

This is a pre-print. Please cite the published version:

Koon, J. (2020). "Options Must Be External," *Philosophical Studies* 177 (5): 1175-1189.

<<https://doi.org/10.1007/s11098-019-01240-0>>.

Options Must Be External

Justis Koon

1. Options-as-Decisions

Walking through Times Square, you see a man draw a gun and point it into a crowd; unbeknownst to you, the man is an actor shooting a film and the gun is a prop. Ought you to cry out in warning? There is a clear sense of the word "ought" in which the answer is "no", you ought not disrupt the shoot and ruin the take. But there is also a clear sense in which the answer is "yes" – the information available to you suggests that the man is about to carry out a terrible crime, which you ought to do everything in your power to prevent, indeed, it would be reprehensible for you to remain silent. This latter sense of the word "ought" is the subjective "ought" of practical rationality. The subjective "ought" takes into account our epistemic limitations, our imperfect knowledge of the world, when prescribing how we should behave.

A complete account of the subjective "ought" has two components, a decision theory and a theory of options. The theory of options determines what qualifies as an option for an agent and sends its candidates to the decision theory, which advises us how to evaluate and rank those options. The subjective "ought" then prescribes that the agent execute whichever option the decision theory ranked best (in the event of a tie, it directs the agent to choose from among the highest-ranked options). Our focus in what follows will be on the theory of options.

One prominent, recent discussion of options for the subjective "ought" comes from Brian

Hedden (2012). Hedden holds that any successful account of options must ensure that they satisfy three constraints (he calls them desiderata), the ability constraint, the belief constraint, and the supervenience constraint. First, we have the ability constraint, (AC):

(AC): Necessarily, ϕ -ing is an option for an agent only if she is able to ϕ .

Second, the belief constraint, (BC):

(BC): Necessarily, ϕ -ing is an option for an agent only if the agent believes she can ϕ .

And, finally, the supervenience constraint, (SC):

(SC): The options available to an agent strongly supervene on her beliefs and desires, that is, necessarily, ϕ -ing is an option for an agent only if, necessarily, any agent with the same suite of beliefs and desires also has ϕ -ing as an option.

To accommodate the possibility that an agent's options may depend on mental states other than her beliefs and desires, let us weaken this last constraint to (SC)*:

(SC)*: The options available to an agent strongly supervene on her total internal mental state, that is, necessarily, ϕ -ing is an option for an agent only if, necessarily, any agent in the same total internal mental state also has ϕ -ing as an option.¹

¹ In his subsequent book, Hedden (2015a: 22-28) adjusts the constraint to make it compatible with externalism about mental content. He also suggests that the supervenience base for options might extend to knowledge, in order to accommodate Williamson's (2000) thesis that knowledge is a mental state. This suggestion is puzzling: if knowledge is a mental state, my mental states supervene on nothing less than the whole of existence. For instance, whether my attitude that nothing can travel faster than the speed of light counts as knowledge or mere belief depends on the entire state of the universe, past, present, and future. Because supervenience is transitive, this turns (SC)* into the trivial requirement that options supervene on all of existence, which no one would think worthwhile to dispute. Hence, if (SC)* is to have any bite at all, it must be restricted to mental states which are in some sense internal.

According to Hedden, all three of these constraints are necessary so that the subjective “ought” of practical rationality may fulfill its three key functions: guiding action, evaluating agents for praise and blame, and predicting the behavior of ideally rational agents. For the subjective “ought” to carry out its action-guiding role, our account of options must guarantee that the agent is always in a position to know what her options are. For it to properly adjudicate blame, we must specify an agent's set of options in such a way that, should the agent fail to execute her best option, she will deserve rational criticism or blame for her choice. And, for the subjective “ought” to serve its predictive role, our account of options must ensure that the best-ranked option matches our expectation of how an ideally rational agent would behave under those circumstances.²

Hedden criticizes several theories of options for violating one or more of these constraints before alighting on his preferred account, which makes an agent's options a set of propositions corresponding to all and only the decisions available to her. Here is his official statement of the view (2012: 352):

Options-as-Decisions: A set of propositions is a set of options for agent *S* at time *t* iff it is a maximal set of mutually exclusive propositions of the form “*S* decides at *t* to ϕ ”, each of which *S* is able to bring about.

