

A Mysterious Case of Missing Value

Earl Conee

University of Rochester

1. The Problem

Our topic is a difficult sort of choice. The choice concerns what to do when different evaluations conflict about what we ought to do. It will be argued that the choice cannot be decided by an objective standard. It will also be argued that this objective indeterminacy has its consolations.

1.1 To highlight the distinctive problem, here is a contrasting example of an objectively resolvable conflict.

E1 You have available two routes to drive home. You know that one route is more scenic and the other route takes less time. This choice poses a conflict between aesthetic value for you and efficiency value for you. Though it is unusual to know this, you also happen to know that nothing of any sort of value except your own interests is relevant to this choice. Thus, you know that it does not make any moral difference which route you take, it makes no legal difference; it makes no difference to others' well being, and so forth. You wonder whether the extra traveling enjoyment is worth the extra time. You ask yourself, "Which route should I take?"

There is an important interpretive question regarding E1: What does your use there of "should" mean? In the view that I take about "should" and "ought" evaluations, they are about optimizing. The view is this:

The Best View (TBV) What "should" or "ought" to be done is the best of the alternatives, by the pertinent standards for being an alternative and for being the best. The pertinent standards are those

that are at stake in the semantically relevant context.¹

We can apply TBV to our story E1. Since only your self-interest is at stake, your question about what route you “should” take is about the values for you. It is about comparing how your well-being is served by your taking the most scenic route, versus how your well-being is served by your taking the quickest route. You ought, for the sake of your self-interest, to take the route that is best for you. So the question can be rephrased, “Which route would be the best one for me?” If it is a tie on this evaluative dimension, then neither is the best and neither is the one you should take.

It might be easy for you to figure out which route benefits you more; it might be difficult or even impossible. The fact about value for you is what answers your question, whatever its epistemic availability to you.

The value for you of the two ways to drive home is the standard that this question makes pertinent in these circumstances. It has an objective answer. The answer is settled by the truth about the

¹Decades ago Aaron Sloman made a proposal similar to TBV, in “Ought’ and ‘Better’”, *Mind* LXXIX no. 315 (July 1970) 385-394. One modest difference is that Sloman takes some uses of “ought” not to be covered by the proposal. He excludes uses for expressing preferences and for proclaiming metaphysical absolutes (*ibid*, Section G). In my view TBV covers all meaningful uses of “ought.” (Though it does not matter here, I take the same idea to apply to “ought to be” sentences, such as predictions: “It ought to be raining by this afternoon.” The meaning accords with TBV. In a typical context the sentence says that rain falling by this afternoon is part of the epistemically best future.) As in any other case, uses of “ought” that voice preferences or make metaphysical proclamations have contexts that set standards for the relevant sort of alternatives and for being the best of them. If their contexts do not do this, then the uses lack determinate meaning.

TBV is a version of what has more recently become a standard semantic view about “ought” sentences. For an overview and partial defense of the standard semantic view see, “The best we can (expect to) get? Challenges to the classic semantics for deontic modals” by Kai von Fintel, web-posted at <http://web.mit.edu/fintel/fintel-2012-apa-ought.pdf>.

greatest value for you. If you know which one would be better for you, then you know which route you should take.

It is important for what follows to emphasize that this sort of value is familiar to us. We understand the notion of being best for a person. There are serious philosophical issues about what is good for a person. Perhaps it is just the person's pleasure, or perhaps factors like improvements in the person's virtue, knowledge, or opportunities for free choices, can also make an alternative better for a person. The question of what is bad for a person is equally non-trivial. But this issue of what determines value for a person is something that we understand well enough to pursue it philosophically. The understanding shows that we have some grip on the value that is at stake. Our grasp is sufficient for us to be able to ponder the issue and to argue about it.

1.2 Now let's consider a more problematic sort of conflict. Our first illustration will be a conflict between self-interest and another person's interest. Here is an example:

E2 Sam is walking through an isolated area of a large city park. He sees a man who is clearly lost. The man seems confused and troubled. Sam is on his way to catch a train. He has no time to spare. If he misses the train, he will miss a golden opportunity to interview for a better job. Sam sees that he ought, for the other man's sake, to stay there to give kind assistance to the man, even though this will result in Sam's missing the train. Sam also sees that he ought, for the sake of his own welfare, to keep going to make it to his interview, helping the other man only by calling call 911 while hurrying onward.

TBV, our best-alternative thesis about "ought" evaluations, makes relevant two bests. Sam's alternative that is best for the other man is for Sam to stay and help him. So by TBV, for the other man's sake Sam ought to stay. Sam's alternative that is best for Sam is for him not to stay in the park, letting him catch the train. So by TBV, for Sam's sake he ought to hurry onward.

