The Higher-Order Dependence Argument

Scenarios like MacAskill and Ord's case of Sophie involve higher-order dependence. In Chicken, Sophie's uncertain about whether her uncertainty (about whether chickens are persons) is descriptive or normative. MacAskill and Ord aim to show that these cases commit Excluders to something implausible. But can they be used to pose a more serious challenge?

Before getting to the argument from higher-order dependence, it's important to show that higher-order uncertainty itself isn't a problem for Excluders. Consider Kagan's higher-order version of Parfit's Miners Case:

Higher-Order Miners Case: We sometimes face mining disasters. As a precaution, we periodically examine possible disaster scenarios and calculate expected utilities for different alternatives. I analyzed a scenario this morning, and now we find ourselves in it. So, for each of our options, I earlier today calculated the expected utility. However, I've now forgotten much of what was calculated. I face three choices, A, B, and C, and remember this: one of A or B had the highest expected utility (+20); the other had a low expected utility (+2). C, meanwhile, was at neither extreme, but had a quite good expected utility (+18). What ought I to do? (2018, p. 155; paraphrased.)⁹

According to Kagan, he should choose C because it has something like the highest 'expected expected value'. An Excluder theory of subjective ought that captures this is:

Maximize Expected Expected Value (MEEV): One ought to perform an action that maximizes expected expected value, EEV, where $EEV(\phi) = \sum_i P(O_i | \phi) EV(O_i)$, $P(O_i | \phi)$ is the agent's credence in outcome O_i conditional on ϕ , and $EV(O_i)$ is the expected value of that outcome, according to MEV.

A theory of subjective ought like MEEV makes Kagan's higher-order uncertainty about the expected utility of his choices relevant to what he subjectively ought to do.

Higher-order uncertainty about what kind of uncertainty one has is different from higher-order uncertainty about the expected values of one's options. But both are kinds of descriptive uncertainty. ('I'm descriptively uncertain whether p ' and 'I'm normatively uncertain whether q ' both express descriptive facts about my mental state.)

⁹ Spencer and Wells (2019) present a similar case (The Fire) and propose extending decision theory to accommodate 'expected expected utility'.

Here's the argument from higher-order dependence. We can represent Sophie's decision in Chicken as follows, assuming that her credences about what kind of uncertainty she has favor normative uncertainty (credence of 0.75) over descriptive uncertainty (credence of 0.25):

Chicken

	States		First-order theory has different results when:		Results of higher-order theory
	Chicken = person (Pr ₁ = 0.5)	Chicken ≠ person (Pr ₁ = 0.5)	'Pr ₁ ' represents descriptive uncertainty (Pr ₂ = 0.25)	'Pr ₁ ' represents normative uncertainty (Pr ₂ = 0.75)	
Eat	-10	+2	X	Y	Z
Don't	0	0	0	0	0

Decision tables can serve two purposes. One is to explain what the agent ought to do and why. Another is to guide the agent. The table is often imagined to be something the agent could complete to determine what to do. If an agent's uncertainty about something is relevant to what the agent subjectively ought to do, then one way of understanding the negative condition on relevance is that: the uncertainty won't prevent the agent from completing a decision table to the degree required to determine what she subjectively ought to do.

There's a problem for Excluders if no plausible Excluder theory allows Sophie to fill in this table enough to find the value of Z. If we assume that chickens aren't actually persons, Excluders adopting a theory like MEEV should complete the table as follows:

Eat	-10	+2	X = -4	Y = +2	Z = +0.5
-----	-----	----	---------------	---------------	-----------------

According to Excluders, the expected value of eating chicken if Sophie is descriptively uncertain about whether chicken are persons is different from the expected value if Sophie is normatively uncertain. If she's normatively uncertain, only the truth about whether chickens are persons is relevant. Given Sophie's uncertainty about what kind of case she's in, the rule MEEV requires that she determine the weighted sum of the two possible expected values, and then perform

the action with the highest resulting values. In this case, it's eating the chicken, with an expected expected value of 0.5.

