Instrumental Reasons[1]

"What reason have I to do *that*?" someone asks. The answer (if there is one) is often: "Doing that is a means to *something else* that you have reason to do."[2] When is this answer correct? What does the existence, or weight, of such an "instrumental reason" depend on?

The focus of these questions—instrumental reasons—should be distinguished from what might be called "instrumental rationality." *Instrumental rationality* is a matter of the coherence, or unification, of the deliberating agent's state of mind. Insofar as one wills an end and believes that something is a necessary means to it, one is, in refusing to will those means, at odds with oneself.[3] By contrast, *instrumental reasons* concern not the coherence of one's deliberation, viewed as a psychological episode, but instead the structure of the subject matter about which one deliberates. If there is reason for one to pursue an end[4] (whether or not one pursues it), and if an action is in fact (or relative to the better information of an advisor or onlooker[5]) a means to that end (whether or not one knows it), then that is itself a reason for one to adopt the means— quite apart from any considerations about one's psychological coherence.[6] This is reflected in familiar patterns of advice. If I know that there is reason for you to lower your blood pressure, and if I know that your taking this medicine will do that, then I am liable to tell you, if I can, that that is a reason for you to take this medicine—whether or not you intend to lower your blood pressure, or know that this medicine will do that. Furthermore, at least on the assumption that what one ought to do depends on the reason there is for one to do it, a sort of transmission of 'ought's from ends to means might be expected. If I know that you ought to lower your blood pressure, I may well conclude that, because of this, you ought to take the medicine.

A burgeoning literature offers many proposals about how this instrumental transmission of reasons and 'ought's works. Common proposals are:

1

Ought Necessity: If one *ought* to *E*, and *M*-ing is a *necessary* means to *E*-ing, then, because of that, one *ought* to *M*[7]

and:

Strong Necessity: If there is reason for one to *E*, and *M*-ing is a *necessary* means to *E*-ing, then that is a reason, *at least as weighty*, for one to *M*,[8]

which in turn entails:

Weak Necessity: If there is reason for one to *E*, and *M*-ing is a *necessary* means to *E*-ing, then that is a reason, of *some* weight, for one to *M*.

However, Ought Sufficiency, Strong Sufficiency, and Weak Sufficiency, which substitute "sufficient means" for "necessary means," have also been suggested.[9]

I begin by considering some objections to Strong Necessity, Ought Necessity, and the Sufficiency Principles: namely, they do not meet two basic desiderata, which are identified in sections 1 and 4. The problem with Strong Necessity, and perhaps also Ought Necessity, is that necessary means may do little to "probabilize" the end, as I explain in sections 2–3. The problem with Sufficiency principles is that sufficient means may be "superfluous" with respect to the end, as I explain in section 5.[10] Section 6 puts forward a positive proposal, titled "General Transmission." Indebted to the work of Raz (2005a and b), it aims to capture, admittedly more pedantically, the ideas animating his "Facilitative Principle" (although I leave explicit discussion of the connections to notes).

One might wonder, however, why we should care what the correct account of instrumental reasons is. So long as we grant that there is such a phenomenon as instrumental transmission, why bother to describe it precisely? Perhaps it is enough to cite the aim, which fuels so much of philosophy, of making explicit the conceptual structures that implicitly

underwrite our thought and practice. But there are two more specific reasons for caring what the principles of instrumental transmission say.

First, such principles have been pressed into service by various treatments of the "normativity" of instrumental rationality: treatments of the demand, or seeming demand, to make our psychology means-end coherent. In particular, some have suggested that there are reasons to be instrumentally rational: for example, that one has "wide-scope" reason (either not to intend the end of, say, starting World War III, or to intend the means of launching a nuclear missile). Others have objected that, according to some principles of instrumental transmission, it would follow, absurdly, that one had "narrow-scope" reason to intend to launch a nuclear missile. While I think that this line of objection is fundamentally correct, its proponents have relied on questionable principles of instrumental transmission, and, accordingly, they have been met with plausible replies. Section 7 shows how General Transmission might put this objection on a surer footing.

Second, once we get clear on what the principles of instrumental transmission would have to look like, we may begin to doubt whether there really is such a phenomenon as instrumental transmission. Section 8 explores this doubt. There I resist a line of thought that suggests, to a first approximation, that General Transmission is just a minor theorem of a certain kind of decision theory and, going further, that the very idea that reasons for means are explained by reasons for ends is a kind of myth.

1. First desideratum: means must probabilize the end
In this section, I discuss the first of two basic desiderata for an account of instrumental transmission. Intuitively, reasons for an end, one's E-ing, should transmit to a means, one's M-ing, to the extent that one's M-ing "probabilizes" one's E-ing, or "makes it likely." If the end for

which I have reason is my putting out the fire, and using a fire hose makes that more probable than using a garden hose, which in turn makes it more probable than blowing weakly on the raging flames, then, intuitively, I have more reason to use a fire hose than a garden hose, and more reason to use a garden hose than to blow.

To a first approximation, we might say that if there is reason for one to *E*, and if there is positive probability, conditional on one's *M*-ing, that one *E*'s, then that is a reason for one to *M* whose strength depends on the reason to *E* and the probability.[11] But this won't do. Suppose a boxer's end is landing a punch. In every possible outcome in which he decides to throw a punch, he "telegraphs" his intention by conspicuously gritting his teeth. In spite of inadvertently warning his opponent, he nevertheless sometimes connects. So the probability, conditional on telegraphing, of connecting is positive. However, intuitively, no reason transmits to telegraphing. This is because at no possible outcome does it "help to bring it about" or "help to make it the case." Instead, telegraphing is merely a by-product of something (deciding to throw the punch) that does help to bring it about. So I formulate the desideratum as:

> *Means Probabilize*: If there is reason to *E* and there is positive probability, conditional on one's *M*-ing, that *one's M-ing, or some part of one's M-ing, helps to bring it about that one E's*, then that is a reason to *M*, whose strength depends on the reason to *E* and on the probability.

"Helping to bring about" need not mean only *causing*. It can also mean *constituting, satisfying preconditions of*, or *preventing things that would prevent*,[12] as well as *helping* to cause, constitute, satisfy preconditions of, or prevent things that would prevent. This means that one's *refraining* from doing something that would prevent one from *E*-ing may help to bring about

one's *E*-ing, insofar as such refraining, in effect, prevents a preventer. We can leave open whether there are still other ways of helping to bring about an end.[13]

Why "or some part of one's *M*-ing"? We don't even need this clause, if we assume that an action helps to bring about at an outcome whatever any part of it helps to bring about at that outcome. But suppose we do not assume this. Consider a case in which (i) giving the patient Drug A brings it about that we save his life at every outcome whether or not we also give Drug B at that outcome, (ii) that (giving Drug A and giving Drug B) does not help to bring it about that we save his life at any outcome, and (iii) that (giving Drug A and giving Drug B) brings it about that we improve his complexion at every outcome. Intuitively, we have at least as much reason (to give Drug A and to give Drug B) as we have to give Drug A. But, without the phrase "or some part of one's *M*-ing," why should this be? There would be very strong reason to give Drug A, because it is instrumental to saving the patient's life, but only much weaker reason (to give Drug A and to give Drug B), which is instrumental only to improving his complexion. If we add "or some part of one's *M*-ing," by contrast, then (giving Drug A and giving Drug B) inherits the instrumental reason of its part, giving Drug A.

The relevant notion of "probability" can be understood in different ways. I favor an epistemic framework, where the possible outcomes (perhaps sets of worlds or underspecified worlds) are ways things might be relative to a relevant body of information, and where the probabilities assigned to those outcomes are likewise fixed by that body of information.[14] I stress that the relevant body of information *need not be that of the agent at the time of acting*. It may be instead the information, for example, of an advisor, or an onlooker. More generally, it may be contextually specified.[15] However, there are other frameworks, which are compatible with much of what I go on to say. For example, on a historical framework, the possible

outcomes are ways things might unfold, compatibly with how they have actually unfolded up to the relevant time (and possibly also constrained by certain physical and psychological laws), and the probabilities (here not epistemic) attached to those outcomes are determined by how things actually are (and the laws in play) at the relevant time.[16] When I say that something is "necessarily" the case, that is equivalent to saying that it is the case at all possible outcomes, however these are understood. I assume that some positive probability attaches to every possible outcome (although, in principle, this may fail when there are infinitely many possible outcomes).

2. Why Strong Necessity has trouble with the probabilization desideratum
If something along the lines of Means Probabilize is correct, it suggests an objection to Weak Necessity. "Weak Necessity may seem plausible, because if one *fails* to take a necessary means, then there is *no* chance of achieving the end. However, it is compatible with that truth that *even if one does* take the necessary means, one has no chance of achieving the end. If so, then those necessary means do not probabilize the end at all. So, if something along the lines of Means Probabilize is correct, then no reason would transmit to such necessary means."[17]

However, there are plausible replies to this objection. First, if there is *no* chance of one's *E*-ing even if one takes the *necessary* means, then there is no chance of one's *E*-ing at all. And if there is *no* chance of one's *E*-ing, then it is not clear that one has *reason* to *E* in the first place. That would follow, at least, from the conjunction of the claims:

Reasons Availability: One has reason to *X* only if *X* is available to one.[18]

and:

Weak Availability: *X* is available to one only if there is some possible outcome at which one *X*'s.[19]

Second, if there is no chance of *E*-ing, then there is a danger that everything, trivially, becomes a necessary means to *E*-ing. If no matter what one won't *E*, then can't one say of anything that, unless one does that, one won't *E*?

Taking these replies on board, I work with the following:

Definition of Necessary Means: One's *M*-ing is a *necessary* means to one's *E*-ing iff:

(i)   in every possible outcome at which one *E*'s: one's *M*-ing helps to bring about one's *E*-ing and if one had not *M*-ed, one would not have *E*-ed, and

(ii)   there is *some* possible outcome at which one *E*'s.

This definition ensures that necessary means probabilize the end to at least some degree. So, Means Probabilize gives us no grounds to doubt that *some* reason transmits to necessary means, at least on this definition. So *Weak* Necessity seems secure.

Yet why should we expect *Strong* Necessity to hold? Granted, since necessary means probabilize the end to some extent, they inherit *some* of the reason. But necessary means might probabilize the end only to some minute degree. Why then should the necessary means inherit, as it were, *all* of the weight of the reason for the end?

Consider Lucky and Unlucky, who occupy parallel universes. Each has an antique sitting on his front porch, which the rain threatens to ruin. A necessary means to saving the antique is taking a taxi back home. There is reason to refrain from[20] taking the taxi; it costs money, say $20. But this cost is outweighed by the value of the antique, say $100. The only difference in their situations is that in Lucky's universe, the rain will be slow in coming, and so he is very likely to get him home in time, if he takes the taxi: say that the probability is .9. In Unlucky's

universe, by contrast, he is extremely unlikely to get him home in time, even if he takes it: say the probability is .1.

Intuitively—and as Means Probabilize implies—Unlucky has *weaker* reason to taxi than has Lucky, given the low probability that the taxi will get there in time. Therefore, if Lucky and Unlucky have the reason of the same strength to save the antique, then we have a counterexample to Strong Necessity. For Strong Necessity would then imply, incorrectly, that Unlucky has reason of the same strength as has Lucky to take the taxi.

As will be argued for a bit more below, Lucky and Unlucky would seem to have reason of the same strength to save the antique. *If* Unlucky *saves* the antique, things will be, in relevant respects, just as they are *if* Lucky saves it. Granted, if Unlucky *tries to save* the antique (say, by taking the taxi), then things are not likely to be as they are if Lucky tries to save the antique. But that seems relevant not to the question of whether there is reason to *save* the antique, but instead to the question of whether there is reason to *try* to save it (say, by taking the taxi).