When there is a conflict between the choices favored by different evaluative standards, like the current conflict between the interests of different people, making a selection on the basis of the

evaluations requires some sort of ranking. One ranking that we sometimes hope to find is *an authoritative and conclusive* ranking of our alternatives. This sort of resolution would give us an objective answer. It would be an answer that is determined by the relevant facts and circumstances, as measured by the relevant value. In other words, we hope to find an answer that is objective in the same way that the answer to the question of which route home is best for you is objective. We might put this aspiration by saying that we would like to find out which of the “ought” evaluations “overrides” the others. We might equally well say that we are seeking the “ought” that “outweighs” the others, or the one that “ought to be taken, all things considered.” By invoking one or another such expression we seek to formulate a question about a comprehensive and authoritative evaluation that resolves the conflict.

Such phrases can seem to us to express what we seek. But what does any of this language really mean?

TBV tells us that any correct “ought” evaluation singles out an alternative that is the best in some way. So if TBV is right, then for any “overriding ought” that could resolve the conflict, there must be a relevant sort of best. Therefore, in order to understand any of these sentences about “ought” evaluations modified by “overriding” or “outweighing,” we must ask: What sort of value determines the relevant way of being best?

This is the value that is missing. We unable to identify what evaluations it makes, because we do not understand what the value is. In fact, its total obscurity is ample reason to doubt that any such value exists. Consequently, any impression we have that we understand the intended “ought” evaluation is a mirage.² We shall

² In section III of “The Ethics of Belief” (*Philosophy and Phenomenological Research*, Vol. 60, No. 3. (May, 2000), pp. 667-695) Richard Feldman defends a similar thesis. He does so in terms of denying the intelligibility of the notion of “just plain ought.” (See note 6 below for a discussion of that phrase.) The principal novel contributions of the present work are to defend its version of the thesis by bringing TBV to bear on the topic, to try out the likeliest candidates for the meaning of “ought, all things considered” and “overridingly ought” and object to the candidates, and to propose consolations for those who acknowledge the semantic emptiness of such phrases.

also see that even if TBV is not correct, we do not understand what makes some sort of ought “override.” This lack of understanding cannot be proven. But we shall observe that the likeliest accounts of the idea do not succeed. We shall then see some easily made mistakes underlying our inclination to think that we do know what this evaluation is.

2. “Override”

2.1 Here is an initially promising idea about the identity of the evaluative dimension that determines which alternatives are “over-ridingly best.” The idea is that this value is the strength of the practical reasons that are involved. This reading sounds good. After all, the question for someone in an “ought” conflict is a practical question about how to act. The conflicting evaluations give differing sorts of reasons for acting in differing ways. It can seem that what needs comparing is the strength of these differing reasons as practical reasons. Thus, the sought-after comprehensive and authoritative evaluation of the reasons seems to be a job for the weights that practical rationality assigns to the alternatives.

The strength of practical reasons turns out not to be what we seek. Practical reason does set a standard. But that standard turns out to be a competing evaluative dimension. It enters into “ought” conflicts on a par with other evaluations. It cannot authoritatively resolve them. An example will bring this out.

E3 What Sally reflectively values most highly is leading a life in which she creates a significant legacy of accomplishment. On reflection Sally esteems this sort of life a great deal more than alternative sorts of lives. She believes with justification that an artistic career as a devoted sculptor would be by far the best way for her to have such a legacy. Sally sees that given her settled reflective values, this is the choice that she has by far the best practical reason to take. But she also foresees that the choice would be disadvantageous to her well-being. The art would be difficult to appreciate and not widely enjoyed. Her life as a sculptor would be financially precarious and it would not be well regarded by her family or friends. Sally also sees that she would have a career that would be more

comfortable, secure, and approved, if she were to become an accountant who had sculpture as a hobby. She realizes that she would enjoy that life considerably more and would have no severe regrets.

Sally's alternative as a dedicated sculptor is markedly favored by her considered values. All of the credible views of what a person's practical reasons favor made this fact decisive. In any credible view of what our practical reasons are, they endorse a decision that is clearly favored by the agent's stable considered values as they evaluate her options from her perspective.³

We get a conflicting result when we apply the standard of self-interest. To optimize her well-being, Sally ought instead to choose the accounting career with a sculpturing hobby. That way of life would be part of the life that she foresees would work out best for her.