Hedden's case for Options-As-Decisions rests on the claim that it is the only account of options that has the potential to satisfy his three constraints. Unfortunately for Hedden, he is wrong about this, as the set of decisions an agent is able to make does not supervene on her mental states. Fortunately for Hedden, this is not a fatal defect for his theory, as no account of options could possibly satisfy all three constraints. As we will see in Section 2, (AC) and (SC)* jointly entail that no non-godlike agent has any options at any time. Because the ability constraint is indispensable, I will argue,

2 If it seems strange that the subjective “ought” should be in the business of prediction at all, remember that standard decision theories are intended as accounts of ideal rationality, so something has gone seriously wrong if a prescription given by the subjective “ought” fails to coincide with the expected behavior of an ideally rational agent.

this leaves us with no choice but to reject (SC)*. In Section 3 I show that the underlying problem lies with the roles Hedden identifies for the subjective “ought.” All three roles ultimately require that options be transparent, that an agent can never be mistaken about what options are available to her, but transparency also proves to be incompatible with the ability constraint, which means that it, too, must be rejected. These failures of supervenience and transparency, in turn, undermine Hedden's rationale for making options mental entities like decisions, opening the door for a more liberal and intuitive theory which permits overt bodily actions to qualify as options as well. In Section 4 I sketch one such account, which conceives of options as exhaustive combinations of atomic movements (EXCAMs).

The supervenience constraint, as I have formulated it, is the minimal thesis of internalism about options.³ It is a crucial component of a broader internalist program, defended by Hedden (2015a; 2015b) and Moss (2015), which aims to show that all the demands that rationality places on an agent can be cashed out in terms of the agent's mental states, regardless of what goes on in the world around her.⁴ But the requirements of rationality in the practical sphere depend on what options are available to us, which means that, if (SC)* fails and options turn out to be external, the whole internalist enterprise will be sunk. Regrettably, this appears to be the case. An agent's options are not fixed by her mental states alone; they depend also on the cooperation of the external world.

2. Frankfurt Demons Inescapable

Let's begin by reviewing Hedden's motivations for imposing the ability and supervenience constraints. The ability constraint takes its cue from the truism that “ought” implies “can.” No matter how desirable it would be for me to put an end to war this instant, intuitively, it is not the case that I ought to do so or that I have the option of doing so. Similarly, although it would be a great boon to

3 Cf. Conee and Feldman (2001), who argue that strong supervenience on the mental is the *sine qua non* of internalism about justification. A note on terminology: I will call any account of options which satisfies (SC)* “internalist,” any account which does not “externalist,” and any account of options which makes them mental entities (or propositions concerning mental entities) “mentalistic.” This is a departure from convention: typically, only internalist accounts may be described as mentalistic, but options present us with a case where mentalism is insufficient for internalism.

4 See also Broome (2013: 250), who proposes that we should restrict options to intentions on the basis of a supervenience argument.

humanity were I to cure cancer last Tuesday, there is no sense in which I ought to cure cancer last Tuesday, and neither does curing cancer last Tuesday seem like it should number among my options. By the same token, “sweep this year's Nobel prizes,” “assassinate Caesar before he crosses the Rubicon,” and “pop off to Arcturus for a holiday” also do not seem like good candidates for options, or the sort of action the subjective “ought” should be in the business of prescribing. The ability constraint offers a compelling, unified explanation for why this is so: none of these actions are options for me precisely because I do not have the ability to carry any of them out, and the scope of our options is limited by the scope of our abilities.

Reflection on Hedden's three roles for the subjective “ought” also lends support to (AC). If what options we have were not constrained by what abilities we have, the subjective “ought” would fail at guiding action, for it would sometimes instruct us to perform impossible feats. It could not be used to adjudicate blame, as it is patently unfair to criticize an agent for failing to perform an action when she lacks the ability to do so.⁵ And it would be hopeless at prediction, because we should never expect a rational agent to perform actions which exceed her abilities. In sum, the tight conceptual connection between options and abilities seems to give (AC) a stronger claim to our allegiance than just about any other constraint we might wish to impose on options.

The justification for (SC)* is a bit more subtle. Consider the following scenario, adapted from Hedden (2012: 349):

Jane's Doppelganger: Jane faces a raging creek, is in fact able to ford it, and, among the things she is able to do, fording the creek is her best choice. Her doppelganger Twin Jane, who is in the same total internal mental state as Jane, faces her own raging creek, but Twin Jane is unable to ford it. Among the things that Twin Jane is able to do, turning back and heading home is her best choice.

5 We must make an exception here for cases where an agent's inability to perform some option is a foreseeable result of the agent's own prior blameworthy actions or negligence. Suppose, for instance, that Lisa causes an accident at t because she is unable to swerve in time to avoid a pedestrian. It might still be fair to blame her for the accident, if the reason she is unable to swerve at t is because she drank too much at dinner, and then chose to drive while intoxicated.