If this is right, then Lucky and Unlucky have the same reason for the end. Means Probabilize implies that Lucky has more reason for his necessary means, since his means have a higher probability of achieving the end. This contradicts Strong Necessity, which says that there is as much reason for the necessary means as there is for the end. So, if this is right, we should reject Strong Necessity.[21]

3. Why Ought Necessity may have trouble with the probabilization desideratum
Is Ought Necessity likewise falsified? Can we say that, while Unlucky ought to save the antique, it is not the case that he ought to taxi?

The argument that Unlucky ought to save the antique is as follows. Lucky ought to save the antique. As far as *saving* the antique is concerned, Unlucky's situation is the same. If Unlucky saves the antique, things will be exactly as they are if Lucky saves it.

The argument that it is not the case that Unlucky ought to taxi rests on the point just made: namely, that Unlucky's reason for taxi-ing weakens as the probability lowers.[22] By contrast, his reason to refrain from taxi-ing—namely, to avoid its cost—remains constant. So, as the probability lowers, a point is reached at which there is more reason to refrain from taxi-ing than to taxi. Assuming:

> *Refrain Reasons-Ought*: One ought to *X* only if there is not more reason to refrain from *X*-ing than to *X*,

it would follow that it is not the case that Unlucky ought to taxi.

> This would also follow from:

> *Relevant Reasons-Ought*: One ought to *X* iff *X*-ing is, from among the relevant, available options, what one has most reason for (or, in the case of ties, *X* is the disjunction of the relevant, available options that one has most reason for).[23]

However, one might worry that Relevant Reasons-Ought proves too much. If one ought to *E*, and the set of relevant alternatives to *M*-ing includes *E*-ing, then it will be the case that one ought to *M* only if *M*-ing is not only necessary, but also sufficient for *E*-ing. But this would mean that it isn't the case that *Lucky* ought to taxi. More generally, if we accept Relevant Reasons-Ought, then Ought Necessity seems underspecified. With respect to *which* set or sets of relevant alternatives does Ought Necessity say that if one ought to *E*, then one ought to *M*? A modest proposal is:

Relevant Ought Necessity: If one *ought* to E relative to {E, refrain from E-ing} and M-ing

is a *necessary* means to E-ing, then, because of that, one *ought* to M relative to {M,

refrain from M-ing}.

And this is falsified by Lucky and Unlucky.

I now address three main lines of reply to this argument against Ought Necessity. (The

reader may skip these and turn to section 4 without any loss of continuity.) Reply 2 and,

especially, Reply 3 parallel, in large part, "possibilist" criticisms of the "actualism" of Jackson

(1985) and Jackson and Pargetter (1986)—or, more accurately, what Goble (1996) calls

"probabilism." Much of my response, then, is, in effect, a defense of probabilism. I doubt that

many possibilists will be convinced; the gulf between basic intuitions may be too wide for that.

But I say something about the debate for purposes of illustration.

Reply 1: The first reply takes aim at the argument that Unlucky ought to save the antique,

on the grounds of:

Refrain Strong Necessity: If there is reason to refrain from M-ing (say, because M-ing

would incur some cost), and if one M-ing is a necessary means to one's E-ing, then that is

a reason at least as strong to refrain from E-ing.[24]

Since Unlucky has stronger reason than Lucky to *refrain* from taking the taxi, he has stronger

reason than Lucky to refrain from saving the antique. So, it is compatible with its being the case

that Lucky ought to save the antique that it is not the case that Unlucky ought to save it.

The first problem with Reply 1 is that Unlucky's reason to refrain from taking the taxi is

that doing so would avoid the cost of the taxi. But *Lucky* would seem to have a reason of

precisely the *same* strength to refrain from taking the taxi: namely, that it would avoid exactly

the same cost. So it is *not* the case that Unlucky has stronger reason than Lucky to refrain from

taking the taxi. (Granted, Unlucky has weaker reason than Lucky to *take* the taxi: namely, its lower probability of saving the antique. But, again, that doesn't imply less reason to *save* the antique, since the probability of Unlucky's saving the antique saving the antique is the same as the probability of Lucky's saving the antique saving the antique: namely, one.) Refrain Strong Necessity would seem, therefore, to transmit the same reason to refrain from saving in both cases. And in both cases there is just as much reason to save. But if the balance of reasons, for saving and refraining, is exactly the same for both Lucky and Unlucky, how then can it be that Lucky ought to save, but not Unlucky?

The second problem with Reply 1 is that Refrain Strong Necessity seems questionable. Are we to accept it because it is a special case of:

> *Strong Necessary Condition*: If there is reason to *X*, and necessarily, one *X*'s only if one *Y*'s, then that is a reason at least as strong to *Y*?

The arguments of the last section against Strong Necessity (as well as the points in note 7) already tell against this principle. The truth in the vicinity, it seems to me, is:

> *Cost Avoidance*: If there is a positive probability, conditional on one's *X*-ing, that one's *X*-ing helps to bring it about that one avoids some cost *C,* then there is reason to *X,* whose strength depends on the cost and the probability.[25]

In other words, the reason that Unlucky has to refrain from saving is that it is a *means* of avoiding the cost of taxi-ing. And how strong this reason is depends on the extent to which refraining from saving helps to bring it about that that cost is avoided. But Cost Avoidance doesn't imply Refrain Strong Necessity. This is because avoiding the end is not always a good means to avoiding the cost of the necessary means.

Indeed, this fact allows us to construct a different kind of counterexample to Ought Necessity. In *Cost-Unlucky*, the taxi is guaranteed to get Cost-Unlucky home in time. Here Cost-Unlucky's misfortune is that the cost of the taxi is now $120: more than the value of the antique. Clearly, there is stronger reason to refrain from taxi-ing than there is to taxi, and so (according to Refrain Reasons-Ought) it is not the case that Cost Unlucky ought to taxi. But is there likewise stronger reason to refrain from *saving* than there is to save? According to Cost Avoidance, it turns on how probable it is, conditional on his refraining, that his refraining helps to bring it about that he avoids the cost of taxi-ing. And this turns on how Cost-Unlucky is likely to refrain from saving, if he does refrain. If—as we would usually expect—Cost-Unlucky is likely to refrain from saving, if he refrains, by refraining from taking the taxi, then, yes, refraining from saving is very likely to avoid the cost of the taxi. But if—to consider a more unusual case—Cost-Unlucky is likely to refrain from saving, if he refrains, by akratically taking the taxi and then dawdling on the front porch, gossiping with his neighbor, until the antique is ruined, then, no, refraining from saving isn't very likely to avoid the cost of taxi-ing. In this more unusual case, Cost-Unlucky might have very weak reason to refrain from saving, since he is so likely incur the cost of the taxi anyway. By manipulating the probabilities, we can make his reason to refrain from saving weaker than his reasons to save. Thus, it would be consistent with Refrain Reasons-Ought, and entailed by Relevant Reasons-Ought (at least relative to {save, refrain from saving}), that he ought to save. As we might think: "It is not the case that Cost-Unlucky ought to take the taxi in the first place, since it costs more than the antique. But given that he is so likely (foolishly) to take the taxi whether or not he saves the antique, he ought at least to save the antique and cut his losses."[26]

Reply 2: A second line of reply is to deny that Unlucky ought to save, because saving is not available to him. This reply would assume:

Ought Availability: One ought to *X* only if *X*-ing is available to one,

and something stronger than Weak Availability, such as:

Strong Availability: *X* is available to one only if it is necessary that there is some 1, 2, …
N such that whether one 1's is entirely under one's control and if one 1's, whether one 2's
is entirely under one's control and if one 2's, … whether one *N's* is entirely under one's
control and if one *N*'s, then one *X*'s.

Saving the antique is not Strongly Available to Unlucky, so long as the probability that if he takes the taxi, he will get there in time, is less than one, which the counterexample, as we have been understanding it, requires.

The first problem for Reply 2 is that Strong Availability is very restrictive. For instance, it would mean that saving the antique is not available even to *Lucky*, and so that it is not the case that Lucky ought to save it.[27] Worse, in conjunction with Reasons-Availability (which was used to defend Weak Necessity), it would follow that Lucky does not even have *reason* to save the antique. Worse still, that would in turn threaten the judgment that Lucky even has reason to *take the taxi*. After all, any such reason would be instrumental, in the service of some end. But of what end? We cannot say that the reason is transmitted from the end of *saving the antique*, since we are denying that Lucky has reason for that end. It only pushes the problem back to say that it is transmitted from reason to *try* to save the antique. That reason is also instrumental, in the service of what such trying might bring about.[28] More generally, Strong Availability tends to crowd out much of the phenomenon of instrumental transmission. Very often the ends for which

we are said to have reason are not strongly available; *something* outside of our control might keep us from achieving them.[29]

The second problem for Reply 2 is that there are counterexamples to Ought Necessity where the end is strongly available. One such counterexample is Cost-Unlucky, for whom saving the antique was strongly available, albeit a bad bargain. Another such example is *Procrastinate*, described by Jackson (1985) and Jackson and Pargetter (1986). Procrastinate has strong reason to review a book. A necessary means to this is accepting the commission to review it. Accepting has one cost: namely, that of contacting the difficult-to-reach review editor in time to accept. Procrastinate is extremely unlikely to write the review, if he accepts. Reviewing the book is Strongly Available. If he accepts the commission, then it is necessary that writing the review will be, at a later time, fully under his control. But otherwise the structure of the case is relevantly similar to Unlucky. Intuitively, Procrastinate has weaker reason to accept than to refrain, and so Refrain Reasons-Ought and Relevant Reasons-Ought imply that it is not the case that he ought to accept. And on Relevant Reasons-Ought, Procrastinate ought to review.

One can rule out these counterexamples by ratcheting up the requirements on availability to:

> *Super-strong Availability*: X-ing is available to one only if it is necessary that there is some Y such that whether one Y's is fully under one's control and if one Y's, then one X's.[30]

But Super-Strong Availability is still more restrictive. It writes off further cases of instrumental transmission: namely, all those in which the reason to take the means is to help to put oneself in a position, at some later time, from which there is something that one can (but might not) do to guarantee the end. (Indeed, for Raz 2005a, this is by definition what "facilitating plans" do.)

Moreover, there are counterexamples that conform even to Super-Strong Availability. In order to save his arm, *Mountaineer* must now cross over a bridge to the other side of a chasm. But if he steps now, he is almost certain to look down while doing so. And if he looks down, he will almost certainly not successfully cross. Indeed, he will almost certainly lose his balance and plummet to his death. All the same, it is entirely within his control to step without looking down. He doesn't know, although we know, that stepping while looking down will upset his balance. Stepping without looking down now is super-strongly available, and he ought to step without looking down now. A necessary means to stepping without looking down now is stepping now, which is also super-strongly available. But it is not the case that he ought to step. Indeed, he ought not to step, since it will almost certainly be his death.

Reply 3: A final reply to these alleged counterexamples to Ought Necessity is to grant that Unlucky, Procrastinate, and Mountaineer ought to achieve the end, but to argue that they ought also to take the necessary means. As for Cost-Unlucky, the reply would grant that he ought not take the necessary means, but would argue that it is also not the case that he ought to achieve the end.

Since Refrain Reasons-Ought and Relevant Reasons Ought, as we have seen, support the opposite conclusions, this approach needs some other account of the relation between reasons and oughts. The likely candidate would be:

Necessary Reasons-Ought: One ought to *X* iff *X*-ing is, or is a *necessary means* to, what, from among *all* the available options one has most reason to do.