Now we have a conflict of "ought" evaluations between the verdicts of practical reason and self-interest. Suppose that Sally sees that one ought to be done to follow practical reason and the other alternative ought to be done for the sake of self-interest. Sally might seek a resolution by attempting to inquire – "Which of these evaluations is overriding in this case?" As an attempt to pose an "override" question concerning conflicting "ought" evaluations, this is as sensible and apt as such attempts ever are. Yet it cannot be that Sally is asking what she rationally ought to do. Sally already knows the answer to that question. She is trying to ask something about how that answer compares to the verdict of self-interest. Thus, practical reason, so conceived, is not the sought-after authoritative resolving evaluation.

We could insist that what we will count as "practical reason" must, among other things, authoritatively decide any conflict among "oughts." In other words, we could insist that what we are

³ This is an implication of views according to which following practical reason consists in maximizing expected utility. It is also an implication of views according to which satisficing can be practically rational, as long as what is favored is good enough and the relevant "utility" is understood to measure the agent's reflective values.

willing count as “favored by practical reason” must be what we would call “overriding.”

Well, we could insist. The trouble would be that, in order to get what we would be demanding, we would have to be able to understand how any account of practical reason could supply an “overriding” sort of verdict. Yet we cannot understand that. Take any of the views according to which practical reason favors whatever is favored by the agent’s own considered values. What sort of “quality” does that evaluation have that could “override” by being more “important,” “authoritative,” or “ultimate” than the values measured by other standards, such as morality and self-interest? What is it about some sort of practical reason that makes it “win out”? In fact, we do not even understand that question.

The point is not that someone in Sally’s position could sensibly wonder whether the standard of practical reason really does “override” the standard of self-interest. The point is not that we understand some further evaluative question that we have Sally asking, but the proposed answer to it remains open to doubt. The putative further question and answer about this “override” are not open; they are empty. The problem is that neither question nor answer is intelligible. We see an opposition between two standards, but we see too little more. We lack any understanding of the sought-after “override” evaluation. Calling this sort of evaluation “practical reason” does not help us to understand what it is.

TBV sharpens our focus on this problem. TBV prompts us to consider the exact way in which an alternative would have to be *best*, in order to be the one that “overridingly ought” to be taken. Our question remains: What is authoritative about that value, whether it is called “the standard of practical reason” or not? What practical reasons clearly do is to give a kind of rationality to choices of conduct. But what we seek is to understand how they could “outweigh” or “have more importance” or “override” reasons from other evaluative dimensions like the moral and the aesthetic. On reflection we do not see what the further feature could be.

2.2 Here is another approach to making sense of the “override.” A perspective with Kantian roots has it that we ought to perform some acts, not for the sake of any specific end, but simply because we are agents. The evaluation is determined by our sheer agency. That might suggest that the standard it sets is somehow the most fundamental one. It might be thought that this fundamental status

makes the standard decisive in resolving “ought” conflicts. To give this broad idea its due, we should consider two interpretations.

2.2.1 The first reading - ‘Override 2a’ - is that a standard that is set by one’s agency overrides because meeting this agency-provided standard is doing what is required in order to continue to be an agent at all. The standard requires taking an alternative that preserves one’s existence as an agent.

If that is the idea, then the standard does not even apply in most conflict cases. In most conflicts, whichever act is performed the person continues to be an agent. So none of the conflicting alternatives would be favored by this agency-provided standard. By Override 2a, then, nothing overridingly ought to be done. Yet “ought” conflicts where the agent’s agency continues, whatever choice is made, can be as severe and as distressing as any.

Occasionally one has a choice where, if an alternative is taken, then one’s agency ceases. One dies or one otherwise becomes utterly incapable of acting. Even in such circumstances it is obscure how some “ought” that favors doing what is required for continued agency would “take priority” or “override.” In extreme circumstances where the conflicting “ought,” the one that favors one’s death, was the “ought” of morality, taking that alternative and dying would somehow be morally worth doing. When an act is that morally momentous, it is entirely unclear what might make staying alive and preserving one’s agency inevitably “win” this sort of competition and “outweigh” the morally necessity act.

2.2.2 The other interpretive possibility that we should consider is that a Kantian reading of “overridingly ought” is not about remaining an agent. Instead it is about some irreducible special requirement that is supposed to be imposed just by virtue of our being agents. This is “Override 2b.” The alleged requirement would make true claims like this: “Simply as an agent, I ought to do this.”

When the Kantian interpretation is understood in this way, we can see that the requirement is subject to a familiar problem. Any such agency “ought” enters the fray as a competitor rather than as an authoritative resolution. As with aesthetic standards, moral ones, legal ones, prudential ones, and the rest, other “oughts” are sometimes in conflict with it. The conflict gives rise to the comparative questions that we are trying to understand. On the present assumptions, an instance of the questions can be asked as follows:

“Simply as an agent, I ought to do one thing. But for the sake of my artistic efforts as a poet, I ought to do another thing, and for my own sake, I ought to do a third thing.” This situation sets up the aspiration for an “overriding ought” as well as any other situation does. It is quite unclear how the fact that as an agent one ought to do something could be any more comparatively decisive than is the fact that doing something is aesthetically best or the fact that doing something is favored by self-interest.

