



BRILL



brill.com/jmp

Some Critical Comments on Zimmerman's *Ignorance and Moral Obligation*

Krister Bykvist

Institute for Futures Studies, Stockholm, SE – 111 36 Stockholm

krister.bykvist@philosophy.su.se

Abstract

In his recent book, Michael Zimmerman continues to defend the Prospective View, according to which moral obligation depends on evidence about both empirical and evaluative factors. In my commentary, I shall first focus on Zimmerman's framework in which different moral theories are defined and distinguished. I argue that Zimmerman fails to formulate a clear and coherent distinction between The Prospective View and the Objective View, which he rejects. Then I turn to the so-called constraint #2, a crucial premise in Zimmerman's master argument against the Objective View. Here I argue that it should be given up so that we can give the right verdict in cases of fundamental moral uncertainty. More specifically, I shall argue that a morally conscientious agent can *rationally* choose the option that is guaranteed to be morally wrong in a Jackson-case of fundamental moral uncertainty. Finally, I shall argue that the Prospective View, in its most recent guise – according to which moral obligation depends on empirical and evaluative evidence the agent has *actually availed* herself of – has very troubling substantive implications that go against all traditional moral theories, as well as an earlier version of Zimmerman's Prospective View.

Keywords

uncertainty – obligation – prospective view – moral conscientiousness

1 Introduction

It is very common in normative ethics to ignore any epistemic shortcomings of the agent. It is often taken for granted that virtue ethics tells you to do

what the virtuous person would do, that Kantian ethics tells you to act on a maxim that could consistently be willed to become a universal law, and that utilitarianism tells you to do what would maximize general happiness. But the obvious problem with these theories as they are stated here is that an ordinary agent will often lack knowledge about whether an action would be done by a virtuous person, whether the maxim of her action can be universalized, or whether her action would maximize general happiness. Is it right to ignore the epistemic situation of the agent when formulating moral theories?

Michael Zimmerman argues that we should reformulate moral theories so that they take into account the epistemic situation of the agent. More specifically, he defends a view he calls the Prospective View, according to which our moral obligations should depend on our *evidence* about both empirical and evaluative factors.

Zimmerman has written about this question before, but his recent book *Ignorance and Obligation* (Oxford: Oxford University Press), goes far beyond the previous publications. He responds to virtually every published objection to his view and shows how difficult it is to defend the traditional normative theories that ignore the epistemic situation of the agent. He also develops his own theory much further and shows how the new version can deal satisfactorily with many of the objections that were raised against the old version.

This book also contains a very interesting discussion about the action-guidingness of moral theories. Zimmerman shows that our common understanding of what it means for a moral theory to be action-guiding is mistaken and confused in many respects, and he does this just by invoking some very simple and plausible-sounding distinctions.

The final chapter relates his general findings to the discussion of moral rights, and argues that moral rights depend on the epistemic features of agents. In defending this idea, he draws a crucial, but often neglected, distinction between rights and desert.

The book has all the Zimmermanian trademarks. It is beautifully written, lucidly organized, and elegantly and honestly argued. As usual in philosophy, one shows one's appreciation of a colleague's work by critically engaging with his or her theories and argument, and this is what I will do here. In my short commentary, I will focus on (i) Zimmerman's framework in which different moral theories are defined and distinguished, (ii) the so-called constraint #2, a crucial premise in Zimmerman's master argument against the Objective View, and (iii) some substantive objections to the Prospective View.

2 Framework Issues: An Ill-defined Tristinction?

According to Zimmerman, the main contenders in normative ethics belong to one of three general moral views: the Objective, the Subjective, and the Prospective View. The official definition of the Objective View is this:

One ought to perform an act if and only if it is in fact the best option one has. (p. 2)

'Ought' expresses overall moral obligation. By 'best', Zimmerman means 'best in terms of what matters morally, insofar as our moral obligations are concerned' (p. 2). He adds the clarification that this notion is meant to capture what 'is pertinent to determining what is obligatory as opposed to what is supererogatory' (p. 5). Later in the book, he calls this value 'obligation-relevant value' (p. 114).