Suppose that the option for which Unlucky has most reason includes saving as a necessary component. This will be some "exhaustive" option, which determines, for every available action that might relevantly affect the outcome, whether or not one performs that action, such as: saving

the antique and donating it to Oxfam and working tirelessly henceforth to end world hunger and so on. But then it will also be true, on this account, that Unlucky ought to taxi as well. Similarly for Procrastinate and Mountaineer. For his part, Cost-Unlucky can't have most reason for an option that includes his saving the antique, since that will always be a net loss. So it follows that it is not the case that he ought to save it.

At this point, intuitions may simply diverge. But mine incline me to affirm both that it is *not* the case that Unlucky/ Procrastinate/ Mountaineer ought to taxi/accept/step and, moreover, that Unlucky/ Procrastinate/ Mountaineer ought *not* to taxi/ accept/ step. (I set aside Cost-Unlucky for simplicity.)[31] Moreover, I find it hard to deny *Refrain Reasons-Ought*: that is, to accept that one ought to do something that one has more reason to refrain from doing. Most importantly, these claims seem to me to track good and bad *advice*, and good or bad *deliberative conclusions,* at least in the cases of Unlucky and Procrastinate. Suppose all the facts are laid out for Unlucky or Procrastinate, and they ask you, or ask themselves, "So, knowing what we know, what ought I to do now?" It seems bad advice to say, "You ought to taxi/accept," and good advice to say, "You ought not taxi/accept. Given that you are so unlikely to save/write, it will be a disaster if you do. You have more reason to refrain." For the same reasons, it seems a bad deliberative conclusion for them to reach themselves, "So I ought to taxi/accept." The correct conclusion, it seems, is to say, "But I ought not taxi/accept. Given that I am so unlikely to save/write, it will be a disaster if I do. I have more reason to refrain."[32]

One might propose theories of error for these intuitions. For instance, one might stress the distinction between *something* that one ought to do and *all* that one ought to do. To say, "He ought to taxi/accept/step," while strictly true, is misleading, since it suggests that this is *all* that he ought to do. A fuller answer would be: "He ought to save/accept and write/step without

16

looking down, and that requires his taxi-ing/accepting/stepping. That is why he ought to taxi/accept/step." Similarly, it sounds odd to say that one ought to do something that one has more reason to refrain from doing, only because "ought" is here heard as "all that one ought," and, so heard, the claims are actually false. It won't be a disaster if one does *all* that one ought, and one does not have more reason to refrain from doing *all* that one ought. Alternatively, with Kiesewetter (forthcoming), one might suggest that the intuitions confuse a "narrow-scope" with a "wide-scope" 'ought.' It is false that if he is unlikely to write/avoid looking down, then he ought accept/step. What is true is that he ought to see to it that (if he is unlikely to write/avoid looking down, then he does not accept/step).

These theories of error, however, still don't seem to deliver plausible advice or deliberative conclusions in Unlucky and Procrastinate. In the episodes of advice and deliberation that we described, all the facts are made plain to all parties. So there can be no lingering confusion between "some of what" and "all of what" one ought to do. And on the wide-scope view, we can never get to the desired answer to the question, "So, what ought I to do now?" The closest we can come is the unhelpful answer, "You ought either to be likely to write in the future, or not to accept."

One might revert to Super-Strong Availability, so that Mountaineer is left as the only relevant case. The advice, "You ought to step," and the conclusion, "I ought to step," seem fine in the Mountaineer's case. For once we give the Mountaineer the full story, or he gives it to himself, he comes to know that he ought to step without looking down if he steps. And so it becomes very likely that if he steps, he will step without looking down. And so "You ought to step" becomes perfectly good advice, and "I ought to step" becomes a perfectly good

deliberative conclusion. Apart from the worry that Super-Strong Availability is too restrictive, therefore, I have no further argument against this reply.

Alternatively, one might deny, on the one hand, that the badness of advice is persuasive evidence of its falsity and, on the other, that the goodness of advice is persuasive evidence of its truth. What makes advice good or bad is, instead, whether one *ought* to or *ought not* to give it. Suppose that if X is now under the agent's control, and if we say to the agent, "You ought to X," then the agent will X. However, if X is not now under the agent's control, then saying to the agent, "You ought to X," will have no effect. What, then, ought we to say to, say, Procrastinate? Clearly, we ought not to say, "You ought to accept," because this will lead Procrastinate to accept, but make him no more likely to write, if he accepts. Indeed, we ought to say, "It is not the case that you ought to accept." This is the sense in which "You ought accept" is bad advice, but "It is not the case that you ought accept" is good advice. The *adviser ought not say* the former and *ought* to say the latter. All the same, this bad advice is true, and this good advice is false.

The first problem with this proposed explanation is that it lets truth and falsity diverge too far from the goodness or badness of the corresponding advice. To be fair, even an *opponent* of Ought Necessity must accept *some* divergence. In particular, he must agree with the proponent of Ought Necessity that some true 'ought' claims will be bad advice. For example, telling Procrastinate, "You ought to accept and write," is bad advice, even though the opponent of Ought Necessity, no less than the present proponent, agrees that this claim is true. This is because it is naturally heard as, "Of the *decisions* that you can now make, accepting and writing is the decision that you ought to make," which, according to the opponent of Ought Necessity, is false. However, the defender of Ought Necessity is arguing not only that some *true* "ought"

claims are *bad* advice, but moreover that some *false* "ought" claims, such as "You ought not accept," are *good* advice. It's more surprising that a good adviser sometimes ought to keep his counsel about what the acknowledged facts imply than that he ought to lie to his advisees about it.

The second problem is that this explanation doesn't work in the case of first-person deliberation. Applied to that case, the idea would be to explain why "I ought to accept" is a *bad* deliberative conclusion by appeal to the thesis that a deliberative conclusion is bad if one ought not to judge it, but good if one ought to judge it. But suppose that Procrastinate will accept only if he concludes, "I ought to accept"—not an unlikely situation. Then, according to Necessary Reasons-Ought, he *ought* to conclude, "I ought to accept." But then "I ought to accept" should be a *good* deliberative conclusion, which is the opposite of what was to be explained.

4. Second desideratum: means must not be superfluous
To recap: We have identified one desideratum for a principle of instrumental transmission: namely, Means Probabilize. Then we observed that it tells against Strong Necessity and, arguably, also against Ought Necessity. Now we turn to the second desideratum, which is revealed by a defect in Means Probabilize, although a defect that does not affect the arguments given so far.

Suppose Dr. Twoways has the end of relieving the patient's pain. Drug 1 alone will do this for sure. If he also gives Drug 2, it will at first neutralize Drug 1, but then combine with Drug 1 in the patient's bloodstream to become Drug 3, which will also relieve the patient's pain for sure. Suppose that at every possible outcome Twoways gives the patient Drug 1. (Either it is known, or fated, that he will give Drug 1.) In this case, it seems that no reason transmits to Twoways's giving Drug 2. However, Means Probabilize implies that reason *does* transmit. The

probability, conditional on giving Drug 2, that Twoways relieves the patient's pain and that giving Drug 2 helps to bring this about, is *one*. In *every* possible outcome in which he gives Drug 2, his doing so *helps to cause* the relief of pain. In every such outcome, Drug 3 is the proximate cause of the relief of pain, and the presence of Drug 3 is a causal consequence of having administered Drug 2.

The reason why no reason transmits, one wants to say, is that in every possible outcome in which Drug 2 helps to bring it about that the pain is relieved, Drug 2 brings this about *superfluously*. In other words, we should replace Means Probabilize with:

> *Means Probabilize Nonsuperfluously*: If there is reason for one to *E*, and there is positive probability, conditional on one's *M*-ing, that one's *M*-ing, or some part of one's *M*-ing, helps to bring it about that one *E*'s *in a nonsuperfluous way*, then there is reason for one to *M*, whose strength depends on the reason for one to *E* and on the probability.

The intuitive idea is that one's *M*-ing is *superfluous* with respect to one's *E*-ing in outcome *W* just when, if one's *M*-ing were "removed" from *W*, one would still *E*. Consider any outcome in which Twoways gives Drug 2. If we "took away" Twoways's giving Drug 2 from that outcome, Twoways's giving Drug 1 would "still be there," and so Twoways would still relieve the patient's pain.

To say "if one's *M*-ing were 'removed' from *W*, one would still *E*" is *not* to say "if one did not *M*, one would still *E*." One difference is that the counterfactually "closest" world, as it were, in which one does not *M* may be one in which *not only* is one's *M*-ing "taken away," *but also something new* is "added in." For example, suppose the end is entering the house. In every possible outcome in which I enter by the front door, it is true that if I had not entered by the front door, I would have entered by the back door. In other words, at the "closest" world at which I do

not enter by the front, not only is my entering by the front "taken away," but also something new, my entering by the back, is "added in." Intuitively, we don't want to count entering by the front door as superfluous. If we did, then no reason would transmit to entering by the front door, and (if the situation is symmetrical) no reason would transmit to entering by the back door, even though (entering either by the front or by the back) would be necessary and sufficient for the end. But it would surely be wrong to say of me, after I enter by the front, not only that there was no more reason for me to do that than to enter by the back (which is true enough), but moreover that, as far as entering the house was concerned, there was no reason for me at all to do it—as one might properly say that I had no reason to tie myself to a stake in the lawn, preventing myself from entering the house.

Another reason why "if one's *M*-ing were 'removed' from *W*, one would still *E*" is *not* equivalent to "if one did *not M*, one would still *E*" is that the "closest" world in which one does not *M* may be one in which only *part* of *M*-ing, not the whole of it, is "taken away." For example, it might be the case at every outcome at which Twoways (gives Drug 1 and gives Drug 2) that, if Twoways did not give both, Twoways would still give Drug 1. Even though, given that Twoways gives Drug 1, giving Drug 2 is superfluous, we *don't* want to say that given that Twoways gives Drug 1, giving *both* is superfluous. Suppose that giving both drugs (and only giving both) will improve the patient's complexion, and that giving both is no more costly than giving Drug 1 alone, since Drug 2 is about to spoil anyway. Intuitively, Twoways has more reason to give both than to give only Drug 1: that will relieve the patient's pain and, without any additional cost, improve his complexion. But if we count giving both as superfluous, then reason to relieve pain will transmit only to giving Drug 1 alone. Since reason to relieve the patient's

pain outweighs the reason to improve the patient's complexion, Twoways will have more reason to give only Drug 1 than to give both.

To capture the intuitive idea, therefore, we should say something like this:

Definition of superfluous means: One's *M*-ing is superfluous with respect to one's *E*-ing in an outcome *W* iff there is some *M** such that:

(i)   one *M**'s in *W*

(ii)   one's *M**-ing in *W* is not a part of one's *M*-ing at W

(iii)   at *W* it is true that, if one did not *M*, then one would still *M** in the same way and one would *E* and one's *M**-ing would help to bring about one's *E*-ing.

Clause (i) saves entering by the front door from superfluity, because when one enters by the front door, one does not also enter by the back door. It is true that when one enters by the front, one "also" (enters by the back or by the front) in some determinate, by-the-front, way. However, in violation of clause (iii), it is not true that if one did not enter by the front, one would still (enter by the back or by the front) *in that determinate, by-the-front, way*.[33] Clause (ii) saves giving both drugs from superfluity, on the assumption that giving Drug 1 is part of (giving Drug 1 and giving Drug 2).