2.3 What else could the “override” be? David Copp has argued that the purported override relation does not exist. In preparation for so arguing he describes what the relation would have to be.⁴ It might be thought that the description at least provides the meaning that we have been seeking for an “overriding,” or “outweighing,” or “all things considered,” sort of “ought”.

Copp describes the purported overriding standard that he argues not to exist as the “supreme” standard. He explains this as being what he calls the “normatively most important” standard. This is “Override 3.” It might be thought that this characterization identifies the relevant override relation.

These descriptive phrases make no understandable designation. We have several contextually definite concepts of “supremacy” and “normative importance.” In various contests such words modify various familiar ratings – legal supremacy, moral importance, motivational supremacy, and the like. But we lack any notion of “absolute supremacy” or “pure normative import.” Copp’s unmodified uses of “supreme” and “normatively most important” signify nothing more definite than does our term “overriding” itself. We cannot understand the “overriding” evaluation in these terms, in the absence of any information about some specific way in which the overriding evaluation is “supreme,” or “of most normative important.” In the absence of this further information, Copp’s descriptions do not give us the understanding that we seek.

⁴ David Copp, “The Ring of Gyges: Overridingness and the Unity of Reason,” *Social Philosophy and Policy* 14 (1997) pp. 86-106. Copp’s argument against an override relation is criticized in ways that seem to me to be effective, in section 5 of Owen McLeod’s “Just Plain ‘Ought,’” *The Journal of Ethics* 6 (2002) pp. 269-291.

2.4 Now let's try out the idea that it is morality that sets the overarching standard. This idea has something going for it. Many of us have a sense that morality has an especially high status, though it is not clear what this height amounts to. Nevertheless, an inchoate sense of a superiority of the moral might be thought to tell us that the moral standard is the one that resolves the conflicts authoritatively.⁵ This is "Override 4."

The trouble for this idea is not new. Morality is in the competition; it is not the authoritative resolution. We can see that an "ought" conflict can remain in need of resolution, even after the moral alternative has been explicitly acknowledged. For example:

E4 Beatrice tells us, "I happen to know something about the politics of my friend's rival for the job that they both seek. But I have no moral excuse to disclose that political orientation to my friend's potential employer. The disclosure would predictably activate a bias in the potential employer, giving my friend the advantage. I morally ought not to do this and induce the employer to make the selection on an immoral basis. Yet it would be very good for my friend to have that job. My friend's well-being is very important to me. And I am sure that my friend would do the job capably." Beatrice asks, "What should I do?"

Someone in a situation like this has a difficult choice to make. Yet it is not difficult for Beatrice to discern what she morally ought to do. In her posing of the question, she explicitly acknowledges the answer to the question of what he morally ought to do. So she cannot be conceiving of her further inquiry as raising that moral question. Yet this is as good an example as any of a request for a further, "overriding ought".

2.5 Here is a further interpretive idea about what we have been calling the "all things considered" or the "overriding" evaluation. We often make unmodified use of "ought."⁶ The new thought is

⁵ I thank Troy Cross for this suggestion.

⁶ In Owen McLeod's "Just Plain 'Ought,'" *op. cit.*, he argues that there is an "ought" evaluation that is thereby fully expressed – what he calls "Just Plain Ought" (JPO). McLeod contends that JPO

that using the term this way sometimes expresses the overlapping of all applicable evaluative standards and this is the “ought, all things considered” or “overriding ought.” That is, we “ought, all

evaluations resolve conflicts between moral, prudential and the other “qualified ought” evaluations. He holds that the qualified evaluations make contributions to a JPO evaluation that are analogous to the contributions that W. D. Ross asserts that *prima facie* duties make toward all things considered moral duties. This idea would be quite helpful if it worked out. It does not. We do not understand a genuinely unqualified JPO evaluation. One problem for our understanding it is this. If TBV is right, then the “ought” in JPO would have to select an evaluative dimension, in order to single out the best alternative on that dimension. It is quite obscure what the dimension might be. It cannot be “just plain value” because there is no such thing. There is value that contributes to self-interest, value that contributes morally, monetary value, epistemic value, and more. But there is no unmodified sort of value. (See section 2.6 below for discussion of the idea that there is a generic sort of value.) Yet if JPO is not about an unmodified sort of value, then it is entirely unclear what value it invokes. If it is any familiar evaluative dimension, then JPO would be another way to make an “ought” evaluation that participates in the conflict. The evaluation could not become an authoritative resolution by being expressible as “just plain ought”, rather than “morally ought”, or “rationally ought” or whatever it was. If it is not any familiar evaluation, then it is quite unclear what it is.