The Objective view is contrasted with the Subjective View, which is defined in this way:

One ought to perform an act if and only if one believes that it is the best option one has (p. 7)

and, his own favourite, the Prospective View, which in its official refined form states that:

One ought to perform an act if and only if it is prospectively best. (p. 32)

Zimmerman does not provide a clear definition of prospectively best, however. He says that the prospectively best option constitutes 'a best bet,' but this can mean different things in different situations. For example, in a situation where the agent knows the actual values of the possible outcomes of the available actions, but does not know the actual outcome of each action, only its probability, the action which counts as the best bet regarding the actual values at stake is the one that maximizes the expected value (p. 35). In a situation where the actual values are known but no probabilities can be assigned to the outcomes, Zimmerman tentatively suggests that maximizing the minimum value of possible outcomes will be the best bet (p. 68). In cases where the actual values are not known, the best bet is to maximize the expectation of *projected* value (p. 36). To calculate the projected value of an option, in a case where you have empirical certainty, you need to take each moral hypothesis

about the moral values of actions and multiply its probability by the value it assigns to the option, and sum up all these products. What makes this a unified account is that the best bet is always the option which would be most reasonable to choose in light of the agent's evidence about both empirical and evaluative matters, *were he morally conscientious* (p. 73). Given these clarifications, we can formulate the Objective, Subjective, and Prospective View, in these more explicit ways:

The Objective View (explicit version)

One ought to perform an act if and only if it is in fact the action that is best in terms of what matters morally, insofar as our moral obligations are concerned.

The Subjective View (explicit version)

One ought to perform an act if and only if one believes that it is best in terms of what matters morally, insofar as our moral obligations are concerned.

The Prospective View (explicit version)

One ought to perform an act if and only if it is the action which would be most reasonable to choose, in light of one's actual evidence about both empirical and evaluative matters, if one were morally conscientious.¹

Presenting the views this way reveals something surprising: the Objective View now looks extremely plausible, for who would argue that we ought to perform an action that is *not* best in terms of what matters morally, if we are talking about what matters morally in the way pertinent to *determining our moral obligation*, that is, what we ought to do? It seems very plausible to say that if there is another action that is better in terms of what matters morally in the way that is pertinent to determining what we ought to do, then that is what we ought to choose. Indeed, to deny this seems almost incoherent. (I say 'seems almost,' since perhaps it is possible that something is obligatory even if actions altogether lack value or their values can be cyclical.)

Furthermore, the Prospective View now looks like a good candidate for an Objective View, contrary to Zimmerman's intentions, for the Prospective View can be seen as an Objective View that assumes that what matters morally, insofar as our moral obligations are concerned, is what would be most

1 I am adding the qualifier 'actual' here and in the following to make clear that what counts is the *actual* agent's evidence, not the possibly different evidence he would have, if he were morally conscientious.

reasonable to choose in light of one's empirical and evaluative evidence, were one morally conscientious. More exactly:

An option *x* is better than an option *y* in the way that morally matters (in the way pertinent to determining what one morally ought to do) iff it would be more reasonable to choose *x* over *y*, given one's actual empirical and evaluative evidence, if one were morally conscientious.

In fact, the Prospective View looks like a *virtue theory* in disguise, since when properly spelt out it tells us to perform the action that would be most reasonable to choose, in light of one's actual evidence, if one had a certain *virtue*: moral conscientiousness. But this means that the distinction between the Objective View and the Prospective View collapses.

There is also something strange about the characterization of the Prospective View itself. Since this view is supposed to be applied even in cases in which the agent lacks knowledge of obligation-relevant value, we get the strange result that considerations that in fact *lack* obligation-relevant value, but are justifiably believed by the agent to have such value, can determine what the agent ought to do. In other words, considerations that in fact have *no* value in terms of what morally matters in the way that is pertinent to determining what we ought to do can nevertheless determine what we ought to do. That seems almost incoherent. (I say 'seems almost' since perhaps it is possible that a factor determines what we ought to do even though it has no value of any kind.)² I think this is what I should have said instead of just complaining as I did before that, according to the Prospective View, the true axiology 'cuts no ice,' an objection Zimmerman discusses in his book, but finds innocuous³ (pp. 61–63).