Suppose we say that Jane has fording the creek as an option, while her twin does not. Then, Hedden claims, neither will be in a position to know what their options are, and the subjective “ought” will be insufficiently action-guiding. What's more, the subjective “ought” will predict that Jane crosses the creek while Twin Jane heads home, when clearly we should expect agents in like mental states to at least begin to perform the same action. Suppose also that both decide to head home. Then Jane will be open to rational criticism for failing to execute her best option, while her twin will not, and this seems unfair – how could two people in the same mental state who perform the same action bear different degrees of blameworthiness? The source of these difficulties seems to be the supposition that Jane and Twin Jane had different options despite occupying the same mental state, and the most obvious fix is to require that pairs of mental duplicates like Jane and her twin always have the same set of options available to them. *Jane's Doppelganger*, Hedden concludes, shows that we must also accept (SC)* as a constraint on options.

A major problem arises for Hedden's constraints when it comes to Frankfurt-style cases.⁶ In Frankfurt cases, an agent is surreptitiously monitored by a would-be Svengali or demon who will manipulate or kill the agent if she deviates from a particular course of action, but the demon never has cause to intervene because the agent perfectly executes the demon's plans of her own accord. *Prima facie*, it seems as though attracting the attention of a Frankfurt demon restricts an agent's options, which calls the supervenience constraint into question, because the presence or absence of a Frankfurt demon does not depend on the agent's mental states alone. Hedden proposes the following strategy for dealing with such cases (2012: 354):

If a Frankfurt demon is monitoring you with an eye toward preventing you from deciding to ϕ , then you lack the capacity to exercise your rational capacities which is necessary in order for you to be subject to the demands of prudential rationality in the first place [...] The rational *ought* thus only applies to agents who are not being disrupted by Frankfurtian demons in this way, and so once we restrict our attention to

⁶ *Locus classicus* is Frankfurt (1969).

agents to whom the rational *ought* applies, which options an agent has will both supervene on her beliefs and desires and be knowable by her.

But Hedden is mistaken in thinking that agents being monitored by Frankfurt demons are immune to the demands of practical rationality. Suppose that I am faced with the following three options (the last of which, as we shall see, is really a pseudo-option):

Option #1: Decide to wisely invest my money. (Best)

Option #2: Decide to buy a dozen lottery tickets. (Intermediate)

Option #3: Decide to jump into a woodchipper. (Worst)

Unbeknownst to me, I am watched over by a guardian angel who cares for my health but is indifferent to my finances, who will intervene just in case she predicts that I will select the third option. She will intervene by directly manipulating my brain before I have a chance to make the disastrous decision, redirecting my synaptic impulses and neurotransmitters so that I deliberate freely between the other two options instead. As it happens, I have no inclination to die horribly, but I do have a bit of a gambling problem, so I opt for #2 and lose a few dollars on the scratch-offs while the angel sits idle.

According to Hedden, the mere presence of the guardian angel serves to absolve me of blame. But this is intuitively the wrong verdict: I ought to have chosen Option #1, and deserve criticism for electing to squander my money on lotto tickets instead. The lesson we should draw from this case is that the presence of a Frankfurt demon eliminates only those options which will trigger the demon's intervention, while leaving us open to rational criticism if we choose poorly from among the remainder. This is already enough to undermine (SC)*, for the third option ceases to be an option for me not because of any internal feature of my mental life but because of extrinsic facts about the world.

Hedden (2015a: 108-109) has since clarified his view on Frankfurt cases. Frankfurt demons now excuse an agent from the requirements of practical rationality only if the demon intends to

intervene in the event that the agent will choose her best apparent option. In such cases, Hedden claims, the subjective “ought” cannot endorse any course of action and so falls silent. On the one hand, it is not true that the agent ought to choose her best apparent option, because she is incapable of doing so. But neither ought she choose her best actual option, because there is a superior apparent option available.

This revision does little to help, however, because an agent divested of her best apparent option by a Frankfurt demon is still acting irrationally should she fail to choose the best option remaining. Imagine that instead of a guardian angel I am watched over by a mischievous imp. The imp is not bloodthirsty, and does not wish on me any grievous bodily harm, but takes great joy in frustrating all of my attempts at self-improvement, so if I seem inclined to spring for Option #1 he will intervene, leaving me with only Options #2 and #3 as genuine possibilities. Nevertheless, it seems as though I would be exhibiting a grave failure of rationality if I decided to pass up the scratch-offs and take a plunge into the woodchipper instead. Even if my best apparent option is unavailable to me, we must still be able to say of my worst remaining options that I ought not choose them. So agents whose best apparent option will trigger the intervention of a Frankfurt demon must be subject to the demands of practical rationality after all. And, just as before, there is no way of reconciling this verdict with the supervenience constraint: in Frankfurt cases, what options an agent has depends on the demon's dispositions, and the demon's dispositions are no part of the agent's mental states.