It is crucial to distinguish the notion of a superfluous means from the notion of an "excessively costly" means. An *excessively costly* means is a means that is more costly than an alternative means that is at least as good a means.[34] Reason *does* transmit to excessively costly means. It is just that the reason transmitted is typically *outweighed* by the excessive cost. While giving both drugs is not superfluous, for instance, it may well be excessively costly. For example, it may use up more scarce medicine than giving Drug 1 alone, which is at least as good

at relieving the patient's pain. While there is as much reason for giving both as for giving Drug 1 alone—either option is equally good a means to relieving the patient's pain—there is stronger reason to refrain from giving both—it uses up more medicine.[35]

5. Why the Sufficiency principles have trouble with the nonsuperfluity desideratum
If Means Probabilize Nonsuperflously is correct, then it poses a problem for the Sufficiency principles. The Sufficiency principles may seem plausible, because if one takes some sufficient means, then one is *sure* to achieve the end. The problem is that compatibly with that truth, one may be sure to achieve the end *even if one does not* take that sufficient means. This can happen, in particular, because the sufficient means are superfluous. In such cases, at least if Means Probabilize Nonsuperfluously is on the right lines, no reason transmits to those sufficient means.

To make the case more exactly, let's use an analogue to our earlier definition of necessary means:

Definition of Sufficient Means: One's M-ing is a *sufficient* means to one's E-ing iff

(i) in *every* possible outcome in which one M's, one's M-ing helps to bring about one's E-ing, and

(ii) there is *some* possible outcome in which one M's.

So understood, a sufficient means can be superfluous in every possible outcome at which it is taken. Indeed, this is exactly what happens with Dr. Twoways. Drug 2 is a sufficient means to relieving the patient's pain. So, according to Weak Sufficiency, Twoways has reason to give Drug 2. However, since Drug 2 is superfluous toward the end of relieving pain in every possible outcome, Twoways intuitively has no reason to give Drug 2. And if Weak Sufficiency fails,[36] then so does Strong Sufficiency. If no reason transmits, then *a fortiori* reason at least as strong does not transmit.

If we assume Refrain Reasons-Ought or Relevant Reasons-Ought, then Ought Sufficiency also fails. Suppose that there is some reason to refrain from giving Drug 2. Then, since there is no reason for giving Drug 2, and some reason to refrain from giving it, Refrain Reasons-Ought implies that it is not the case that Twoways ought to give Drug 2.

However, the untenability of Ought Sufficiency is not very surprising. It is also falsified, in less exotic cases, by sufficient means that are excessively costly, even if not superfluous. Suppose we ought to visit my folks for Thanksgiving. Options *A* and *B* are sufficient means in the sense defined. But *A* involves a 36-hour, $5,000 flight with six layovers, where *B* involves only a six-hour, $150, direct flight. Presumably, it is not the case that we ought to take *A*.

One might object to this argument against Weak Sufficiency, on the grounds of:

Reasons Non-Necessity: One has reason to *X* only if it is possible that one does not *X*. Since is necessary that Twoways relieves the patient's pain (by giving Drug 1), Reasons Non-Necessity implies that he has no reason to do so. So Weak Sufficiency does not imply that he has reason to give Drug 2. However, Reasons Non-Necessity seems implausible. Knowing that I will, or even that I am fated to, refrain from committing suicide in the next minute, for example, doesn't undermine my judgment that I have reason to refrain. In any event, there are counterexamples to Strong and Ought Sufficiency, if not Weak Sufficiency, compatible with Reasons Non-Necessity. Suppose that Drug 1 fails in one out of every *N* cases, whereas Drug 2 works in every case. Even if it is necessary that Twoways gives Drug 1, it is not necessary that he relieves the pain. So, compatibly with Reasons Non-Necessity, there is reason to relieve the pain. Nevertheless, only weak reason transmits to the sufficient means of giving Drug 2, because in *N*–1 cases, giving Drug 2 is superfluous. By making the transmitted reason suitably weak (by the right choice of *N*), we can construct counterexamples to Strong and Ought Sufficiency.[37]

6. A positive proposal: General Transmission
Having criticized Strong and Ought Necessity, and the Sufficiency principles, I now suggest a

positive proposal:

> *General Transmission*: If there is reason for one to *E*, and there is positive probability,
>
> conditional on one's *M*-ing, that one's *M*-ing, or some part of one's *M*-ing, helps to bring
>
> it about that one *E*'s nonsuperfluously, then that is a reason for one to *M*, whose strength
>
> depends on the reason for one to *E* and the probability, *so long as the reason for one to E*
>
> *is not explained by an application of General Transmission to reason for one to achieve*
>
> *some distinct E'*.

Observe, first, that General Transmission meets our desiderata; it is, essentially, Means

Probabilize Nonsuperfluously plus a proviso. Second, General Transmission is informative: it

says something not only about whether *some* reason is transmitted, but also about *how much*

reason is transmitted and what it depends on.[38] Third, General Transmission is comprehensive: it

says whether reason transmits to *any kind* of means (not simply, say, necessary or sufficient

means).[39] Finally, General Transmission implies, as a special case, the one principle that

survived criticism: namely, Weak Necessity.

One might find it a shortcoming of General Transmission that it tells us nothing about the

instrumental transmission of 'ought.' But I suspect that this is as it should be. The problems

with Ought Necessity and Sufficiency discussed earlier suggest that there may be no interesting

generalities about the instrumental transmission of 'ought.' Perhaps all we can say, in general, is

that reasons transmit in the way that General Transmission describes, and that 'ought' depends

on reasons (perhaps in the way that Refrain Reasons-Ought and Relevant Reasons-Ought

suggest). Whether, in any given case, the reason transmitted from an end to means will ensure

that the 'ought' that accompanies the end will also accompany the means may just depend on the specifics of the given case.

The reason for the italicized "Intransitivity Proviso" is a problem explored by Millsap (ms. a) (who credits it to Simon May and Kenny Easwaran) and Bedke (2009: 679 n. 12) (who credits it to Jamie Dreier). In an earlier version of Millsap's example, Kenny has reason to prepare a display on the life of Marie Antoinette, which requires that he bake a loaf of bread and a cake. General Transmission might well imply that Kenny has reason to bake a loaf of bread. One means to that is to bake a normal loaf. Another is to bake a mega-loaf, which is very likely to leave no flour left to bake a cake, and so to prevent Kenny from preparing the display. Applying General Transmission with E = baking a loaf, it might well be that Kenny has the same reason to bake a mega-loaf as to bake a normal loaf. However, applying General Transmission with E = preparing the display, Kenny will presumably have much less reason to bake a mega-loaf as to bake a normal loaf. Intuitively, it is the latter conclusion, not the former, which is correct. The Intransitivity Proviso reflects this, by saying that reason for a means to a means is transmitted directly from the "ultimate" end, not from the means to which it is a means.[40]

Some might urge a second, "Silencing Proviso": "…and the reason for one to refrain from M-ing is not overwhelmingly more weighty than the reason that there is, or would otherwise be, for one's M-ing." This would accommodate an objection that is often raised against transmission principles. Suppose the end is slightly improving college policy, which, at least at first glance, is something that I have reason to do. The only roadblock is the kindly old don who is sure to veto any proposed change. Suppose there's a positive probability, conditional on poisoning the don, that I change college policy and poisoning the don helps to bring this about in a nonsuperfluous way. General Transmission implies that there is *some* reason for me to

poison the don—although obviously overridden. But some will say that reason for the means is "silenced." Yes, when means carry some moderate cost, reasons for them may be overridden. But when means are overwhelmingly costly or objectionable in this way, there isn't even any reason to be overridden. Transmission in such cases is simply blocked.[41]

Like many others, however, I find the objection unpersuasive. I grant that it is odd to *say* that there is reason to poison the don. And I grant that a fully virtuous agent would not *treat* the improvement to college policy as a reason for poisoning the don. That thought would not so much as enter into his deliberations. But I doubt that we should conclude from these observations that there is no such reason. As Millsap (ms. a), Raz (2005b: 3), and Bedke (2009: 684–686), observe, a point made by Schroeder (2004, 2005, and 2007) seems to explain why it is odd to *say* that there is reason to poison the don. It is a violation of the Maxim of Relation of Grice (1989) to say that there is a reason, when, as one knows, it is very weak or vastly outweighed. And, to entertain the thought, at least in the context of live deliberation, does indeed indicate a vice. (Doesn't it go without saying, or thinking, that the agent shouldn't kill the don, in which case entertaining thoughts about the reason for doing so is idle? So why is he entertaining thoughts about them? Is the verdict somehow not obvious to him? Or is he somehow tempted to defy it?) But, as Seidman (2005) observes, the viciousness of entertaining the thought is compatible with its truth. It similarly indicates a vice of one kind to entertain lascivious but true thoughts, during wedding vows, that sex with the officiant would be pleasurable, and a vice of another kind to entertain distracting but true thoughts, during oral argument, that it is time to treat oneself to a new judicial robe. Thus, I find it more natural to conclude that, when taking some means is opposed not simply by reasons that would be *moderately overriding*, but moreover by *overwhelmingly strong* reasons, this makes it odd (or

27

vicious) to cite (or contemplate) reason for the means, than to say that it somehow blocks transmission of reason. But for those who are troubled by the objection, the option of adding the Silencing Proviso is there.[42]

7. Implications: No reasons to be instrumentally rational
At least for the sake of argument, suppose that the cases against Strong and Ought Necessity, and against the Sufficiency principles, and for General Transmission, have succeeded. What then follows for appeals to principles of instrumental transmission to show that "narrow-scope" reason or 'ought' for the consequent F "detaches" from "wide-scope" reason or 'ought' for a material conditional (one E's $\supset$ one F's)? Recall that detachment threatens the claim of Broome (1999) that, in general, one ought to be instrumentally rational: that is, that one ought to see to it that (one intends at t to E and believes at t that intending at t to M is a necessary means to E-ing $\supset$ one intends at t to M). For it scarcely seems to follow from the facts that one intends to start World War III and that one believes that intending to launch this missile is necessary for starting World War III, that one has "detached," "narrow-scope" reason to intend to launch the missile, let alone that one ought to.

Setiya (2007) and Schroeder (2009), revisiting an observation of Greenspan (1975), advocate:

> *Ought Detachment*: If one ought to make it the case that (one E's $\supset$ one F's) and one
> cannot alter the fact that one E's, then one *ought* to F.

According to Ought Detachment, Broome's (1999) claim implies that if one could not alter the intention and means-end belief above, one ought to intend to launch the missile. However, Setiya's and Schroeder's arguments for Ought Detachment rely on Ought Necessity, which we have questioned. So it not clear that we have grounds for accepting Ought Detachment.

To a similar purpose, Raz (2005a) argues for:

Weak Detachment: If one has reason to make it the case that (one *E*'s ⊃ one *F*'s), then

one has reason to *F*,

which he takes to follow not from the Facilitative Principle, but instead from the more basic and general idea that: "People have reason to do what will bring them into conformity with the reasons that apply to them" (2005b, 3). Bedke (2009) argues that something similar follows from his "Instrumental Principle." If they are right about this, then we need no auxiliary assumptions, about unalterable intentions or anything else, to refute Broome (1999). If there were reason to be instrumentally rational, there would *always* be reason to intend to *M*, for any *M* whatsoever.

But, first, the case *for* Weak Detachment rests on questionable principles. Since Bedke's Instrumental Principle entails Weak Sufficiency, it would seem to be falsified by the phenomenon of superfluity. The "more basic" idea of Raz's quoted above also seems falsified by superfluity. Presumably we do not have reason to do what *superfluously* brings us into conformity with what we have reason to do.