These problems could be avoided if the kind of value invoked in the definitions above was not so closely tied to what we ought to do. We could, for instance, invoke *final or intrinsic value of outcomes* and define the views as follows:

The Objective View (intrinsic value version)

One ought to perform an act if and only if it has the intrinsically best consequences.

2 The problems stated here would be equally damaging for an account that invoked 'reason' instead of 'obligation-relevant value.' For example, to say that we morally ought to do something we have no moral reason to do seems as odd as to say that we morally ought to do something that lacks obligation-relevant value.

3 Krister Bykvist, 'Objective versus Subjective Moral Oughts,' *Uppsala Philosophical Studies* 57, (2009), p. 50.

The Subjective View (intrinsic value version)

One ought to perform an act if and only if one believes it has the intrinsically best consequences.

The Prospective View (intrinsic value version)

One ought to perform an act if and only if it is the action which would be most reasonable to choose, in light of one's actual evidence about the probabilities of the possible outcomes of one's options and about what has intrinsic value, if one were morally conscientious and only cared about what has intrinsic value.

The obvious problem with these reformulations is that they become too *consequentialistic* in spirit. After all, the Objective View is supposed to capture more than act-consequentialism, and the Prospective View is supposed to be something a stern non-consequentialist can accept.

Another way to avoid the problems would be to invoke *different senses* of ought and start with a definition of objective ought and then define the Subjective View and the Prospective View in terms of the objective ought.

The Subjective View (subjective ought version)

One subjectively ought to perform an act if and only if one believes that it is best in terms of what matters morally in the way pertinent to determining what one objectively ought to do.

The Prospective View (prospective ought version)

One prospectively ought to perform an act if and only if it is one's best bet regarding what one has evidence has value in terms of what matters morally in the way pertinent to determining what one objectively ought to do.

The obvious problems with these reformulations are that it is not at all clear how we should define the Objective View. To say that the objective ought is determined by objective features is not to say much, since, as Zimmerman himself points out, 'objective' is such a vague and ambiguous term (p. 8). If we want to say, with Zimmerman, that traditional moral views, such as Kantianism and Virtue theory, are forms of the Objective View, we need to invoke a notion of objectivity that is flexible enough to accommodate subjective features of the agent, such as intentions and character traits.

It is also clear that Zimmerman would not be happy with these reformulations, since he does not want to commit to the idea that 'ought' has these different senses. Indeed, he seems inclined to reject the idea that there are these different senses of ought (p. 22).

Zimmerman himself does provide a different characterization of the tris-tinction between the three views later in the book (pp. 114–115). This character-ization invokes the notion of ‘suiting one’s situation’:

The Objective View (reformulated)

One ought to perform an act if and only if that act *suits one’s situation as a whole* more than any other option one has.

The Subjective View (reformulated)

One ought to perform an act if and only if one believes that one’s act *suits one’s situation as a whole* more than any other option one has.⁴

The Prospective View (reformulated)

One ought to perform an act if and only if that act provides one with a prospect of doing what *suits one’s situation as a whole* better than any other option one has.

I do not think these reformulations help, at least not as Zimmerman understands them, since they are not supposed to provide alternative char-acterizations of the three views; they are only introduced to *extend* the ear-lier characterizations to prima facie moral obligations. Roughly, a prima facie moral obligation is an act that (actually suits/is believed to suit/has the best prospect of suiting) *some* aspect but not all aspects of the agent’s situation. Indeed, Zimmerman explicitly defines suitability in terms of value for he says that ‘[t]o say that an act is suitable in this way is simply to say that it serves some actual value at stake’ (p. 114), and later, on the same page, he clarifies that this value is ‘obligation-relevant actual value.’

Could Zimmerman avoid the definitional problems by forgetting about obligation-relevant value and instead simply treating the notion of suiting one’s situation as primitive? That is doubtful, since it is not clear what should count as part of the situation. For example, isn’t the agent’s empirical and evaluative evidence also part of the situation? After all, it is natural to say that some actions suit the agent’s evidence better than others and thus have more going for them. But if we do this, then, again, the Prospective View becomes a version of the Objective View, and the crucial distinction collapses. To say that the situation comprises only objective features will fall prey to the previously noted problem of the vagueness and ambiguity of ‘objective’ and the problem

4 Zimmerman does not explicitly state the reformulation of the Subjective view, but it is clear from the text that this is the formulation he has in mind.

of including the traditional moral theories, many of which give great weight to subjective features, in the Objective camp.