This way of understanding options in Frankfurt cases does have the unhappy consequence that an agent who freely passes up what appears to be her best option can sometimes escape rational criticism. If I decide to purchase the lottery tickets of my own volition, believing, falsely, that I also had the option of investing my money, we would like to say that I am still blameworthy for my choice. Notice, though, that Hedden's proposal has the strictly worse consequence that an agent whose best apparent option is taken away by a Frankfurt demon can never come in for blame no matter what she does. The ability constraint obliges us to remove whichever actions will trigger the demon's intervention from the agent's set of options. Hence, if we are to retain the constraint, the only question left is whether we should require that the agent ought to choose her best remaining option, or whether

we should instead, as Hedden suggests, absolve her of blame no matter how foolishly she behaves. We must issue the agent an indulgence in either case; better, then, to tailor the scope of the indulgence as narrowly as possible, and insist that the agent at least execute the best option she can.⁷

An even bigger problem for (SC)* is lurking in the vicinity. Curiously, Hedden restricts his attention to cases where the agent herself is being monitored by a Frankfurt demon, but it is Frankfurt cases involving the agent's duplicate which he should have been most worried about. This is because, for all non-godlike agents and all decisions, it will always be possible that the agent has a duplicate sharing her total internal mental state who is incapable of making any of the decisions available to the agent. This might be because of the presence of a Frankfurt demon, but it might also be, more prosaically, because the duplicate is a microsecond away from being flattened by an unseen bus or struck by lightning. (AC) entails that these duplicates have no options, so if (SC)* also holds and we have no options not available to our duplicates, it follows that we never have any options at all.

In fact, (AC) and (SC)* jointly entail that no non-godlike agent ever has any options, however options are conceived. Here is one way of formalizing this argument:

- (1) Plenitude of the Dying: Necessarily, for all agents and all times t , possibly there exists a mental duplicate of that agent at t who ceases to exist immediately after t .
- (2) Impotence of the Dead: Necessarily, for all times t , any mental duplicate who ceases to exist immediately after t cannot exercise any apparent options available to her at t .
- (3) Consequence of (AC): Necessarily, for all times t , any mental duplicate who cannot exercise any apparent options available to her at t has no options at t .

⁷ This resolution is not entirely satisfactory – there still seems to be something amiss about saying that an agent who fails to execute what she takes to be her best option is acting as she subjectively ought to, even if, as it turns out, she could not have done any better. In the next section, we will see that this problem is not confined to Frankfurt cases, and that it is insoluble. What we are really demanding here is that an agent's options always be transparent to her, but if our options depend on what abilities we have, this is impossible, for no non-godlike agent has transparency over her abilities.

As consolation, we might still be able to capture the intuition that the agent who passes up her best apparent option is blameworthy for acting as she does if we accept an account of blameworthiness that divorces it from the agent's abilities (Fischer and Ravizza [1998], writing on moral responsibility, suggests how this might go). However, pairing this approach with a subjective “ought” governed by the ability constraint forces us to say that an agent can be subject to blame even if she executes her best option, that is, even if she behaves exactly as she ought to.

- (4) From (1), (2), and (3): Necessarily, for all agents and all times t , possibly there exists a mental duplicate of that agent with no options at t .
- (5) Consequence of (SC)*: Necessarily, for all agents and all times t , an agent has no options at t if possibly there exists a mental duplicate of that agent with no options at t .
- (6) Conclusion: Necessarily, for all times t , all agents have no options at t .

(1) is a corollary of standard principles of modal plenitude:⁸ for every agent, there is always some possible world containing a mental duplicate of that agent who perishes in the next instant, by squashing, immolation, mountain lion predation, or what have you. This leaves (2) as the only auxiliary premise open to serious challenge. Why should we accept (2)? Because it seems that to exercise an option an agent must execute some kind of action – at minimum, a mental action like forming an intention or making a decision – and no non-godlike agent can execute any action in an instant. For all beings located in time and caught up in the causal machinery of the world, carrying out an option is a temporally extended process, and it is beyond our powers to have an option while simultaneously exercising numerically the same option.⁹

In a different context, Hedden (2015a: 147-149) acknowledges that, on his view, any being unable to make decisions in no time flat falls short of ideal rationality. He is willing to bite the bullet, writing, “Ideally rational agents would not require time to come to satisfy the requirements of rationality. We do require time to do so, but that is because we are only imperfectly rational, even if blameless for our cognitive limitations.” But this is no help with the present argument. For the objection is not that (AC) and (SC)* together guarantee that any non-godlike agent is in some respect irrational, it is that they jointly entail that any agent who cannot execute her options in an instant

⁸ See, e.g., Lewis (1986: 86-92).

⁹ Strictly speaking, (1) and (2) need to be qualified to rule out outré cases involving duplicates who are time-travelers or who themselves have godlike powers. The needed changes are as follows:

(1)* Necessarily, for all agents and all times t , possibly there exists a non-time-traveling, non-godlike mental duplicate of that agent at t who ceases to exist immediately after t .

(2)* Necessarily, for all times t , any non-time-traveling, non-godlike mental duplicate who ceases to exist immediately after t cannot exercise any apparent options available to her at t .