A second problem for JPO does not depend on TBV. It is very hard to see how a JPO could be decisive in an authoritative way. It could be comprehensive. But all of the particular evaluations are comprehensive. For example, aesthetic “ought” evaluations take into account moral factors. Mostly moral factors are aesthetically neutral. But at times something that makes an act morally better, perhaps a graceful and modest way of accepting a gift, makes it aesthetically better as well. The problem concerns how the JPO evaluation, whatever it was, could subordinate the rest. A response like the following seems frequently to be available, entirely sensible, and unanswerable: “Granted, I just plain ought to do A. But for my own good, it would be better for me not to do A. Why does the JPO ‘win’ this comparison, or any other? In fact, what does ‘victory’ mean here?”

things considered” or “overridingly ought” to do something when doing it is our best alternative according to every evaluation that rates our alternatives, moral, prudential, legal, and the rest.⁷ This is “Override 5.”

One asset of Override 5 is that it gives us thoroughly understandable statements. They just state that all “ought” evaluations coincide in favoring one and the same alternative. So we understand them on the Override 5 interpretation. Another strength of Override 5 is that on occasions when all of the standards that we consider do coincide on one alternative, it is particularly easy for us to affirm it as the alternative that “ought” to be taken. Suppose that I realize that it is a fine time for me to do some housecleaning, but I am procrastinating. You might tell me, “Given how you like a clean house, it is in your own best interest to do the house cleaning. Additionally, you promised your wife you’d do it and the house will look its best. You have no competing moral tasks to accomplish. There’s just no question about it, you ought to do it.” When we use “ought” like that, we need not have any particular standard in mind. Yet clearly we do understand what we are saying. So maybe we are then applying this coinciding idea.

It is not finally credible, however, that the coinciding interpretation is the meaning of “ought, all things considered” or “overridingly ought” as we are trying to understand them. Override 5 has an unbelievable implication concerning the sorts of conflicts that we have been discussing. Whenever the evaluations conflict, the standards do not all coincide. So in any case of conflicting evaluative standards, the coinciding interpretation implies that none of the alternatives “ought” to be taken. In any conflict the answer to the question, “What I ought I to do?”—understood by Override 5 as the question of what all standards favor—would always have the same answer, namely, “Nothing.” And it would be easy for us to see that this is the correct answer. All that we would have to do is to note the conflict and understand that what we mean requires unanimity. The unanimity would be rare. Any person’s interests constitute an evaluative dimension. So whenever an alternative does not do best by any one person, it would not turn out that any alternative “ought, all things considered” or “overridingly ought” to be taken. But it is not

⁷ I thank Toby Handfield for this suggestion.

believable either that this negative answer to the intended question holds in all conflicts or that we can find out the negative answer just by noting the conflict.

The Override 5 meaning is a good idea about other phrases, such as “unambiguously ought” and “incontestably ought.” Those phrases suggest that all evaluations agree. In contrast, what we seek from an “ought, all things considered” or an “overridingly ought” evaluation is, among other things, one that selects alternatives in at least some conflicts. The general meaning of the term “override” supports this aspiration. Generally, the overriding consideration takes a side—it favors some authoritative verdict.⁸

2.6 Another interpretive approach combines all evaluations. Instead of requiring them to coincide, the new interpretation relies on a “generic value” that they are to share. Alternatives are to have some quantity of this generic value. It is to be the sum of all of the contributions that they make to each specific sort of value. Applying TBV, the alternative that “overridingly ought” or “ought, all things considered” to be taken, is the one (if any) that is relevantly the best: it has the most generic value.

Here is the thought in somewhat more detail. Value is to be a generic magnitude. Its specifications are to include all of the evaluative types that can enter into “ought” conflicts. Thus, some generic value would be moral value, some would be aesthetic value, some would be value from self-interest, some would be epistemic value, and so on. The generic value of an alternative is to be the sum of the quantities of all of its specific sorts of value. The alternative that someone “ought, all things considered” to take is the one that has most generic value. This is “Override 6.”

Here is an illustration of how Override 6 is supposed to work. An alternative that has great moral value by being heroic would thereby have great generic value. If its alternatives have less generic value from morality and from all other sorts of value, then on the Override 6 reading it is the alternative that “overridingly ought” to be taken. However, the heroic alternative could have even greater negative generic value by being terribly costly to the agent’s self-interest, while some other alternative is generically neutral. If so, then the heroic choice is not generically most valuable and on the

⁸ See, for example, the definition of “override” in Dictionary.com: <http://www.dictionary.com/browse/override>

Override 6 reading it is not the one that “overridingly ought” to be taken.