Second, there is an argument *against* Weak Detachment. Rippon (2011: 12) notes that Weak Detachment seems to prove too much. There are arguably some wide-scope reasons: for example, that I have reason (I am in France ⊃ I speak French). Now suppose that I'm not in France; I'm in the USA, surrounded by English monolinguists. Weak Detachment implies that I have reason to speak French. But intuitively I have no reason at all to do this.

Nevertheless, I believe that we can derive from General Transmission analogues to Ought and Weak Detachment that vindicate Setiya's and Schroeder's, and Raz's and Bedke's, lines of argument. Let us make two plausible assumptions. First, one's *F*-ing helps to bring it about that

(one E's ⊃ one F's). Second, when the probability in General Transmission is one, then the

reason for the means is of the same weight as the reason for the end. Then General Transmission

implies:

> *Strong Detachment*: If one has reason to make it the case that (one E's ⊃ one F's), and, in
>
> every possible outcome, one E's, then that is a reason *of the same weight* to F.

This is because, when the antecedent holds, F-ing helps to bring about the end nonsuperfluously

in *every* outcome in which one F's. Strong Detachment implies that if in every possible outcome

one has the intention and means-end belief above, and if one has reason to be instrumentally

rational, then one has at least as much reason to intend to launch the missile, which seems absurd

enough.[43]

> General Transmission also validates:
>
> *Restricted Weak Detachment*: If one has reason to make it the case that (one E's ⊃ one
>
> F's), and there is positive probability, conditional on one's F-ing, that one's F-ing helps
>
> to bring it about in a nonsuperfluous way that (one E's ⊃ one F's), then one has some
>
> reason to F.[44]

This seems to suffice for Raz's and Bedke's purposes. It entails that if at some possible outcome

intending to launch the missile is not superfluous to being instrumentally rational, and if one has

reason to be instrumentally rational, then one has reason to intend it.

Moreover, Restricted Weak Detachment handles Rippon's counterexample, even if we

grant the (perhaps dubious) claim there is a wide-scope reason to see to it that (I am in France ⊃

I speak French). When we suppose, in entertaining Rippon's counterexample, that I am not in

France, we are in effect shifting to a context in which there is no (epistemic[45] or historical)

probability that I am in France. But, in that case, the material conditional (I am in France ⊃ I

speak French) is already guaranteed to be satisfied, which makes speaking French a *superfluous* means to satisfying it (as Rippon 2011: 15 n. 31 in fact observes). And *Restricted* Weak Detachment does *not* say that reason transmits to the consequent when satisfying the consequent is a *superfluous* means to satisfying the conditional. So it does *not* follow from *Restricted* Weak Detachment that, when I am not in France, I have reason to speak French. In other words, Rippon's case simply exemplifies the general point that reasons don't transmit to superfluous means, as General Transmission says.

In reply, it might be said that if we add the Silencing Proviso of section 6 to General Transmission, then Strong Detachment and Restricted Weak Detachment will inherit it. Since the reasons to refrain from intending to launch this nuclear missile are presumably overwhelming, no reason transmits to intending to launch it. Alternatively, one might appeal to Schroeder's pragmatic point (discussed in the last paragraph of the previous section) to question the intuition that it is absurd for reason to transmit to intending to launch it. However, other bootstrapping cases will escape this reply: namely, cases in which the there is weaker reason for intended end than for some alternative, but not so much weaker that the reason for the intended end is "silenced," or becomes pragmatically odd to cite. Presumably, we cannot insist that whenever there is stronger reason for an alternative, the reason is silenced, or becomes pragmatically odd to cite, on pain of denying that there are ever overridden reasons, or pragmatically acceptable mention of them.

8. Conclusion: Is instrumental transmission a myth?
I close by noting a further implication of General Transmission: namely, that it may pave the way for the conclusions, first, that, strictly speaking, there is no such thing as instrumental

transmission and, second, that whatever truth there is in the vicinity is just a consequence of a certain kind of decision theory.

We have been supposing that there are reasons for ends. But where do these come from? On a value-based view, at least *some* reasons for ends are presumably provided by *valuable* states of affairs. And on a desire-based view, *all* reasons for ends are provided by *desires* for states of affairs. So says:

General Production: If S is a valuable (or desired) state of affairs, and if there is positive probability, conditional on one's X-ing, that one's X-ing, or some part of one's X-ing, helps to bring about S "nonredundantly"—that is, helps to bring about S when S would not have obtained if one had not X-ed—then that is a reason for one to X, whose strength depends positively on the value of (or strength, depth, etc. of desire for) S and the probability.

The need for "nonredundantly" can be seen from a variant of the Twoways case. Here a *different* doctor has already given the patient Drug 1. Has Twoways reason to give Drug 2? Suppose that the valuable state of affairs is simply that the patient's pain is relieved, it does not matter who brings it about. (In some cases the value of a state of affairs may depend on who brings it about, but this is not such a case.) Although there is some probability that Twoways helps to bring it about that the patient's pain is relieved, there is no probability that Twoways helps to bring this about nonredundantly. So General Production says, as seems intuitive, that he has no reason to give Drug 2. Crucially, nonredundancy differs from nonsuperfluity. The *redundancy* of an action with respect to a state of affairs is a matter of whether that *state of affairs would have obtained* (by one's agency, by someone else's agency, by natural causes, etc.) even if one had not performed that action. By contrast, the *superfluity* of an action with respect to an end is

(roughly) a matter of whether *one would have achieved a given end* (by, specifically, one's own agency) even if one had not performed that action.

"Now assume that General Production explains *all* reasons for ends," one might challenge. "Then wouldn't it *also* explain *all* reasons for means, without needing any help from General Transmission? Suppose there is reason for an end. By assumption, this will be because it helps to bring about some valuable (or desired) state of affairs. But then any means to that end will also be likely to help to bring about that valuable (or desired) state of affairs. And if so, won't General Production imply there will be a reason for those means too? General Production, thus, explains General Transmission. The latter becomes, as it were, minor theorem of a kind of decision theory.[46] Indeed, to be more accurate, instrumental transmission is not so much *explained* by General Production, as explained *away*. For now the reason for the means is simply *that it helps to bring about a valuable state of affairs*. *Reason for the end* appears nowhere in the explanation. It simply drops out. Of course, we can still say that whenever there is reason for the end, there is also reason for the means. But we can also say, with equal justice, that whenever there is reason for the *means*, there is reason for the *end*. The idea of instrumental transmission—that there is reason for the means *because there is reason for the end*—is an illusion."

This conclusion—that instrumental transmission is a myth—may be less damaging to the argument of this paper than it might at first appear. We could change General Transmission to accommodate it. We would rewrite "…then *that* is a reason…" as "…then *there is also a* reason …" So revised, General Transmission would still be a *useful* truth. It would still help, for example, to clarify the debate over "detachment." And it would be a *neglected* truth, overlooked, in particular, by proponents of extant transmission principles. The fact that this

revised General Transmission would be a minor theorem of decision theory would not change this.

However, I have doubts that the challenge succeeds. First, at least on a value-based view, the initial assumption of the challenge is far from obvious. On a value-based view, one might deny that all reasons for ends are provided by the value of states of affairs, and so explained by General Production.[47] For example, I may have reason to honor or respect something of value, whether or not doing so brings about a valuable state of affairs.[48]

Second, even if all reasons for *ends* were explained by General Production, there would still be reasons for *means* that were *not* explained by General Production. The trouble is that, paradoxically, something can be a *means to the end* of helping to bring about a valuable (or desired) state of affairs nonredundantly (and so, intuitively, reason can transmit to it) without *itself* helping to bring about that state of affairs *nonredundantly* (and so without, as far as General Production is concerned, any reason transmitting to it). Suppose that the valuable (or desired) state of affairs is my being in the house, regardless of how this comes about. In every outcome in which I enter by the front, if I hadn't, I would have entered by the back. As we saw in section 4, entering by the front is *not a superfluous* means. So, according to General Transmission, I have reason to enter by the front, as seems intuitive. However, my entering by the front is a *redundant* means, because in every outcome in which I enter by the front, even if I hadn't entered by the front, the state of affairs of my being in the house (it matters not how) would have obtained. So, according to General Production, I would have *no* reason to enter by the front, which seems false.[49] In sum, even when one's end is one's helping to bring about a valuable or desired state of affairs (whose value or status as desired does not depend on how it comes about), we still need to think of means as helping to bring about *that end*: *one's* bringing

about that state of affairs *oneself*. We can't bypass the reason for end, and treat the means simply as something that helps to bring about the obtaining of that valuable or desired state of affairs. So, contrary to the challenge, reasons for ends don't drop out of the explanation of reasons for the means.

This suggests, to my mind, that we should keep General Transmission in its original form and, more generally, that instrumental transmission is not a myth. But what is perhaps more remarkable is not that this challenge can be countered, but that it could be raised in the first place. That is itself surprising—or at least surprised me. If nothing else, it testifies to the importance of getting clear about the principles governing (to put it neutrally) the relations between reasons for ends and means. Whatever comes of the challenge, I hope that this paper is some means in furtherance of that end.

References

Anderson, E. (1993), *Value in Ethics and Economics* (Harvard).

Anscombe, G.E.M. (1957), *Intention* (Basil Blackwell).

Bedke, M. (2009), "The Iffiest Oughts," in *Ethics* 119: 672–698.

Bedke, M. (ms.), "Ends to Means"

Bratman, M. (1987), *Intentions, Plans, and Practical Reason* (Harvard).

Bratman, M. (2009), "Intention, Practical Rationality, and Self-Governance," in *Ethics* 119: 411–443.

Broome, J. (1999), "Normative Requirements," in *Ratio* 12: 398–419.

Broome, J. (2002), "Practical Reasoning," in J. Bermùdez and A. Millar (eds.), *Reason and Nature: Essays in the Theory of Rationality* (Oxford), 85–111.

Broome, J. (2005), "Have We Reason to Do as Rationality Requires?—A Comment on Raz," in
Journal of Ethics and Social Philosophy, Symposium 1:1–8.

Broome, J. (ms.), "The Facilitating Principle."

Brunero, J. (2010), "Self-Governance, Means-End Coherence, and Unalterable Ends," in *Ethics*
120: 579–591.

Cariani, F. (2009), *The Semantics of 'Ought' and the Unity of Modal Discourse*, Ph.D.
Dissertation, University of California, Berkeley.

Darwall, S. (1983), *Impartial Reason* (Cornell).

Darwall, S. (2006), *The Second-Person Standpoint* (Harvard).

Dreier, J. (2001), "Humean Doubts about the Categorical Imperative," in E. Miglram (ed.), *The
Varieties of Practical Reasoning* (MIT), 27–48.

Finlay, S. (2009), "Oughts and Ends," in *Philosophical Studies* 143: 315–340.

Finlay, S. (2010), "What *Ought* Probably Means, and Why You Can't Detach It," in *Synthese*
177: 67–89.

Goble, L. (1996), "Utilitarian Deontic Logic," in *Philosophical Studies* 82: 317–357.

Grice, H. P. (1989), "Logic and Conversation," in *Studies in the Way of Words* (Harvard), 22–40.

Hall, N. (2004), "Two Concepts of Causation," in J. Collins, N. Hall, and L.A. Paul (eds.)
Causation and Counterfactuals (MIT), 225–76.

Horty, J. (2001), *Agency and Deontic Logic* (Oxford).

Hubin, D. (1999), "What's Special about Humeanism," in *Noûs* 33: 30–45.

Jackson, F. (1985), "On the Semantics and Logic of Obligation," in *Mind* 94: 177–195.

Jackson, F. and Pargetter, R. (1986), "Oughts, Options, and Actualism," in *The Philosophical
Review* 95: 233–255.