(1)* will still be safely underwritten by principles of modal plenitude, and the argument remains valid, so I have suppressed this complication in the text.

never has any options at all. However plausible we find Hedden's claim that all finite beings fall short of perfect rationality, it is not at all plausible that all humans lead lives wholly without options, for this would mean that the subjective "ought" could play no role in guiding our actions, in meting out praise for our prudence or criticism for our blunders, or in predicting how we will behave. Surely, the foremost constraint on any account of options is that it at least sometimes allow some of us to have them. So we cannot accept the argument's conclusion.

Our choice of which premise to reject thus narrows to (3) and (5), which are consequences of (AC) and (SC)*, respectively. This means one of the constraints must go, and it had better not be (AC). Without (AC), the subjective "ought" would be indifferent to our physical limitations, and would continually demand that we cure cancer last Tuesday, put an end to war this instant, and win slews of Nobel prizes. We would, moreover, be open to blame should we fail to satisfy these insane demands. A decision theory which prescribes that we perform all manner of humanly impossible feats is a useless one.¹⁰

So it seems that we have no alternative but to reject the argument's fifth premise, and (SC)* along with it. Let me offer, now, a more in-depth explanation for why the argument succeeds. Philosophers often remark upon the fixity of the past and the openness of the future; less commonly noted, but just as evident upon reflection, is that the present, too, is fixed. In other words, there is nothing that any earthly creature is able to do at t which can alter the state of the world at t – our power to affect the world begins in the future, not the present. As a result, if there is going to be any connection between our options and our abilities, what options we have must also depend on facts about our future. But the future is full of unforeseen perils, which means there is always some hapless mental duplicate of ours, somewhere out there in logical space, who is an instant away from an unexpected demise. There is nothing this duplicate is able to do in the quantum of life she has left, so her set of available options is empty. The supervenience constraint, however, precludes us from having

10 The belief constraint weakens the sting of rejecting (AC), but only a little. Even if I mistakenly believe that I have the ability to cure cancer last Tuesday, it is not the case that I ought to or that I should be blamed for failing to do so. Similarly for decisions: if I will be struck by lightning a microsecond hence, it is not the case that I ought to decide to do anything or that I can rightly be criticized for failing to make the best decision, no matter what I believe. Indeed, without (AC), we are saddled with the ghoulish consequence that any agent who dies unexpectedly deserves blame for failing to posthumously execute her best option.

any options not available to our duplicates, so our duplicate's impending doom ends up vacating our set of options as well. The upshot is that, if we accept both (AC) and (SC)*, we are left with no options at any time. Since we cannot possibly endorse this conclusion, and the only weak link in the chain of reasoning that led us there is the supervenience constraint, the supervenience constraint has to go.

Once we reject (SC)*, there is no longer much reason to think that our options are restricted to decisions. The only rationale Hedden gives for his favored theory of options is that it is the sole account compatible with the supervenience constraint, but, as we have seen, it is not and supervenience must fail in any case. So I conclude, first, that options need not be decisions, and second, that options must be external.

3. The Opacity Of Options

Hedden is right that the failure of options to supervene on our mental states will impair to some degree the action-guiding, evaluative, and predictive roles of the subjective “ought.” But this cannot be avoided by retreating to a mentalistic account of options, for the exercise of mental abilities, too, may be thwarted by a hostile environment. This suggests that the difficulties we encountered with the supervenience constraint are merely a symptom of a deeper problem afflicting the three roles Hedden has identified. In fact, as we will see, the only way the subjective “ought” could perfectly fill any of the three roles is if an agent's options are totally transparent to her, that is, if her options and her beliefs about what options she has are always in exact alignment. Following convention, let's say that options are transparent just in case the following two theses are true:

Self-intimation: Necessarily, whenever ϕ -ing is an option for an agent, she believes that ϕ -ing is an option for her.

Infallibility: Necessarily, whenever an agent believes that ϕ -ing is an option for her, ϕ -ing is an option for her.

In more familiar terms, self-intimation requires that agents always be aware of what options they have, while infallibility demands that an agent can never falsely believe she has an option when she does not. How, though, do transparency and its sub-theses connect to Hedden's constraints? Self-intimation and (BC) are more or less equivalent; they differ only inasmuch as (BC) speaks of an agent's belief that she can ϕ , while self-intimation speaks of an agent's belief that she has ϕ -ing as an option. Self-intimation and infallibility, meanwhile, together entail (SC)*.¹¹ This will be important in what follows, but first, we will need to characterize in the abstract the problems that arise when an agent's options are not transparent to her.