Override 6 does not succeed as an interpretation. It depends on the existence of a thoroughly obscure sort of value. To illustrate the obscurity, epistemic contributions would have to be comparable in “generic value” to aesthetic contributions. But we have no idea of what might make it the case that, say, the aesthetic good of knitting a handsome scarf is of more “generic value,” or less, or an equal amount, to the epistemic good of, say, deriving of a moderately difficult theorem. Furthermore, in order to cover all potential “ought” conflicts, generic value must subsume many other disparate values, such as financial value, communal value, and entertainment value. It is not credible that all of the mentioned sorts of value are commensurate types of one value. We have no notion of such value. It is just as opaque to us as “generic physical magnitude” would be, if this were alleged to be a quantity that includes as commensurate types all specific physical magnitudes such as distance and mass. Yet Override 6 needs some such notion of value. So Override 6 does not give us an intelligible reading of “overridingly ought.”

2.7 One more interpretive idea deserves consideration. We would like to employ an override relation that we understand, in order to resolve the conflicts. We do not seem to have an existing candidate that pans out. Perhaps we can create by stipulation the meaning of “override” that we seek. Morality matters a lot to most of us. Exploiting this fact, we can stipulate that what we “overridingly ought” to do is always what we morally ought to do.

This stipulation does not give us what we seek. A stipulated meaning only changes our expressive resources. It does not create a new evaluation. Yet for “overridingly ought” to do what it is supposed to do, we need a new one. We have seen that morality and the other familiar evaluations do not work. Putting the point another way, we have not been looking merely for some way to use “overridingly ought” that we understand and that makes a definite selection in a conflict. We seek a meaning of “overridingly ought” that does these things and more. We seek a sort of override that assesses alternatives authoritatively across the various evaluative dimensions. Yet we have no grip on an authoritative comprehensive evaluation. So we do not have available for any purpose, stipulative definition or otherwise, a meaning of the sort that we seek.

3. Taking Stock

We have not found anything that we can recognize as an authoritative standard to resolve “ought” conflicts. The reasons from different evaluations are irredeemably disparate. The trouble is not that we have no way to compare them. The trouble is that we have too many ways. We have one way for each type of evaluation of our alternatives. For any case, each aesthetically substantial factor has some moral significance – morally positive, negative, or neutral; each morally substantial factor has some level of aesthetic significance – aesthetically positive, negative, or neutral; and so on. Applying one or another of these familiar ratings is all that we can do in order to make comparative evaluations. We wish that there were some comprehensive ultimate standard. But finally we do not understand what it would be.

4. Explanations of an Illusion

Many of us have thought that we understood phrases like “ought, all things considered” and “overridingly ought.” The present view is that we do not understand them. That is mysterious. Our impression of understanding is not some shallow error. It is not like mistaking for a meaningful English word a vaguely familiar bit of gibberish like “brillig” or “slithy” from Lewis Carroll’s nonsense poem, “Jabberwocky”.⁹ The view is that phrases like “ought, all things considered” and “overridingly ought” as wholes do not combine intelligibly. Though the semantic obstacle is considerably less blatant, it is otherwise quite similar to trying to understand the phrases, “ought, no things considered” and “indeterminately ought.” The modifiers thwart the applicability of the verb, rather than restricting it intelligibly. Yet some of us have an impression of understanding. A credible defense of the view requires

⁹ Noteably, the word “chortle” from “Jabberwocky” did become a meaningful English word. It seems to have acquired a meaning by sounding enough like the already meaningful words “chuckle” and “snort” to take on something like their conjoined meaning. (See the Online Etymology Dictionary:

<http://www.etymonline.com/index.php?term=chortle>).

Our problem phrases with “ought” are not like that. No combined meaning is in the offing.

explaining the illusion behind an apparent grasp of “ought, all things considered” and “overridingly ought.” Two explanations are available.

4.1 First, we can mistakenly think we have a concept on the basis of attempting to single out the concept using properties that we intend it to have. Suppose that someone finds a conflict between her self-interested alternative and her moral one. She is moved to ask what she “ought, all things considered” to do. The person may be conceiving of the evaluation that she seeks as the one that accomplishes two things. First, it “weighs” the moral and prudential considerations and every other specific sort of evaluative consideration. That is, it measures their strength as reasons along some dimension that is inclusive with respect to these other sorts of evaluations. It does this without any indeterminacy that prevents selecting a best alternative.