However, Sophie doesn't know what the expected value of eating chicken would be if hers is a case of normative uncertainty. She might know that it's either -10 or +2, but that's all. She can't fill in the values of Y or Z. The Excluder theory MEEV thus doesn't make Sophie's higher-order descriptive uncertainty relevant to what she subjectively ought to do. And this is because it doesn't make Sophie's (possible) normative uncertainty about whether chickens are persons relevant to what she subjectively ought to do. This is a problem for Excluders because Sophie's higher-order uncertainty does seem relevant to what she subjectively ought to do, and it seems that the Exclude View should make it relevant because it's descriptive uncertainty.

6. Responding to the Dependence Arguments

Excluders have two ways to resist these arguments. One is to reject the counterfactual version of the positive condition on relevance. Another is to drop the common assumption that normative uncertainty is *uncertainty about a normative fact* and descriptive uncertainty is *uncertainty about a descriptive fact* (e.g., Hedden 2016, p. 102; Hicks 2018, p. 162).

To block all three arguments, Excluders must do both. To get around identity and higher-order dependence arguments, Excluders need different definitions of normative and descriptive uncertainty. To get around straight evidential dependence arguments, Excluders need a different gloss for the positive condition on relevance.

6.1. Splitting Definitions

Excluders can block two of the three dependence arguments by adopting the following definitions, which I'll call '**splitting definitions**':

You're **normatively uncertain** about p when your uncertainty either is, or is the result of, uncertainty about a fundamental normative fact.

You're **descriptively uncertain** about p when your uncertainty isn't, and isn't the result of, uncertainty about a fundamental normative fact.

It might seem odd to say that your uncertainty about p can be normative even if p is a descriptive fact, or that your uncertainty about q can be descriptive even if q is a normative fact. But my aim is just to show what these definitions can do for Excluders.

6.2. Blocking Identity Dependence Arguments

The Splitting Definitions can block arguments against the Exclude View from cases like Superintelligent AI if Tia's uncertainty about whether to allow the injection counts as normative.

I'll assume that we can distinguish normative from descriptive facts. I'll also assume that we have some idea how to distinguish fundamental from non-fundamental normative facts. To adopt an example from Hicks (2018, p. 163), you and I might be uncertain about whether to tell a lie. I'm uncertain because I don't know if I can pull it off. You're uncertain because you don't know if lying is intrinsically wrong. My uncertainty about whether I ought to tell the lie doesn't involve my being uncertain about a fundamental normative fact; yours does.

Here's a possible test for reliance on fundamental normative facts:

Test 1: Are You Uncertain About a Fundamental Normative Fact?

Would your uncertainty about whether p go away if you were informed of all descriptive facts? If not, you're uncertain about a fundamental normative fact.

But if Tia were informed of all the descriptive facts (including the fact about what the superintelligent AI would say), she'd no longer be uncertain about whether to allow the injection. This test and the splitting definitions imply that Tia isn't normatively uncertain about anything, which is absurd.

A better test has another part:

Test 2: Are You (Still) Uncertain About a Fundamental Normative Fact? If you're uncertain of any normative facts n (other than p —where p is the same proposition considered in Test 1), what descriptive facts would you no longer be uncertain about if you were to be certain of these? Remove all these descriptive facts from consideration. Of the remaining descriptive facts, would your uncertainty about whether p be resolved if you knew them all? If not, you're uncertain about a fundamental normative fact.

Since Tia's uncertainty about what the AI would say would disappear if she knew whether to allow the injection, it gets set aside. Plausibly, her uncertainty about whether to allow the injection would remain even if she were to know all the other descriptive facts. According to this second test, she's uncertain about whether to allow the injection because she's uncertain about some fundamental normative fact. Therefore, the splitting definitions classify her as normatively uncertain about whether to allow the injection.

Excluders can thus appeal to the splitting definitions to explain why they don't have to treat the two ways of representing Tia's decision differently. They can hold that if her uncertainty about whether to allow the injection is normative, even the apparently descriptive states describing how the AI might answer get treated like states of normative uncertainty by a rule like MEV. Excluders can maintain that what Tia subjectively ought to do just depends on the moral facts, and that her uncertainty between theories and her uncertainty about what the AI would say are irrelevant.