There is one very natural way to draw the distinction between the Objective, Subjective, and Prospective views, which is very clear, but it is doubtful that it will satisfy Zimmerman's demands. The idea is to start with a moral view that states that a certain feature F is ought-making:

The Objective View (F-relative)

One ought to perform an act if and only if the act has feature F.

One then defines the *F-relative* Subjective View in this way:

The Subjective View (F-relative)

One ought to perform an act if and only if one believes that the act has feature F.

The F-relative prospective view is defined in this way:

The Prospective View (F-relative)

One ought to perform an act if and only if it maximizes expected F-ness.

Or in this more general, F-relative virtue-ethical way:

One ought to perform an act if and only if it would be reasonable to choose, given one's actual empirical evidence and evidence about what has F, if one were morally conscientious and only cared about F-ness.

The intrinsic value versions of the Objective, Subjective, and Prospective views can be seen as F-relative views, where F is identified with the intrinsic value of outcomes. The obvious limitation of these definitions is that the Prospective View becomes too F-focused. Remember that the Prospective View, according to Zimmerman, is supposed to apply even in cases of evaluative uncertainty, which includes cases where the agent is not sure whether F is morally important (no matter what F is). The F-relative definition of the Prospective View will thus not be acceptable to Zimmerman.

The conclusion is that, as it stands, Zimmerman's distinction seems ill-defined, and it is unclear how to best resolve its definitional problems, at least if we want to satisfy his own demands on these definitions. This does not make it impossible to critically engage with his views, however. For, no matter how the definitional problems are resolved, it is at least clear which

verdicts Zimmerman thinks the Prospective View implies for certain particular cases. It is also clear that these verdicts are not always the same as the ones given by traditional moral theories, such as utilitarianism, consequentialism, Kantianism, and Aristotelianism. Another crucial difference between these views and Zimmerman's is that his view takes into account *evaluative uncertainty*. Moral obligation depends partly on the agent's evidence about what has value (in some yet to be specified sense).

3 Constraint #2, the Crucial Premise

The crucial premise in Zimmerman's master argument against the Objective View is this:

Constraint #2

It is necessarily the case that, if one acts morally conscientiously, then one does not deliberately do something that one (fully) believes to be morally wrong.

This constraint sounds plausible, but it is easy to conflate it with the following constraint, which no one could reasonably reject:

The uncontroversial constraint

It is necessarily the case that, if one acts morally conscientiously, then one does not deliberately do something that one (fully) believes to be morally wrong *rather than something one (fully) believes to be morally right*.

It is easy to conflate the two constraints, since, *typically*, if an agent believes an action to be morally wrong, he also believes there to be an alternative morally right action. The so-called Jackson-case would be an exception, assuming a moral view that does not take into account empirical uncertainty but only actual outcomes when determining moral obligation. In this case, the agent knows that the guaranteed second-best action is wrong, but she does not have a full belief about which action is right because she does not know which alternative action is right in terms of being best for the patient. She knows that treatment B will provide a partial cure, and that of treatments A and C, one will kill the patient and the other will fully cure her, but she does not know which is the killer and which is the savior. Table 1 illustrates this case (where the numbers represent the patient's well-being on an interval scale):

TABLE 1

	S1 (p = 0.5)	S2 (p = 0.5)
A	10	-100
B	6	6
C	-100	10

Is Zimmerman's more restrictive constraint true? Myself and Peter Graham have argued that in the Jackson-case we should say that one would act morally conscientiously by choosing the medication that is guaranteed to give only a partial cure – the guaranteed second-best option in terms of the actual well-being of the patient.⁵ The idea is that morally conscientious agents are not just interested in doing what they believe is morally right and avoid doing what they believe to be morally wrong; they also care about the *severity or seriousness* of wrong-doings.⁶ More specifically, they are not willing to risk a serious moral wrong for the sake of the possibility of doing right, if both the risk and the severity are sufficiently great. In the Jackson-case, the morally conscientious agent does not know which act will in fact fully cure the agent and thus be best for the patient. He knows that either the A or the C option will fully cure the patient, but he also knows that one of them will kill the patient. Choosing either A or C would thus be risking a major wrong, and a morally conscientious agent is not willing to do this. So he chooses option B, which he knows is a minor wrong, in order to avoid risking the major wrong. My own view is that choosing B is what is *rational* to do, given the morally conscientious agent's beliefs about