Kenny, A. (1966), "Practical Inference," in *Analysis* 65–75.

Kiesewetter, B. (forthcoming), "Instrumental Normativity: In Defense of the Transmission Principle," in *Ethics*.

Kolodny, N. (2007), "How Does Coherence Matter?," in *Proceedings of the Aristotelian Society* 107: 229–263.

Kolodny, N. (2008), "Why Be Disposed to be Coherent?," in *Ethics* 118: 437–463.

Kolodny, N. and MacFarlane, J. (2010), "Ifs and Oughts," in *Journal of Philosophy* 107: 115–143.

Kolodny, N. and MacFarlane, J. (ms.), "Ought: Between Objective and Subjective."

Korsgaard, C. (1997), "The Normativity of Instrumental Reason," in G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason*, (Oxford), 215–254.

Korsgaard, C. (2009a), *Self-Constitution* (Oxford).

Korsgaard, C. (2009b), "The Activity of Reason," in *Proceedings and Addresses of the American Philosophical Association* 83: 23–43.

McDowell, J. (1998a), "The Role of *Eudaimonia* in Aristotle's Ethics," in *Mind, Value, and Reality* (Harvard), 3–22.

McDowell, J. (1998b), "Virtue and Reason," in *Mind, Value, and Reality* (Harvard), 50–73.

McDowell, J. (1998c), "Are Moral Requirements Hypothetical Imperatives?" in *Mind, Value, and Reality* (Harvard), 77–94.

Millsap, R. (ms. a), "Three Kinds of Reasons Transmission."

Millsap, R. (ms. b), "How to Avoid an Explosion of Reasons."

Millsap, R. (ms. c), "The Balancing Theory of Ought and Reasons Transmission."

Nagel, T. (1970), *The Possibility of Altruism* (Princeton).

Parfit, D. (ms.), "What We Together Do."

Price, A. (2008), *Contextuality in Practical Reason* (Oxford).

Raz, J. (2005a), "The Myth of Instrumental Rationality," in *Journal of Ethics and Social Philosophy* 1: 1–28.

Raz, J. (2005b), "Instrumental Rationality: A Reprise," in *Journal of Ethics and Social Philosophy*, Symposium 1:1–19.

Raz, J. (2012), *From Normativity to Responsibility* (Oxford).

Rippon, S. (2011), "In Defense of the Wide-Scope Instrumental Principle," *Journal of Ethics and Social Policy* 5.

Ross, A. (1941), "Imperatives and Logic," *Theoria* 7: 53–71.

Scanlon, T.M. (1999), *What We Owe to Each Other* (Harvard).

Scanlon, T.M. (2004), "Reasons: A Puzzling Duality?" in R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.)*, Reasons and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford), 231–246.

Scanlon, T.M. (2007), "Structural Irrationality" in G. Brennan, R. Goodin, F. Jackson and M. Smith (eds.), *Common Minds: Essays in Honor of Philip Pettit* (Oxford), 84–103.

Schroeder, M. (2004), "The Scope of Instrumental Reason," in *Philosophical Perspectives* 18: 337–64.

Schroeder, M. (2005), "Instrumental Mythology," in *Journal of Ethics and Social Philosophy, Symposium* 1: 1–12.

Schroeder, M. (2007), *Slaves of the Passions* (Oxford).

Schroeder, M. (2009), "Means-End Coherence, Stringency, and Subjective Reason," in *Philosophical Studies* 143: 223–248.

Seidman, J. (2005), "Two Sides of Silencing," in *Philosophical Quarterly* 55: 68–77.

Setiya, K. (2005), "Is Efficiency a Vice?" in *American Philosophical Quarterly* 42: 333–339.

Setiya, K. (2007), "Cognitivism about Instrumental Reason," in *Ethics* 117: 649–673.

Snedegar, J. (2013), "Contrastive Semantics for Deontic Modals," in M. Blaauw (ed.),
    Contrastivism in Philosophy (Routledge), 116–133.

Stegenga, J. (2012), "Probabilizing the End," in *Philosophical Studies* 165: 95–115.

Vranas, P. (2011), "New Foundations for Imperative Logic: Pure Imperative Inference," in *Mind*
    120: 369–446.

Wallace, R. J. (2006): "Normativity, Commitment, and Instrumental Reason" (with postscript),
    in *Normativity and the Will* (Oxford), 82–120.

Way, J. (2010), "Defending the Wide-Scope Approach to Instrumental Reason," in
    Philosophical Studies 147: 213–33.

Wedgwood, R. (2006), "The Meaning of Ought," in *Oxford Studies in Metaethics* 1: 127–160.

Wedgwood, R. (2011), "Instrumental Rationality," in *Oxford Studies in Metaethics* 6: 280–309.

Yalcin, S. (2007), "Epistemic Modals," in *Mind* 116: 983–1026.

Department of Philosophy at the University of Texas, the Department of Philosophy at Harvard University, the Workshop on Reasons and Rationality at the London School of Economics, and the 2011 Central Meeting of the APA.

[2] Nagel (1970: 34) observes that it is a "perfectly general property" of reasons for action "that they transmit their influence over the relation between ends and means."

[3] As Kant claimed in the *Groundwork*, and as the idea has been developed (in markedly different ways) by Korsgaard (1997), Bratman (1987), and Broome (2002).

[4] This formulation is agnostic on whether reasons for ends derive from our desiring those ends, or from the relation of those ends to things of independent value. However, desire-based theorists may deny, against Hubin (1999), that their theory is a combination of a principle of instrumental transmission and the principle that reasons for ends are provided by desires. Instead, they may say, there is just one principle, a principle of, if you will, instrumental transmutation: if one desires the end, then one has reason to take the means. See the discussion of General Production, in section 8, for a doubt about this.

[5] See note 15 and the preceding text.

[6] See, for example, Bedke (2009), Bratman (2009), Darwall (1983: 16, 46–48 and 2006), Hubin (1999), Raz (2005a and 2005b), Kolodny (2007 and 2008b), Price (2008: 7, 138–139), Schroeder (2009), Setiya (2007), Vranas (2011) (who distinguishes it from the kinds of support relevant to the validity of imperative inference), and Way (2010). Passages in Scanlon (2004) and Wallace (2006) assume or grant its correctness.

[7] "If you should do E, all things considered, and doing M is a necessary means to doing E, you should do M, all things considered" (Setiya 2007: 660). "If X objectively ought to do A, and to do A X must do B, it follows that X objectively ought to do B" (Schroeder 2009: 239). The principle

is often implicitly invoked: "because you promised to do A and cannot do A without doing B, you ought to do B" (Korsgaard 2009b: 38).

Ought Necessity is also entailed by certain deontic logics and semantics for 'ought.' In broad outline, and expressed in the idiom of this paper, these views hold that "One ought to *X*" is true iff one *X*'s at every possible outcome that is "ideal." (See Wedgwood 2006 for a particularly forceful defense of such views.) The argument in section 3 against Ought Necessity is, in effect, one objection to such views. (Interestingly, Horty 2001 sec. 3.4.2 presents a similar argument against such views, although not against Ought Necessity as such.) Moreover, since such views transmit 'ought's not only to necessary means, but also to necessary conditions, they may be subject to further objections. For instance, there's Ross's (1941) paradox: that "One ought (to post the letter)" entails "One ought (to post the letter or to burn it)," an entailment whose strangeness, as Cariani (2009) shows, stubbornly resists pragmatic explanation. There's also our telegraphing boxer of section 1. If in every possible outcome in which he throws the punch he telegraphs it, then "He ought to throw it" entails, counterintuitively, "He ought to telegraph it."

One might think that the "end-relative" account of 'ought' presented by Finlay (2009) and (2010) is also committed to Ought Necessity, since that account implies that, whatever else one ought to do in a given context, one ought to take necessary means to the "end" specified by that context. However, the *E* in Ought Necessity need not be this contextually specified end itself. If the contextually specified end is, say, maximizing expected value (see Finlay 2009: 326 n. 26), then it might be the case for Unlucky, discussed in section 3, that he ought to save the antique (=*E*), but not the case that he ought to take the necessary means of taking the taxi.

[8] Note that the explanatory connection indicated by the phrase, "then because of that," in Ought Necessity is here indicated by "then that is a reason."

"Reasons for me to make something my end are, owing to the hypothetical imperative, equally reasons for me to take the necessary means to it" (Darwall 1983: 16). "If one has conclusive reason to believe that one will *E* only if one *F*s, then one has reason to *F* that is at least as strong as one's reason to *E*" (Kolodny 2007: 251). "If *X* has an objective reason to do *A* and to do *A X* must do *B*, then *X* has an objective reason to do *B* of equal weight to *X*'s objective reason to do *A*" (Schroeder 2009: 245). "If *R* is a practical reason in favor of *X*, *X* is attainable by the agent, and *M* is a necessary means to or necessary constitutive element of *X*, then *R* is a practical reason in favor of *M*" (Bratman 2009: 424). "If you have a reason to do *A* and doing *B* is a necessary means to doing *A*, you have a reason to do *B* which is at least as strong as your reason to do *A*" (Way 2010: 225). See also Millsap (ms a).

[9] Way (2010: 224) and Bedke (2009: 687 n. 10) endorse Weak Sufficiency. The "(M/E)" principle of Dreier (1997: 38) has a similar structure. The "logic of satisfactoriness" of Kenny (1966) validates an analogue of Ought Sufficiency for "fiats."

[10] Anscombe (1957: 36) anticipates these problems when she observes that the intention expressed by, "I do *P*, so that *Q*," can be "contradicted" either by saying, "But *Q* won't happen, even if you do *P*" or by saying, "But it will happen whether you do *P* or not." In my terms, *P* can fail to inherit the reason to *Q* because *P* fails to "probabilize" *Q*—*Q* won't happen even if you do *P*—or because *P* is "superfluous" with respect to *Q*—roughly, *Q* will happen whether or not you do *P*.

[11] Alternatively, one might suggest that the weight of the reason depends on:

(1) the *difference* of the probability, conditional on one's *M*-ing, that one *E*'s *less the probability that one E's*;

or:

(2) the *difference* of the probability, conditional on one's *M*-ing, that one *E*'s *less* the probability, conditional on one's *not M*-ing, that one *E*'s;

or:

(3) the *difference* of the probability, conditional on one's *E-ing*, that one *M*'s *less* the probability*, conditional on one's *not-E-ing,* that one *M*'s (i.e., (2), with *E* and *M* changing places);

or:

(4) the *difference* of the probability, conditional on one's *M*-ing, that one *E*'s *less* the probability*, conditional on the worst one could do with respect to E-ing,* that one *E*'s.

I reject (1) on the grounds that it implies that the probability of one's *M*-ing can affect one's reason to *M*. At the limit, if one is sure to *M*, then one has no reason to *M*, since the conditional probability equals the corresponding unconditional probability.

(2) also gives the wrong results. Suppose the only relevant end is keeping the patient alive. Giving drug *A* alone improves the patient's chances of living by 98 percentage points, and giving drug *B* in addition improves them by a further percentage point. Intuitively, the doctor has *no less* reason to give both *A* and *B*, which has a 99% chance *exactly* of keeping the patient alive, than to give the patient at least *A*, which has *at most* a 99% chance. However, giving both *A* and *B* rather than not giving both *A* and *B* (assuming the doctor is very likely to give *A* but not *B*, if he does not give *A* and *B*) raises the probability by something approaching only one percentage point, whereas giving at least *A* rather than not giving at least *A* raises the probability

by something approaching 98 percentage points. So the current proposal implies, counterintuitively, that the doctor has *less* reason to give both *A* and *B* than to give at least *A*.