6.3. Blocking Evidential Dependence Arguments

The splitting definitions and tests classify my uncertainty about what to do in Sainly Friend as descriptive. Test 1: I'm uncertain about whether to go to the shelter or to Zoe's neighbours'. If I were certain of all the descriptive facts, I'd know where Zoe is, so I'd no longer be uncertain about where to go. Test 2: If I were certain of which normative theory were true, then I still wouldn't be completely certain of where Zoe is. No descriptive facts get excluded. If I were informed of all the remaining descriptive facts (including the fact about where Zoe is), I'd no

longer be uncertain about where to go. So, my uncertainty about where to go doesn't involve my being uncertain about any fundamental normative facts.

But my normative uncertainty about which theory is true might still be relevant to what I subjectively ought to do according to Excluder theories of subjective ought (like MEV). My normative uncertainty doesn't prevent me from knowing what I ought to do according to Excluder theories, so it meets the negative condition for relevance. It also meets the counterfactual version of the positive condition. If my credences about theories A and B were different, that could change what I subjectively ought to do, because my credences about where Zoe is would be different.

Excluders, therefore, must object to this understanding of relevance. I suggest that they hold, instead, that the positive condition is met if (and only if) *the best algorithm or decision procedure for determining what one ought to do according to T appeals to S*. Call this a 'procedural version of the positive condition', and the resulting account of relevance 'procedural relevance'.

Procedural relevance does as well as counterfactual relevance at explaining why your uncertainty about which mineshafts are old and new is irrelevant to what you subjectively ought to do in the Extra Info Miners Case. (The best algorithm or decision procedure for determining what you subjectively ought to do in that case won't appeal to this state of uncertainty.) And with procedural relevance, Excluders need only claim that the best decision procedure is an application of an Excluder theory like MEV. Since MEV doesn't appeal to my normative uncertainty about which theory is true—it only appeals to my descriptive uncertainty about where Zoe is—my normative uncertainty in Sainly Friend isn't procedurally relevant.¹⁰

6.4. Blocking Higher-Order Dependence Arguments

¹⁰ Excluders might explain, further, that in cases of evidential dependence, the descriptive fact evidentially *screens off* the normative fact. For example, in Sainly Friend, the probability that I subjectively ought to go to the shelter given that Zoe's there and Theory A is true is the same as the probability that I ought to go to the shelter given that Zoe's at the shelter. Once I know where Zoe is, learning which theory is true provides no further information about where I ought to go.

The argument from higher-order dependence assumes that uncertainty about what kind of uncertainty one has is descriptive uncertainty, since facts about what kind of uncertainty one has are descriptive facts. The splitting definitions call this into question.

In Chicken, Sophie is uncertain about whether she's normatively or descriptively uncertain about whether chickens are persons.

Chicken

	States		First-order theory has different results when:		Results of higher-order theory
	Chicken = person (Pr ₁ = 0.5)	Chicken ≠ person (Pr ₁ = 0.5)	'Pr ₁ ' represents descriptive uncertainty (Pr ₂ = 0.25)	'Pr ₁ ' represents normative uncertainty (Pr ₂ = 0.75)	
Eat	-10	+2	X	Y	Z
Don't	0	0	0	0	0

The higher-order dependence argument aims to show that Sophie can't fill in the values of Y and Z on any Excluder theory of subjective ought, and that, therefore, no Excluder theory makes Sophie's higher-order descriptive uncertainty relevant to what she subjectively ought to do.

Excluders might opt to revise their view to hold that only first-order descriptive uncertainty is relevant to what subjectively ought to be done, and that both normative and higher-order uncertainty are irrelevant. However, while many Excluders would be happy with this (e.g., Weatherson 2019; Lasonen-Aarnio 2014), it's an unnecessary concession. Excluder views according to which descriptive higher-order uncertainty is relevant to what we subjectively ought to do aren't *incoherent*, though they might be rejected on other grounds.