Setting aside Hedden's constraints for the moment, suppose that self-intimation fails, and an agent's best option is one she has no inkling is an option for her. Then the subjective "ought" will not effectively guide that agent's behavior, because she will be in the dark about what she ought to do. She will also be subject to blame for failing to execute her best option, which seems unfair, precisely because it is unfair to blame someone who is unaware that she had a better course of action available.¹² And the subjective "ought" will fail at prediction, too, because there is little chance the agent will execute her best option if she does not realize it is an option for her.

Now suppose instead that an agent's beliefs about her options are fallible, and she falsely believes that ϕ -ing is her best option when it is not an option for her at all. Then the subjective "ought" will be insufficiently action-guiding, as the agent will be mistaken about what options she has. It will also make the wrong prediction about the agent's behavior, because, if she is rational, she will attempt to execute what she takes to be her best option and fail, while the subjective "ought" predicts that she will instead carry out her best actual option. And, should the agent choose her best actual option rather than her best apparent option, she will be unfairly absolved of blame for her decision –

11 *Proof:* Suppose that M 's options fail to supervene on her mental state, and ϕ -ing is an option for her but not for her mental duplicate, N . By the definition of mental duplication, either both M and N believe that they have ϕ -ing as an option, or neither does. If both do, N 's belief is false, and infallibility fails. If neither does, M has ϕ -ing as an option without believing that she does, and self-intimation fails. In either case, options are not transparent. Therefore, non-supervenience entails non-transparency; by contraposition, transparency entails supervenience.

12 We must again make an exception for cases where the agent's ignorance is a foreseeable consequence of her own prior blameworthy actions or negligence. We might also wish to extend this exception to include cases where the agent's ignorance is a product of her own epistemic irrationality; I discuss this possibility further in the next section.

the agent deserves criticism for passing up what she took to be her best option, even if, for reasons beyond her ken, it turns out that this was not a genuine option for her.

Let's return to Hedden's case of Jane and her somewhat less-capable doppelganger, which he intended to illustrate the need for (SC)*. Jane, recall, was able to ford the creek, and, of the actions she could perform, fording the creek was her best choice. Twin Jane, meanwhile, was unable to ford the creek, and, of the actions she was able to perform, heading home was her best choice. Suppose we say that both Jane and her twin believe they have the option of crossing the creek. Then the subjective "ought" will fail in its action-guiding role, because Twin Jane is mistaken about what options she has. It will also fail in its predictive role, because it prescribes that Twin Jane ought to head home, when we should instead expect her to try to ford the creek and fail. And the subjective "ought" will fail in its evaluative role, too, because, should Twin Jane decide to head home rather than try to cross the creek, she will unfairly avoid blame for choosing what was, from her perspective, a sub-optimal action.

If you carefully compare the problems the subjective "ought" runs into in the case of Twin Jane to the problems which arise for any agent who is fallible about what options she has, you will see that they are the same. It is Twin Jane's false belief about her options that ultimately explains why the subjective "ought" fails to properly guide, evaluate, or predict her actions. Jane's presence in the thought experiment is superfluous, and the subjective "ought" would encounter the same difficulties even if, *per impossible*, Twin Jane's options did supervene on her mental states, so long as she is mistaken about what options are available to her.¹³

The reason why it seems initially plausible that *Jane's Doppelganger* supports the supervenience constraint is that transparency entails (SC)*, which means that any case where (SC)* is violated must perforce be a case where an agent lacks transparency over her options. So it is easy to examine the case, see that it makes trouble for the subjective "ought" if we say that Jane has the option of fording the creek while Twin Jane does not, and attribute this to the fact that Jane and Twin Jane are mental duplicates with different options. But it is the failure of transparency doing all of the work

13 If we alter the case to instead stipulate that neither Jane nor her twin believes she has the option of crossing the creek, Jane's plight will then illustrate the problems the subjective "ought" encounters in any case where an agent's options are not self-intimating. In this version of the case, Twin Jane's presence in the thought experiment would be superfluous.

in *Jane's Doppelgänger*, not the failure of supervenience, and (SC)* only turns out to be useful insofar as it helps us to secure transparency. What Hedden's case really shows is that we have a strong, pretheoretical intuition that options should be transparent.

We saw in the last section that (SC)* could not be reconciled with the ability constraint; the same holds *a fortiori* of the logically stronger transparency requirement. What options an agent has depend on what abilities she has, but facts about what abilities she has depend on facts about her future. Hence, we could be infallible about our options only if we had some way of acquiring error-proof beliefs about the future, which, of course, we do not, as it is always a live possibility for any agent that she is a second away from being struck by an unseen bus or having her jugular torn out by a wild beast. No account of options could make them transparent, which means the subjective “ought” could never perfectly fill its action-guiding, blame-adjudicating, or predictive roles. Hedden's project, it turns out, was doomed from the start.