Stegenga (2012) proposes option (3), along with a variant structurally analogous to a ratio of likelihood functions. While Stegenga's application of confirmation theory to these issues is welcome, I have doubts that his proposal suits the practical context. We don't seem to practically deliberate in terms anything like (3). And some of the implications that Stegenga suggests intuitively support (3) seem to me to tell against it. For example, no (as opposed to little) reason to buy a lottery ticket transmits from reason to get rich. And, to take an extreme version of a case that he considers, no reason transmits to studying from reason to pass the exam, if one is sure to study, as was the case with (1).

Bedke (ms.) pursues option (4). It has an appealing simplicity, but I worry that it sacrifices relevant structure. To anticipate (and to vary slightly) an example from section 4, suppose that the end is my entering, and no matter what, either I enter by the front or I enter by the back. Then the worst I can do by way of entering is just as good as the best I can do. (4) implies, counterintuitively to my mind, that I have no reason to enter by the front. (Note that (1)–(3) also have this implication, with undefined probabilities in the case of (3).) There seems to me an intuitive difference between my reason to enter by the front in this case and my reason, say, to immobilize myself by tying my ankle to a stake in the front lawn in an altered case in which that is possible. If you ask me, "Why did you do that? What reason had you to?" I could give an intelligible answer in the first case, but not in the second. However, I suspect that my approach and Bedke's (4) can be made to deliver the same 'ought'-judgments, at least relative to a given Reasons-Ought principle of the kind introduced in section 3. Moreover, the choice

between these approaches does not affect the criticism of the Necessity and Sufficiency

principles in sections 2, 3, and 5, or the implications discussed in sections 7 and 8.

[12] See Hall (2004) for examples of preventers of preventers and for an argument that, on one

conception of causation, they are not causes. It might seem that "helping to bring about" as I am

understanding it here is just the other conception of cause that Hall distinguishes: roughly,

counterfactual dependence among distinct items. But this would not cover cases of, say,

constitutive means, which are not suitably distinct from the end.

[13] Bedke (ms.) and Stegenga (2012) consider cases like the telegraphing boxer, but resist adding

a "causal" condition. I believe that the counterexamples that, according to Bedke, tell against a

"causal" condition can be addressed by the point that preventing preventers can help to bring

about and by the nonsuperfluity condition of section 4. Stegenga suggests that the "causal"

condition is not necessary, so long as we specify the relevant "population" appropriately: in this

case, not this boxer, but all boxers. But this seems to me misguided. Why shouldn't our advice

by tailored to *this* boxer given what we know about *his* idiosyncracies?

[14] Insofar as we use the epistemic framework, we may need some response to Newcomb's

problem. Adding the phrase "helps to bring about" is, in itself, no response, since both one-

boxing and two-boxing help to bring about the relevant payoffs. One response, which would not

radically affect the discussion, would be to screen off whatever information one's M-ing gives

about the probability of factors that (i) bear on whether one's M-ing helps to bring about one's E-

ing, but that (ii) one's M-ing itself does not help to bring about. One method for doing this

would be to take a partition of the possible states, S_i, that one's M-ing does not help to bring

about. (These are not quite states of "nature," since they may depend on other things, besides M-

ing, that one may independently do. We are not assuming that M-ing is an "exhaustive" action in

the sense described in section 3.) Then we sum, over i, the products of the probability,

unconditional on one's M-ing, of S_i, and the probability, conditional on M-ing in S_i, that one's M-ing (or a part thereof) helps to bring it about that one E's in S_i. (Note that this formula still needs

the phrase "helps to bring it about that." Without it, the formula might deliver a *positive*

probability, on telegraphing, of connecting, and so imply, wrongly, that there is reason to

telegraph.) Other responses to Newcomb's problem, however, might more radically affect our

discussion.

[15] As for the question of *how* context determines the relevant body of information, I favor the

view of Kolodny and MacFarlane (ms.), that the truth of an occurrence of an 'ought' or 'reason'

sentence depends on the information relevant at a context of *assessment*: to a first approximation,

the information of the person—agent, advisor, onlooker, etc.—who is considering the

occurrence.

[16] See Belnap, Perloff, Xu (2001) and Horty (2001). On such a framework, "outcomes" would

be more naturally understood in terms of histories, instead of worlds. If agents are conceived as

choosing "freely" in a strong sense, probabilities will be unavailable whenever the occurrence of

the end depends some further "free" choice. Nevertheless, "sure-thing" reasoning may still be

possible, as Horty shows.

[17] Raz (2005a: 7) makes essentially this point with the case of buying a ticket to visit one's

grandmother on Easter Island when the airline workers are on strike.

[18] For discussion of similar principles, see Bart Streumer, "Reasons and Ability," in this volume.

[19] When "possible" is interpreted as "epistemically possible," this implies, troublingly, that when

it is known that one will do something horrible, there is no reason for one to refrain from doing

it. So, on the epistemic framework, one may wish to replace "epistemically possible" with

"epistemically possible is historically and nomologically possible." This has a somewhat surprising implication, though. If it is known that one will not *X*, although one historically and nomologically could *X*, then whether one has reason to *X* is assessed *not* relative to what is known, but instead relative to a *weaker* body of information: roughly, the "closest" body where it is not known that one will not *X*.

[20] One might also naturally describe the cost as a "reason not to taxi" or a "reason against taxi-ing." For simplicity, I assume that reasons not to *X* and reasons against *X*-ing are just equivalent to reasons to refrain from *X*-ing. This assumption may well be false, but I don't think it affects the argument in any substantive way. See Bedke (ms.) for discussion.

[21] Insofar as Schroeder (2009: 246) takes his "General Reason Transmit" ("If *X* has… reason to do A and *X*'s doing *B* would facilitate her doing *A*, then *X* has… reason to do *B* of weight at least proportional to *X*'s… reason to do *A*, and to how well her doing *B* would facilitate her doing *A*") to entail Strong Necessity, it also seems to face this problem.

[22] Way (2010: 225 n. 32) reports an objection to Ought Necessity along these lines from John Broome.

[23] For a recent survey and defense of such views, see Snedegar (2013).

[24] This is similar to the principle "Negative" in Millsap (ms. c).

[25] This might be a special case of General Production of section 8.

[26] Compare the "Combo" example of Millsap (ms. c).

[27] One might suggest replacing "necessary" with "sufficiently likely" in Strong Availability. But, among other difficulties, this means that in order to know whether Lucky ought to save the antique, he first needs to decide whether he will take the taxi. This seems to get instrumental reasoning exactly backwards.

[28] One might reply that the instrumental reason derives directly from the *value of the state of affairs* of the antique's being saved, instead of from any reason for the *end* of one's saving it. To anticipate just one of two problems with this reply, which will be discussed in section 8, there may not be any such valuable state of affairs. It might be the case, to alter the scenario, that the reason to save the antique is not that one's saving it would *itself* be a valuable state of affairs, but instead that one's saving it would honor something, or someone, of value, e.g., the memory of Unlucky's grandmother, who schlepped the antique from the old country. (Hence the awkward example of saving an antique.)

[29] Independent reasons for accepting Strong Availability might seem to be provided by the following puzzle. Suppose ten miners are trapped in one of two shafts, we know not which, with the floodwaters rising. If we block shaft A, all ten will survive if they are in A, whereas all will die if they are in B. If we block shaft B, the reverse. If we block neither shaft, nine will survive and one will die wherever they are. It seems that we ought to block neither, which will save nine wherever they are. But why ought we not instead block *whichever shaft the miners are in*? After all, *that* is sure to save *all ten*, wherever they are. As Kolodny and McFarlane (2010) observe, something like Strong Availability would give us an answer. Because there is nothing under our control, that we can knowingly do, that is sure to bring it about that we block the shaft the miners are in. So doing so is not available to us.

However, I think it is right to say both "We ought to block neither shaft" and "We ought to block the shaft that they are in." Each is correct relative to a different set of alternatives. Where the question is, "Which, among the *decisions* we can now make, ought we to make?" blocking neither shaft is the decision to make. Deciding to block the shaft that they are in, without knowing which this is, would have very bad results. Where the question is, "Which,

among the *actions* we can now perform, ought we to perform?" blocking the shaft that they are in is the action to perform.

Alternatively, there is a proposal suggested by remarks of Kiesewetter (forthcoming) and Simon Rippon in correspondence:

> *X* is available to one at *t* only if: if one does everything one ought to do up until the time, *t'*, when one would *X*, then at *t'* *X*-ing is fully under one's control.

In many cases, this collapses into Super-Strong Availability. For example, in the case of Buridan's ass, before the ass sets out, neither eating from the left bale, nor eating from the right bale is available to him (at least if retracing steps is not allowed). In heading for the left bale, the ass does everything that he ought to do, but he cannot then eat from the right bale. Same if he heads for the right bale. In any event, in the Mountaineer case, below, stepping while not looking down still counts as available by this proposal.

Some complain that this "actualist"-leaning view gives agents oughts that they cannot jointly satisfy. (See Kiesewetter (forthcoming), for example.) This complaint is not quite accurate, though, if we accept *Relevant Reasons-Ought*. *Relative to a fixed question*, it does *not* give agents oughts that they cannot jointly satisfy. It is only across *different* questions that it gives agents incompatible oughts. In itself, this is not so troubling. If I ask: "Of the holidays in the mountains, which ought I to make? And of the holidays at the beach, which ought I to make?" I should expect oughts that I cannot jointly satisfy.

Of course, this leads to another reservation, about the Relevant Alternatives View: namely, that *all* of the alternatives are *always* relevant. Yes, one can ask questions like the

questions about the holidays above in an idle, hypothetical spirit. But when it comes to *actual, engaged choice*, a *serious, conscientious* deliberator or advisor won't exclude *any* alternatives.

But this isn't so. Invited to review, for example, Procrastinate needs to make a decision now. He needs to determine, of the things that he can decide now, which he should decide. This includes fewer than all the available alternatives. For example, it excludes *actually* writing in the future, which is not a present decision. However, it seems a perfectly serious, engaged, conscientious question for him to ask.

To consider all of the available alternatives, even all of the things that one can decide now, would be overwhelming. Usually, we reduce the options to some more manageable set, assuming that further distinctions don't matter. Thus, to take a second example, the Mountaineer might ask himself, seriously, conscientiously, in an engaged way, "Ought I to step or not?" Or he might ask us this, in a context in which we can give him only a yes or no answer. Now, he is mistaken in assuming that further distinctions don't matter. The further distinction between stepping-while-looking-down and stepping-while-not-looking-down does matter. Still, he was engaged, serious, and conscientious in asking his question. (It is this feature of the case that leads me to present it instead of other "synchronic" cases in the literature.)

[32] It may sound better to conclude, in the third person, "Procrastinate ought to accept." If so, I suspect, it is because it is easier to hear the third-person claim as being not in the mode of advice or deliberation, but instead in the mode of criticism or assessment: as equivalent to, "If Procrastinate were ideal, he would accept."

[33] Of course, this raises a question about how "ways" should be individuated. But I see no reason to doubt that they can be individuated so as to lead to plausible results.

[34] More precisely: *M*-ing is an *excessively costly* means to one's *E*-ing iff there is some *M** such that:

    (i) it is possible for one to *M** without *M*-ing,

    (ii) the probability, conditional on one's *M**-ing without *M*-ing, that one's *M**-ing, or some part of one's *M**-ing, helps to bring about one's *E*-ing nonsuperfluously is at least as high as the probability, conditional on one's *M*-ing, that one's *M*-ing, or some part of one's *M*-ing, helps to bring about *E*-ing nonsuperfluously (which is positive),

    (iii) there is stronger reason to refrain from *M*-ing than from *M**-ing.