5 Krister Bykvist, 'Objective versus Subjective Moral Oughts,' pp. 53–54, and Peter Graham, 'In Defence of Objectivism about Moral Obligations,' *Ethics* 121 (2010), pp. 98 ff. See also, Krister Bykvist, 'How to do wrong knowingly and get away with it,' *Uppsala Philosophical Studies* 58, (2011), and Krister Bykvist, 'Evaluative uncertainty and consequentialist environmental ethics,' in L. Kahn and A. Hillel (eds.), *Environmental ethics and consequentialism* (London: Routledge, 2014).

6 Note that this is not to say that rightness comes in degrees, which would be an odd claim. Nor is it to say that wrongness *itself* comes in degrees. It is just to say that the severity or seriousness of wrongness comes in degrees, which seems very plausible. Note that this does not rule out Zimmerman's account of wrongness, according to which what is wrong is what must not be done. It is just that among those options that must not be done some are more severe wrong-doings than others.

the situation, and his moral preferences (where 'rational' expresses the kind of belief-and-preference relative rationality employed in standard decision theory). Choosing B is also (to some extent) a *morally* virtuous action in the sense that it *expresses* a virtuous motivation, i.e., being moved by one's moral preferences and evaluative beliefs. Even if the agent herself lacks this moral motivation, choosing B can be seen as virtuous (to some extent), in an *external* sense, for it is an action a certain kind of virtuous person *would* choose, namely a person who is moved by her moral preferences and evaluative beliefs.

Now, Zimmerman thinks that whether to call choosing B the morally right choice or the rational choice is 'a merely verbal quibble' (p. 45), and he goes on to say that:

I frankly do not care whether the requirement to give John drug B (in the Jackson-case) is said to be moral, rational, reasonable, or something else, as long as it is recognized that it exists and that it is this requirement that is of ultimate moral concern to the conscientious person. (p. 45)

I doubt that this quibble is merely verbal. My own account, unlike Zimmerman's, has the virtue of being applicable in certain cases of *fundamental moral uncertainty*, cases where we are not certain about which moral theory is correct.⁷ Here what is morally right and what is rational can come apart. Table 2 illustrates a schematic case where the agent is uncertain about whether M₁ or M₂ is the true moral theory (M₁ and M₂ can be any kind of moral view, Objective, Subjective, Prospective, or what have you):

TABLE 2

Actions	Moral hypotheses	
	M ₁	M ₂
A	right	right
B	right	major wrong

7 Krister Bykvist, 'How to do wrong knowingly and get away with it,' and Krister Bykvist 'Evaluative uncertainty and consequentialist environmental ethics.'

Suppose M_1 is true. Then both A and B are morally right. (To avoid complications let us also assume that none of these actions is supererogatory.) But B is not a rational choice for someone who cares about morality, for it is not rational to risk doing wrong without any possible gain when one can avoid the risk altogether. Indeed, doing so would be rationally reckless. More specifically, an agent who cares about morality will prefer right-doings to wrong-doings and be indifferent between different right-doings (at least when the right-doings have the same moral value, which we have assumed for simplicity). This means that the morally concerned agent will prefer the outcome (A, M_2) to (B, M_2) and be indifferent between (A, M_1) and (B, M_1). But then, given the agent's preferences and beliefs, the rational choice is A, not B, since A *weakly dominates* B in the sense that, no matter whether M_1 or M_2 is true, A will lead to an outcome that is weakly preferred (strictly preferred or seen as indifferent to) the outcome of B. So, we have a case where rational choice and moral rightness come apart: B is morally right (since we have assumed that M_1 is true) but not the rational choice.