We will still want our account of options to ensure that we are as seldom mistaken about what options we have as possible. So it is open to Hedden to retrench and insist that, even if options can never quite be transparent and the subjective “ought” never quite fill its roles, a mentalistic account of options should be preferred because it will make options more nearly transparent than a fully externalist account would. This argument might succeed against crude externalist accounts which permit actions like “ford the raging creek” or “walk to Damascus” to qualify as options, but nothing compels an externalist to countenance options as coarse-grained as these. An account which restricts options to simple bodily movements, for instance, will come at least as close to achieving transparency as Options-as-Decisions. Indeed, we may often be in a better position to know what simple bodily movements are available to us than what decisions we are able to make, so the aim of making options as nearly transparent as possible offers no support for mentalism, either.

4. Externalism About Options

We have seen that (AC) and (SC)* together have the unacceptable consequence that no non-godlike agents have any options at any time, and, therefore, because (AC) is indispensable, that any

plausible account of options must reject (SC)*. We have also seen that the underlying motivation for adopting (SC)* – the conviction that an agent's options should be transparent to her – likewise proves to be incompatible with the ability constraint. Let us turn now to an account of options which embraces externalism about options rather than losing itself in pursuit of these will-o'-the-wisps.

In addition to complying with the ability constraint,¹⁴ there are several other important desiderata we should ask our account of options to satisfy. First, it should ensure that an agent's options are all pairwise mutually exclusive, because otherwise the subjective “ought” will sometimes yield inconsistent prescriptions. Suppose that an agent's options are not mutually exclusive, and her worst option is to ϕ while her best option is to ϕ -and- ψ . It will then be the case that she ought not ϕ , because ϕ -ing is worst for her, but it will also be the case that she ought to ϕ , because ϕ -ing-and- ψ -ing is best for her and she must ϕ in order to ϕ -and- ψ .¹⁵ Clearly, though, the subjective “ought” would be of little use if it placed inconsistent demands on us. For similar reasons, our account should guarantee that there is always a privileged, unique way of specifying an agent's options, lest we find that one way of carving up an agent's options makes ϕ -ing her best choice while another way of carving up her options does not, in which case the subjective “ought” will again give conflicting advice on whether or not she should ϕ .¹⁶

We also want an agent's set of options to be maximal, in the sense that every action which is at

14 What of (BC), the belief constraint? If forced to choose, I believe we should replace it with a normative epistemic constraint on options, something along the following lines:

(EC): Necessarily, ϕ -ing is an option for an agent only if she ought to believe she can ϕ , given her evidence.

The chief reason for preferring (EC) to (BC) is that (EC) allows us to criticize agents who fail to execute their best option out of ignorance in cases where the agent's ignorance is caused by her own negligence in gathering or responding to evidence. I do not think we are forced to choose between the two constraints, however. We should instead follow Parfit (2011: 162-163) in decomposing the subjective “ought” into a belief-relative “ought” and an evidence-relative “ought,” where (BC) governs the former and (EC) the latter.

The dialectic is further complicated by the fact that Hedden has since disavowed the belief constraint. In a footnote (2015a: 100, fn. 9), he reveals that he could find no way around the objection that decision-theory cannot accommodate an agent's uncertainty about what her options are. This means that an agent who believes but is not certain she has the option of walking a narrow tightrope suspended across a canyon, and is in fact able to walk the tightrope, might come in for blame for failing to execute her best option, even though it seems quite reasonable for her to choose some other course of action in light of the risk. Pollock (2002) grapples with this problem.

15 This is true, in any case, in standard deontic logic. Some heterodox accounts of deontic modals (e.g. Cariani 2013) do not have this consequence.

16 As Hedden (2012: 357) points out, this problem is at the root of Chisholm's paradox, for which see Chisholm (1963) and Jackson and Pargetter (1986).

once compatible with the constraints and inconsistent with each of the agent's other options is included in the set. If an agent is able to φ and chooses to φ , but φ -ing does not officially number among her options, the subjective “ought” will draw a blank and be unable to appraise the rationality of the agent's action at all. Moreover, we want our account of options to give some sense of their granularity. It should tell us both how dissimilar two courses of action must be before they qualify as different options for an agent, and how much time must elapse after the agent executes one option before a new suite of options becomes available to her. And, finally, if our account of options makes them into philosophical constructs of some kind, we wish to know how they connect to the sorts of actions picked out by ordinary language, because, in practice, it will be these quotidian actions which serve as the everyday loci of decision and evaluation.