[35] Here's another example, due to Aaron Bronfman. Suppose on my way home, I could leave the route to help stop a boulder from destroying my house, then returning to my route and finishing the trip home. However, there are already enough people there to stop the boulder; my additional help won't make it any more likely that my house is saved. Taking the detour is not superfluous. I do have reason to do it: namely, that it will move my body to locations from which I will then be likely to proceed home. But it is presumably excessively costly. Just staying on route is at least as good as means to getting me home and takes less effort.

[36] Broome (ms.) argues against Weak Sufficiency with an example that, I believe, exploits superfluity. An itchy donkey has the end of scratching its itch. If it walks to the scratching-post, which it hates to do, that will help to bring it about that it scratches its itch. However, the donkey can also easily, effortlessly scratch the itch with its hind leg. In that case, Broome suggests, it has no reason to walk to the scratching-post. But this judgment seems so natural, I think, only because we implicitly assume that, because the donkey can so easily scratch its itch with its hind leg, it is almost sure to do so. In that case, its walking to the scratching-post, in addition to its hind-legging it, is superfluous (as well as excessively costly). Alternatively, we

might imagine that the donkey will hind-leg it only if it does not walk. In that case, walking is not superfluous (but still excessively costly). It has just as much (itch-alleviation) reason *to* walk as to hind-leg it (although more reason to refrain from walking than from hind-legging it).

[37] For rather subtle reasons, Joseph Raz's:

> *Facilitative Principle*: When we have an undefeated reason to take an action, we have reason to perform any one (but only one) of the possible (for us) alternative plans that facilitate its performance (Raz 2005a: 6)

also has trouble with superfluity. This is so, despite the fact that his "but only one" restriction seems intended to ban transmission to superfluous means. So understood, the restriction would presumably say (in the idiom of this paper) that if in all possible outcomes, one takes some facilitating plan A, then no reason is transmitted to taking a second facilitating plan B. (Alternatively, the restriction might be understood as a ban not on superfluous means, but instead on excessively costly means. It would then say that if A and B are each facilitating plans, no reason is transmitted to the facilitating plan consisting of the *conjunction* of A and B, at least if that conjunction is more costly than either alone. But this restriction would be misplaced, since, as Raz 2005b: 2–4 rightly stresses, reason transmits to excessively costly means.) The problem is that this restriction bans *more* than merely superfluous means. The condition might be met even though B is not superfluous. B might be a failsafe, in case A doesn't come to fruition. Intuitively, in such a case, it seems that at least some reason to take B is transmitted from the end. So, unless more is said about what a facilitating plan is, the "but only one" restriction seems too strong.

One might reply that by "facilitating plan," Raz means "sufficient means." Taking a second *sufficient* means, when one has already taken one sufficient means, *is* always superfluous.

However, Raz writes that the Facilitative Principle is only "roughly speaking" about sufficient means to an end or "what we may crudely and inaccurately describe as means sufficient for its realization" (2005b: 9).

Alternatively, one might reply that by "facilitating plan," Raz means "*unimprovable means*." Unimprovable means need not be *sufficient*; they may fail to achieve the end. But they fail to achieve the end *only when this couldn't have been helped*: only when taking additional means would not have achieved it either. More rigorously, a means, *M*-ing, to *E* is *unimprovable* iff there is *no* means *M** such that if one *M*'s in every possible outcome, there is some possible outcome which one's *M**-ing helps to bring about one's *E*-ing in a nonsuperfluous way (and the way in which one *M*'s at that outcome is not the way in which one *M**'s at that outcome). This would validate the "but only one" restriction. If one is sure to take an unimprovable means, then *any* further (putative) means are (at best) superfluous at every possible outcome. The problem is that if facilitative plans are identified with unimprovable means, then the Facilitative Principle transmits reason to facilitating plans that *are* superfluous. Suppose that, in every possible outcome, one takes some means, *M**-ing, that is improvable and so not a facilitating plan. *M**-ing makes taking some further means, *M*-ing, which *is* unimprovable and so a facilitating plan, superfluous in every possible outcome. If one additionally *M*'s, one will still be taking only one facilitating plan overall. So the Facilitative Principle allows reason to transmit to additionally *M*-ing, even though it is superfluous in every possible outcome. (For an illustration, suppose that in every possible outcome, one gives the patient Drug *A* (our *M**). Drug *A* will cure the patient if the patient has Disease 1, as is likely. However, administering Drug *A* is not unimprovable, and so not a facilitating plan. If one additionally gives the patient Drug *B*, this will not interfere with the effect of Drug *A* and will cure the patient on the off chance that the patient has Disease 2. If

one additionally gives the patient Drug *C* (our *M*), this will combine with Drug *A* to make Drug

D, which has exactly the same effect as Drug *A*. However, giving Drug *C* is unimprovable, let us

suppose, since it precludes Drug *B* from having any effect. So giving Drug C is a facilitating

plan. According to the Facilitative Principle, under the present interpretation, one has reason to

give Drug C, in addition to A. But this seems wrong, since giving Drug C is superfluous in

every possible outcome.)

[38] This is in contrast to the Weak principles, Bedke's (2009: 678) "Instrumental Principle" ("One

has reason to take the means to what one has ultimate reason to do"), and Raz's Facilitative

Principle. The Strong principles and Schroeder's General Reason Transmit are more

informative. However, while Schroeder says that the strength of the reason depends on "how

well" the means facilitates the end, he leaves it largely open how "how well" is to be understood.

General Transmission might be seen as offering one understanding.

[39] To illustrate, suppose an aimless teenager has not adopted any plan that contains taking the

SAT as a step. Nevertheless, a parent may know that the teenager has reason to take the SAT

exam as a means to a college education (perhaps partly because the parent knows that, in time,

the teenager will begin to take his future seriously and take other steps, such as filling out

application forms). Taking the SAT is neither a necessary means (since some schools permit the

ACT instead), nor a sufficient means (since other partial means, such as submitting the

application are necessary). By contrast, Bedke's Instrumental Principle and Schroeder's General

Reason Transmit apply to such means.

Raz's Facilitative Principle, however, seems not to. Taking the SAT does not fit the

description of a "facilitating plan": "Facilitating steps can come in ordered sequences, each

constituting one way of bringing us to the point where we can take the action we have reason to

take. We could call each a (possible) plan of how to get to a point at which we can take that action" (2005a: 5). (Moreover, if taking the SAT were a facilitating plan, then other partial means, such as submitting the application, would also count as facilitating plans to that end. But then the "but only one" restriction of the Facilitative Principle would imply, incorrectly, that the teenager does not both have reason to take the SAT and have reason to mail the application.) Nor is it covered by Raz's supplemental claim that "we have only conditional reason to take [steps within a plan], the condition being that we have adopted and are pursuing the plan" (2005a: 6). The teenager has not adopted, and is not pursuing, any plan that has taking the SAT as a step. To be sure, Raz observes that reason can transmit to such means, which "keep options open" for not yet adopted plans (2005b: 8). The problem, to be pedantic, is that letter of his formulations doesn't actually allow for this. Furthermore, General Transmission applies to means, such as taking the left bale as a means to taking either the left or the right bale, that cannot be steps even within adopted facilitative plans, since such means do not put the agent in a position to perform the action, but instead are ways of performing the action. Again, the point is pedantic, since Raz (2005a: 13 n. 18) clearly recognizes transmission to such "means," governed by the more basic and general principle that, as I recount in section 7, he uses to argue for Weak Detachment.

[40] This is more or less Bedke's solution. Millsap (ms. a) observes that caution about what we count as necessary or sufficient means will solve many versions of the problem. It is also worth noting that, if we add the Intransitivity Proviso, then we need not worry about the problem of "explosion"—that if there is reason for something, then there is reason for anything—that Millsap (ms. b) notes is a consequence of certain transmission principles, and that similarly appears to worry Raz (2005a: 12 n. 17). However, if we add the Proviso, then many of the

examples used in the text may not be entirely apt, since the ends in question are themselves not "ultimate," although, in particular cases, this may not lead to problems of the kind that motivate the Proviso.

[41] For discussion of silencing, see McDowell (1998a, b, and c); and Price (2008: 12 n. 9, 184 n. 72). Setiya (2005) raises this silencing objection against something like Weak Necessity, whereas Broome (2005) appears to raise it against Weak Sufficiency and Raz's Facilitative Principle. Broome (2005) argues that no reason transmits from the end of alleviating one's hunger to the overwhelmingly (prudentially, if not also morally) objectionable, but not superfluous, means of killing oneself. In Broome (ms.), he claims to make the same point with a different example: the itchy donkey of note 36. But the donkey makes a different point. Its walking is superfluous, but not overwhelmingly objectionable. What may lead Broome to assimilate the two examples is that, in both, the means are excessively costly. But that may be a red herring.

[42] Kieran Setiya suggests another response: namely, that even if there are intuitive cases of silencing, the Silencing Proviso is unnecessary, because such cases are already handled by the Intransitivity Proviso. We resist the conclusion that I have reason to poison the don, for example, only because we think that the *ultimate* end for which I have reason is not changing policy, but instead changing policy *in a morally permissible way*, which poisoning the don (like baking the mega-loaf) does nothing to further. Whether this ingenious suggestion will handle all the intuitive cases of silencing, however, I am not sure (in part because I don't share many of the relevant intuitions).

[43] While sympathetic to Broome (1999), Bratman (2009) avoids this result by arguing that in such exceptional cases, the general reason to be instrumentally rational lapses, because it derives

from reason to be "self-governing," which is not possible in such cases. See Brunero (2010) for criticism.

[44] Against Weak Detachment, Bratman (2009: 422) suggests that reason in favor of one's *F*-ing "involves an implicit comparison of [one's *F*-ing] with its available alternatives," including refraining from *E*-ing. Thus, it might be that refraining from *E*-ing, rather than *F*-ing, is "what there is reason to do." But this reply seems to the point only if "there is reason to *X*" means, in the idiom of this paper, "at least as much reason for *X*-ing as for any alternative to which we are implicitly comparing *X*-ing" (a usage which makes "overridden reason" an oxymoron). That isn't the usage at issue here.

[45] Compare Yalcin (2007) and Kolodny and MacFarlane (2010).

[46] General Production may differ from standard decision theories, however, in at least two respects. Decision theory is typically presented as a theory of *rationality*, and so focuses on epistemic probabilities fixed by *the agent's information at the time of acting*. By contrast, judgments about *reasons* often depend on other probabilities (e.g., such as those fixed by the information of an advisor, or the non-epistemic historical situation). Second, decision theory often focuses on "exhaustive" actions of the sense described in section 3. By contrast, ordinary judgments about reasons rarely, if ever, involve exhaustive actions. Moreover, focusing on exhaustive actions threatens to crowd out the very phenomenon of instrumental transmission itself. For it ignores other actions from which to transmit reason *to*, or to which to transmit reason *from*, exhaustive actions. (See Wedgwood's insightful 2011 for an analogous point about instrumental rationality.)

[47] Compare Anderson (1993: Ch. 2); Scanlon (1999: Ch. 2); and Raz (2012: Ch. 12).

[48] This isn't to deny that one can treat the state of affairs of honoring the thing of value as itself a valuable state of affairs and still get many of the right "results." But it is still a distortion, even if useful for certain modeling purposes.

[49] Insofar as decision theory cares only about *ranking* actions, it does not need to distinguish this sort of case, in which two actions are equally ranked because there is the reason of *equal strength* for either, from a case in which two actions are ranked equally because there is *no* reason for either.