Unlike Zimmerman's theory, my account is also applicable in certain Jackson-cases of fundamental moral uncertainty, where we can compare the severity of wrongness across the considered moral theories.⁸ Table 3 illustrates such a case (where the numbers represent degrees of severity of wrongness):

In this situation, it is rational for the agent to prefer B to A, given his beliefs and preferences, since the possible (probability-weighted) gain, from

TABLE 3

Moral hypotheses		
Actions	M_1 ($p = 0.5$)	M_2 ($p = 0.5$)
A	0 (right)	-100 (major wrong)
B	-6 (minor wrong)	-6 (minor wrong)
C	-100 (major wrong)	0 (right)

8 What about fundamental *rational* uncertainty, cases where you are not certain which fundamental rationality principle to use? There are two main options here, which has to be considered no matter whether you accept Zimmerman's Prospective View. The first is to assume that what is rational to do is not sensitive to uncertainty about what is rational. You rationally ought to do what you rationally ought to do, and that is it. Another option is to distinguish between different levels of rationality: what is rational (at level 1) when one is not uncertain about what is rational (at level 1), what is rational (at level 2) when one is uncertain

major wrong to minor wrong, outweighs the possible (probability-weighted) loss, from right to minor wrong. More generally, as a rational person he prefers bringing about the prospect $(x, 0.5, y)$ rather than $(z, 0.5, u)$, if his preference for x over z is much stronger than his preference for u over y . For the same reason it is also rational for him to prefer B to C , (since C 's prospect is just a permutation of A 's). But then we get the result that for a morally conscientious agent it is rational to prefer B to all alternatives in a Jackson-case with fundamental moral uncertainty. In other words, it is rational for him to prefer doing something that he knows is morally wrong. In fact, in contrast to the dominance case above, in this case it is *impossible* to act both morally and rationally right.

So, it is pretty clear, I think, that the choice between calling the second-best option in the original Jackson-case the rational choice or the morally right choice is not a mere verbal quibble, for the rational choice approach can be applied to cases of fundamental moral uncertainty, where what is rational and what is moral will not always coincide.

Furthermore, the fact that Zimmerman's constraint #2 rules out that the morally conscientious person can conscientiously choose B in the Jackson-case with moral uncertainty, tells against this constraint.

Note also that my account does not deny that the morally conscientious agent has an ultimate concern for moral rightness. On the contrary, the morally conscientious agent's fundamental preferences track moral rightness and degrees of severity of wrongness in the ways laid out above.

4 Substantive Objections to the Prospective View

Even though Zimmerman does not provide a clear definition of the Prospective View, it is pretty clear what it will imply for certain cases. He is well aware of the fact that some of these implications are counter-intuitive. For example, he grants that the Prospective View will have the implication that, in some situations, the morally right thing to do is to bring about a horrific outcome

about what is rational (at level 1), but not uncertain about what is rational (at level 2), what is rational (at level 3) when one is uncertain about what is rational (at level 2) but not uncertain about what is rational (at level 3), and so on. Note that this does not lead to a troubling infinite regress for normal rational agents, since no such agent can simultaneously entertain uncertainty at an infinite number of levels. For more on this, see Krister Bykvist, William MacAskill, Toby Ord, (forthcoming).

(p. 61).⁹ Even though Zimmerman is aware of the general problem, I do not think it has been made sufficiently clear exactly how counter-intuitive his new version of the Prospective View is, and how radically it differs from traditional moral theories in this respect. Furthermore, once we have established the counter-intuitiveness of his theory by using a very telling example, we are in a better position to engage with his methodological point that we should not solely rely on the counter-intuitiveness of a theory in deciding whether to do adopt it.

The substantive objection to Zimmerman's theory can be made more powerful by adding more details to the situation suggested above. Let us call this richer version of the situation the *Horrible Action Case*. Suppose that the agent has the option of doing something without any gains for others – to torture people for fun, commit ethnic cleansing, or destroy a whole country, for example. He knows all the details of the horror of the (unique) outcome of this action. But he would take great pleasure in the horrors because he is morally vicious. So, if he were to perform the action, he would do that from the vicious motive of enjoying the sight of the horrors, and the maxim he would act on could not be universalized. He also has *available* strong evidence that what he is doing is the morally worse thing he could possibly do and that it would be deeply morally wrong. Can it still be morally obligatory for him to perform the horrible action?