Here's what Excluders can say instead. Sophie is either normatively or descriptively uncertain about whether chickens are persons. If she's normatively uncertain and chickens aren't actually persons, here's how to complete the table:

Eat	-10	+2	X = -4	Y = +2	Z = +2
-----	-----	----	---------------	---------------	---------------

This is because only the normative truth that chickens aren't persons matters. Sophie can't figure out the values of Y and Z, and so she can't figure out what she subjectively ought to do. An Excluder theory of subjective ought like MEEV makes her higher-order uncertainty irrelevant. But that's okay. According to the splitting definitions, her higher-order uncertainty is normative and therefore shouldn't be relevant to what Sophie subjectively ought to do.

Why is her higher-order uncertainty normative? Sophie is uncertain about whether to eat chicken. If she were informed of all the descriptive facts, this uncertainty wouldn't go away. By hypothesis, she's normatively uncertain about whether chickens are persons. So, her higher-order uncertainty also counts as normative uncertainty.

If Sophie is descriptively uncertain about whether chickens are persons, then her higher-order uncertainty about what kind of uncertainty she has is descriptive. If she were informed of all the descriptive facts—including the descriptive fact about what kind of uncertainty she has—she'd no longer be uncertain about what to do. And Test 2 doesn't change this because she isn't normatively uncertain about anything.

The higher-order dependence argument tried to show that, because MEEV couldn't do it, no Excluder theories could make Sophie's descriptive higher-order uncertainty relevant to what she subjectively ought to do. But there are other Excluder theories. Consider the following:

Maximize Expected Expected Value* (MEEV*): In cases of descriptive higher-order uncertainty, one ought to perform an action that maximizes expected expected value, EEV. In all other cases, one ought to perform an action that maximizes expected value, EV.

According to MEEV*, if Sophie is descriptively uncertain about whether chickens are persons, then she can fill in the values of Y and Z. Y is the probability₁-weighted sum of the values of the first two columns if she eats the chicken ($0.5(-10) + 0.5(2) = -4$); Z is the probability₂-weighted sum of the values of the second two columns ($0.25(-4) + 0.75(-4) = -4$).



Eat	-10	+2	X = -4	Y = -4	Z = -4
-----	-----	----	---------------	---------------	---------------

On such a theory, one’s state of higher-order uncertainty is relevant to what one subjectively ought to do so long as it’s descriptive. If one’s first-order uncertainty is normative, then the higher-order uncertainty doesn’t end up being relevant to what one subjectively ought to do, but that’s okay because it also doesn’t end up being descriptive.

6.5. The Outcome for Excluders

I’ve shown that Excluders can block the dependence arguments in section 5. To do so, they must adopt the Splitting Definitions and procedural relevance. There are often costs to wriggling out of a tight spot by re-defining terms. But in this case the commitments may only help to represent the Excluders’ debate with Includers more accurately.

The splitting definitions should be easy to defend because fundamental normative uncertainty is a crucial component of the thought experiments used by Includers and Excluders.^{11,12} Adopting procedural relevance looks more like a concession. The Exclude View ends up being: ‘descriptive uncertainty, but not normative uncertainty, is *procedurally* relevant to what we subjectively ought to do’. Excluders must admit that normative uncertainty can sometimes matter (in some sense) to what we subjectively ought to do, because it can be counterfactually relevant.

Nevertheless, it’s plausible that this characterization of the Exclude View explains what’s really at issue. Excluders might argue that it’s procedural relevance—not counterfactual relevance—that could lead to ‘normative fetishism’ or aiming at what’s right for the wrong reasons. And that it’s procedural relevance—not counterfactual relevance—that causes a

¹¹ See, for example, Harman’s (2015, p. 59) case of Bob. The assumption that ‘the conservative speaker does not challenge any non-moral facts Bob already believes’ ensures that Bob is uncertain about a fundamental normative fact.

¹² Podgorski recognises the importance of fundamental normative facts, although he makes D-Subjectivism (his version of the Exclude View) committed to the standard way of defining normative and descriptive uncertainty (p. 45).

problem of intertheoretic value comparisons¹³ for Includer theories of subjective ought. Excluders can still deny that we ever subjectively ought to hedge our bets morally (or normatively), and claim this is the crux of their disagreement with Includers.¹⁴

Conclusion

‘Excluders’—those who think that descriptive uncertainty, but not normative uncertainty, is relevant to what one (subjectively) ought to do—have the extra challenge of keeping the two kinds of uncertainty cleanly divided. In a case of dependence, one’s state of descriptive uncertainty seems to depend on a state of normative uncertainty in such a way as to make it impossible for the descriptive uncertainty to be relevant to what one ought to do without the normative uncertainty being relevant as well. Cases of dependence thus threaten to make the Excluder’s position untenable—even by the Excluder’s own lights. I’ve shown how challenging the problem is and how, nonetheless, a successful division can be made.