One promising approach to options is to conceive of them as exhaustive combinations of atomic movements (EXCAMs). An atomic movement takes the smallest physiological unit over which the agent has conscious control at t – for humans, this will typically be an individual muscle or a mental faculty – and specifies, in minute detail, how the agent will dispose it over a small interval following t . An exhaustive combination assigning one atomic movement to each physiological unit the agent controls constitutes an option; a complete set of options includes every exhaustive combination compatible with (AC) and whatever other constraints we wish to introduce.¹⁷

By individuating options as narrowly as possible, Options-as-EXCAMs guarantees that they will always be mutually exclusive, maximal, and uniquely specified. What's more, we can allow the granularity of atomic movements to be set by our built-in cognitive limitations by way of the ability constraint. If human fine motor skills are not so precise that we can elect to move our hands six micrometers forward rather than seven, these do not register as distinct atomic movements for agents like us. Similarly, because we are incapable of making decisions in successive instants, we can let the interval between when an agent begins to execute an option and when a fresh suite of options becomes

17 Note that the ability constraint is applied twice in constructing an agent's set of options – first, to limit the atomic movements included in combinations to those under the agent's control, and second, to restrict the set of combinations which count as options for an agent to those the agent is able to perform. This redundancy is necessary because an agent may be able to perform either of a pair of atomic movements individually but not both together.

available to her be the period of time before she first becomes psychologically capable of revoking any atomic component of the earlier decision. Generally, where the agent is sufficiently resolute and the background environmental conditions sufficiently stable, it will be safe to extend the options indefinitely far into the future and thereby recover the more familiar sort of action from everyday life and toy decision theory cases.

The chief drawback of this proposal is that it leads to a dizzying proliferation of options. Every tiny difference in how we choose to position our bodies or direct our minds carries with it hundreds or thousands of new options, and a complete set of options for an agent at a time will be longer than we could ever hope to enumerate. It may be, however, that any account of options will exhibit this defect when fleshed out in sufficient detail.¹⁸ A mentalistic account like Hedden's certainly will, as each way of positioning each part of our body also corresponds to a vast array of different decisions we can make. So we may have no choice but to grow accustomed to this profligacy if we wish for our theory of options to reflect the combinatorial network of causal pathways connecting our minds to the world. An agent who starts walking towards Damascus left-foot-first and an agent who starts walking towards Damascus right-foot-first are really performing different actions, and any account which elides this distinction has not truly succeeded at supplying us with a maximal set of options.

5. Conclusion

Options could not supervene on our mental states, for what options we have is constrained by what abilities we have, and none of our abilities are immune to the vicissitudes of a hostile world. Options could not be transparent, for what options we have is constrained by what abilities we have, and none of our abilities are immune to the surprises of an uncertain world. These failures of supervenience and transparency undermine Hedden's rationale for limiting options to decisions, and

¹⁸ We might try to tame this excess by individuating options by their moral characteristics, for instance, by combining all EXCAMs with the same expected utility into a single option. This proposal, though, has a fatal defect: two actions can be as different as you like, intuitively speaking, yet still share all of the same moral features. As an example, I might save five dogs from being euthanized by volunteering my nights off at the pound, or I might instead save five dogs from euthanasia by working long hours of overtime and donating the proceeds to the ASPCA, and it does not seem appropriate to conflate these two courses of action into a single option. Because ties like these are bound to occur from time to time, individuating options by their moral features will inevitably require us to run together actions which should be kept distinct.

pose a serious challenge to his program of making all the demands rationality places on us a function of our mental states alone. A more promising approach, Options-as-EXCAMs, rejects any arbitrary distinction between actions of the mind and actions of the body. Options-as-EXCAMs extends our options to every joint and crease in the causal structure of the world over which we exercise any conscious control.

References

- Broome, J. (2013). *Rationality Through Reasoning*. Malden: Wiley-Blackwell.
- Cariani, F. (2013). "'Ought' and Resolution Semantics," *Nous* 47 (3): 534-558.
- Chisholm, R. (1963). "Contrary-to-Duty Imperatives and Deontic Logic," *Analysis* 24 (3): 33-36.
- Conee, E. and Feldman, R. (2001). "Internalism Defended," *American Philosophical Quarterly* 38 (1): 1-18.
- Fischer, J. and Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, H. (1969). "Alternate Possibilities and Moral Responsibility," *The Journal of Philosophy* 66 (23): 829-839.
- Hedden, B. (2012). "Options and the Subjective Ought," *Philosophical Studies* 158 (2): 343-360.
- Hedden, B. (2015a). *Reasons Without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.
- Hedden, B. (2015b). "Time-Slice Rationality," *Mind* 124 (494): 449-491.
- Jackson, F. and Pargetter, R. (1986). "Oughts, Options, and Actualism," *Philosophical Review* 95 (2): 233-255.
- Lewis, D. (1986). *On the Plurality of Worlds*. Malden: Blackwell.
- Moss, S. (2015). "Time-Slice Epistemology and Action Under Uncertainty," in J. Hawthorne and T. Gendler (eds.) *Oxford Studies in Epistemology* 5.
- Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.
- Pollock, J. (2002). "Rational Choice and Action Omnipotence," *Philosophical Review* 111 (1): 1-23.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.