The second feature of what the “ought, all things considered” is intended to do is to provide an authoritative factually determined resolution of the conflict. That is, the sought-after evaluative dimension is further conceived as one that solves the problem of choice that the conflict poses with authority and objectively. Putting these things together, the overriding ‘ought’ is intended to be the one that provides an objective authoritative ranking of the alternatives that resolves the conflict.

Thinking about these uses of “all things considered, ought” or “overridingly ought” in some such way can seem to identify some particular evaluative standard. But that impression is illusory. Further reflection reveals that this much is not enough to give the question a meaning. The phrases describe some understandable evaluation only if it singles out some understandable sort of “weight” and “authority.” What sort is it? It is not any familiar specific weight like moral weight. Moral weight is only one among the factors that are supposed to be given “their due.” We are now seeking what this other ‘weight’ is, the one that is supposed to determine the due contributions to an authoritative conclusive resolution. For the same reason, it is not the weight of practical reasons, or reasons of any other familiar kind. The fact is that we cannot further identify the ‘weight.’ And what could be its ‘authority’? It is not authoritative by being morally decisive. Again, it is not just the moral evaluation stated over again. We have seen that morality does not arbitrate as intended- rather, only some “due” contribution from moral considerations enters in. For the

same reason, it is not ‘authoritative’ by making any other familiar evaluation. Thus, the abstract constraints of accomplishing some “comprehensive weighing” and “conclusive authoritative resolution” do not yield any understandable evaluation.

4.2 TBV helps to identify a second source of illusion that can obscure the lack of content to “all things considered ought,” “overridingly ought,” and the like. TBV tells us that for any meaningful use of “ought,” something contextual provides a standard for being relevantly best. The provision can be explicit, as when we say, “morally ought.” Very often the standard is not explicit. If TBV is correct, then a tacit contextual designation of a standard plays a frequent role in how we understand “ought,” on the frequent occasions when we do get it. Since our comprehension of ‘ought’ claims often relies on this, we are disposed to find contextual cues that fix the standard. So, when a use of “ought” is not explicitly modified, we are inclined to seek some familiar standard that the context somehow calls for. We usually succeed. For instance, when someone asks, “How ought I to open this package?” it is usually clear that the question concerns being most efficient, or it is clear that it concerns best protecting the contents. For another example, when someone asks, “How ought I to identify the heirs in my will?” it is usually clear that the question is one of maximal legal efficacy. In general, usually the relevant evaluative dimension is readily specified if we consider the matter.

This interpretive inclination is not disengaged when “all things considered” or “overriding” is added to “ought.” We expect context to do its usual work. This is true despite the fact that, as we have seen, we could not actually find any standard that it selected if we tried to do so. We might even think we could find it if we sought it carefully enough. The serious interpersonal matters at stake in some conflicts tend to bring morality to mind as a prime candidate. On other occasions, our concern for an agent makes a prudential standard the one that looms large in our thinking. In still other cases, the desire for a clear-cut conclusion to a conflict inclines us to prefer an evaluation that decisively favors some one alternative. In these ways it can appear to us that we can specify a definite evaluative measure for the “ought, all things considered,” “overriding ought,” or the like.

Further thought undercuts this appearance. We can explicitly acknowledge the implications of the particular standard that we are inclined to favor, and then try again to ask the problematic sort of

question. Suppose that we are deeply concerned about the welfare of some agent. We are inclined to resolve a conflict by recommending whatever is most prudent for her to do. We can intelligibly affirm that the alternative that we recommend definitely outweighs the others on the prudential grounds that matter most to us. But we still cannot fathom the “overriding ought” question: “Is this priority that we ourselves have also objectively correct? Does it identify the alternative that she ought, all things considered, to take? Is it what she overridingly ought to do?” We do not understand what any such question asks.

5. A Consolation

This lack of comprehension is not appealing. We sometimes find that we cannot take every alternative that in some way or other we ought to take. On these occasions, it can be discouraging to think that we do not even understand how the conflict could be objectively resolved. We still have to make a choice.

It is some consolation that the emptiness of ‘overriding ought’ does not deprive those who confront such conflicts with a rational basis for deciding. Our reflective concerns give good reasons to make choices. Usually we care on reflection more about some of the conflicting evaluations than others. This concern can be capable of surviving scrupulous critical scrutiny. Choosing in accordance with this sort of priority is a way to make a rational choice.