Acknowledgements: Research on this paper was supported by ARC discovery grant DP170101394 and the ANU Humanising Machine Intelligence Grand Challenge project. The author also thanks Katie Steele and Brian Hedden for helpful comments on drafts.

School of Philosophy

Australian National University

¹³ Lockhart 2000, Ross 2006, Sepielli 2009, Gustafsson and Torpman 2014, Hedden 2016, and Hicks 2018 all discuss the problem of intertheoretic value comparisons.

¹⁴ For example, Weatherson 2014, Nissan-Rozen 2015, and Tarsney 2017 all frame the debate in this way.

References

- Gracely, Edward. (1996). 'On the Noncomparability of Judgments Made by Different Ethical Theories,' *Metaphilosophy* 27(3), pp. 327-332.
- Gustafsson, Johan, and Olle Torpman. (2014). 'In Defense of My Favourite Theory,' *Pacific Philosophical Quarterly* 95(2), pp. 159-174.
- Harman, Elizabeth. (2015). 'The Irrelevance of Moral Uncertainty,' in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* 10, pp. 53-79. Oxford: Oxford University Press.
- Hedden, Brian. (2016). 'Does MITE Make Right: On Decision-Making Under Normative Uncertainty,' in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* 11, pp. 102-128. Oxford: Oxford University Press.
- Hicks, Amelia. (2018). 'Moral Uncertainty and Value Comparison,' in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* 13, pp. 161-183.
- Kagan, Shelly. (2018). 'The Paradox of Methods,' *Politics, Philosophy & Economics* 17(2), pp. 148-168.
- Lasonen-Aarnio, Maria. (2014). 'Higher-Order Evidence and the Limits of Defeat,' *Philosophy and Phenomenological Research* 88(2), pp. 314-345.
- Lockhart, Ted. (2000). *Moral Uncertainty and Its Consequences*. Oxford: Oxford University Press.
- MacAskill, William. (2014). 'Normative Uncertainty.' (PhD dissertation, University of Oxford.)
- MacAskill, William and Toby Ord. (2020). 'Why Maximize Expected Choice-Worthiness?' *Noûs* 54(2), pp. 327-353.
- MacAskill, William, Krister Bykvist, and Toby Ord. *Moral Uncertainty*. (2020). Oxford University Press.
- Nissan-Rozen, Ittay. (2015). 'Against Moral Hedging,' *Economics and Philosophy* 31(3), pp. 349-369.
- Parfit, Derek. (1988). 'What We Together Do.' (Unpublished Manuscript.)
- Podgorski, Abelard. (2020). 'Normative Uncertainty and the Dependence Problem,' *Mind* 129(513), pp. 43-70.
- Ross, Jacob. (2006). 'Rejecting Ethical Deflationism,' *Ethics* 116(4), pp. 742-768.
- Sepielli, Andrew. (2018). 'How Moral Uncertainty Can Be Both True and Interesting,' *Oxford Studies in Normative Ethics* 7, pp. 98-116. Oxford: Oxford University Press.
- Smith, Holly. (2010). 'Subjective Rightness,' *Social Philosophy and Policy* 27, pp. 64-110.
- Spencer, Jack and Ian Wells. (2019). 'Why Take Both Boxes?' *Philosophy and Phenomenological Research* 99(1), pp. 27-48.
- Tarsney, Christian. (2017). 'Rationality and Moral Risk: A Moderate Defense of Hedging.' (PhD dissertation, University of Maryland.)
- Weatherson, Brian. (2014). 'Running Risks Morally,' *Philosophical Studies* 167(1), pp. 141-163.
- Weatherson, Brian. (2019). *Normative Externalism*. Oxford: Oxford University Press.