This rational assessment can be expressed with “ought.” We can say that choosing a certain alternative is what someone ‘ought’ to do, in that it best meets the person’s reflectively stable priorities. As with all of the uses of “ought” that we understand, this one applies a particular standard. The problematic comparative question is not thereby given content, much less answered. To see this, we can note that one who has ascertained the reflectively preferred choice can raise the problematic question with as much apparent sense as usual. Consider:

E5 Gauguin Prime is a fictional artist who we stipulate to be reflective in a certain way about a major choice. In preparation for the choice, Gauguin Prime critically scrutinizes his priorities. Upon due reflection he finds that his settled values favor devoting himself to artistic endeavors, even at the

expense of fulfilling family responsibilities. Gauguin Prime is not comfortable with taking that alternative. He observes, “A thorough devotion to my painting is favored by my considered values. I see that this is not what I morally ought to do and that it is not best for the people closest to me.” He asks, “What ought I to do, all things considered?”

This attempt to pose a question that would resolve the conflict is as understandable as any typical “ought, all things considered” inquiry. Yet Gauguin Prime would not be seeking to rate his alternatives by his underlying reflective priorities. That rating is already established for him as part of his basis for attempting to ask a further question. The question seeks a further sort of standard. Reflective values that can rationalize choices do not furnish that standard. Nothing does. It is an illusory question that could not have a correct answer.

6. Objection

An objection should be considered. Its discussion should add clarity to the view advocated here.

6.1 It might be thought that there must be an override relation that can resolve “ought” conflicts, because we know its application in some cases. For instance, we know how it applies in the following case.

E6 Etiquette favors Maria’s making a polite and graceful departure from a dinner party. But Maria has just recalled a late evening appointment to mentor a troubled teenager. She will clearly miss this appointment unless she catches the next subway. She can just barely make it. So although Maria ought, as a matter of etiquette, to depart politely and gracefully, she ought, as a matter of morality, to rush away.

The objection alleges that in cases like this it is quite clear that Maria overridingly ought to do the moral thing. We know that she ought, all things considered, to do that. We know that the moral “ought” overrides here. The objection concludes that since we have a suf-

ficient grasp of an overriding relation to discern its application in such cases, the relation exists and we have some awareness of it.

6.2 In a case like this we do succeed in making sense of an override relation between the alternatives. But in such cases different evaluative dimensions do not compete. In our example of Art, the two evaluations are fundamentally moral. Evaluations of etiquette often assess minor matters of morality. The “weights” of reasons that we discern in examples like E6 compare a morally significant factor with a morally minor factor. The relevant ‘override’ is just a matter of something clearly having greater moral significance. This gives no evidence for the existence of a comprehensive override relation that compares different evaluative dimensions.

This reply might seem to take advantage of the particular sort of conflict in our example. It does not. Only in cases like this one is an apparent conflict between evaluations resolvable by a well-understood sort of override. In these cases the appearance of a resolution is not deceptive. The deception rather consists in an appearance of separate evaluative standards. On the reading of the case that allows a clear resolution, only one familiar standard is actually at work.

Some examples in which justice and morality appear to conflict are similar. The pursuit of justice is also part of the moral endeavor. So the moral standard governs the conflict and there can be a moral resolution.

When the conflicting “ought” evaluations are clearly on different evaluative dimensions, moral, vs. aesthetic, vs. self-interest, vs. the interest of another, and so forth, we never see what the agent “overridingly ought” to do. Applying TBV, we would have to discover the “overridingly best” thing for the agent to do. We cannot make sense of that. Suppose that self-interest slightly favors one choice while morality heavily favors another. We can know that there is this contrast in magnitudes—a small favoring on the one side and a large one on the other. This contrast can lead us to recommend the moral alternative. The cost to the agent is small and we take morality seriously. But this recommendation is not based on our apprehending any authoritative override relation and employing it. Without overlooking anything relevant we could have recommended the self-interested choice instead, with equal sense and equal conviction, if we had cared dearly for the agent while caring less about morality. Clearly we have preferences for some evaluations in some cases. What we never have is a grasp of

a comprehensive and authoritative sort of quality by the measure of which our favored choice is best.

7. A Concluding Consolation

The alternatives in a conflict are not rated by any resolving objective standard. This absence can be found liberating. Our selection is left open by the evaluative facts. The choice is not beyond criticism. It can be variously good or bad. For some examples, it can be good by being generous, wise, healthy, courageous, tactful, considerate, or sensible; it can be bad by being cruel, selfish, malicious, imprudent, tasteless, biased, or frivolous. Nonetheless, in making the choice we cannot be rightly accused of missing some objective fact about what overrides. We are free to choose a value to favor, with no danger that the choice makes an objective evaluative mistake.¹⁰

¹⁰I am grateful for the cordial and helpful discussion of a draft of this work at the SUNY Brockport Center for Philosophic Exchange.