All traditional moral theories I know of would deny that this action could be morally obligatory in the Horrible Action Case. But different theories would focus on different factors from the long list just given when explaining why the action is wrong and not obligatory. For example, consequentialists would focus on the actual or expected consequences (which are the same in this case); virtue theorists would focus on his vicious motives or vicious character traits; and Kantians would focus on his evil maxim, or the fact that he would be treating people merely as means to his evil ends.

Even Zimmerman's earlier version of the Prospective View would agree that the horrific act is morally wrong in the Horrible Action Case because it would not be the best bet, given the agent's *available* evidence about what is morally better. However, in the light of Holly Smith's criticism, which pointed out that his earlier version could not avoid saying that the guaranteed second-best option is morally wrong in some Jackson-style cases, Zimmerman revised the view so that what we ought to do depends on *availed* empirical and evaluative

9 This is also discussed in Gideon Rosen's review of Zimmerman's book, in *Notre Dame Philosophical Reviews*, 2015-03-09.

evidence, that is, evidence which is actually believed by the agent (or believed to a degree that is proportioned to the strength of the evidence).¹⁰

As far as I can see, Zimmerman's new version of the Prospective View does not rule out that the horrible action in the Horrible Action Case is morally obligatory. For it is enough that the availed evidence, the available evidence believed by the agent, implies that the best bet, relative to the availed evaluative evidence, is to perform the horrible action. This can happen if the availed evaluative evidence is very misleading or incomplete. For an example of the latter, suppose that the agent in the Horrible Action Case has *only* availed himself to some evidence that supports the claim that the pleasure he would enjoy if he were to perform the horrible action would be good for him. If no other evaluative evidence is availed, then it seems he ought to perform the horrible action, since this is his best bet, relative to the actually availed evaluative evidence.

This, I must say, seems to me like a pretty good knockdown objection to the Prospective View. The agent's horrible action in the Horrible Action Case looks like a *paradigm case* of moral wrong-doing. But Zimmerman can of course still say that *something* is morally wrong in the situation above. For example, the agent is vicious, and he also fails to do what is actually best. But he still has to deny that this paradigm is a case of moral wrong-*doing*, which is a big bullet to bite.¹¹

Zimmerman attempts to turn the tables and points out that a more objective view will also have unwelcome consequences. For example, consider a theory that says that the doctor ought to give the patient the medicine that in fact will fully cure the patient because this is the option which is best for the patient. But to choose this option would be to be *extraordinarily reckless*, Zimmerman thinks. So, we seem to have to choose between theories, none of which lack unwelcome implications.

Zimmerman thinks that this shows that it is a mistake to choose a theory here *solely* on the basis of whether it has unwelcome implications since all theories have unwelcome implications (p. 61). That might be true, but this does

10 Holly Smith, 'The "Prospective View" of Obligation,' *Journal of Ethics and Social Philosophy*, Discussion Note, (2011), pp. 1–8.

11 Wouldn't my account say that it can be *rational* for the agent to perform the horrific action, given his beliefs and immoral preferences? Yes, but it is no news to be told that immoral agents can perform rational actions in the belief-and-preference relative sense invoked by standard decision theory. What is newsworthy – indeed shocking – is to be told that deeply immoral agents can be morally obligated to deliberately perform morally horrific actions.

rule out that we can still reasonably compare different theories on the basis of their counter-intuitiveness. A moral theory that is less counter-intuitive is preferable, in one important respect, to one that is more so.

This means that, even if it is true that the more objective view supports extraordinarily reckless actions in some situations, this has to be compared to the degree of counter-intuitiveness that the Prospective View is saddled with. On the face of it, the extraordinary recklessness dwarfs in comparison the extremely high degree of counter-intuitiveness that afflicts the Prospective View in the case presented above.

Furthermore, it is important to note that the recklessness-objection can be weakened by considering a moral view that assigns some kind of negative value to recklessness. A virtue theory could say that it is a vice to act recklessly and that this makes the doctor's action wrong; a consequentialist who only counts actual outcomes could say that imposing great unwanted risks to a patient is intrinsically bad, and an expected value consequentialist, like Jackson himself, could say that the doctor ought to maximize expected value, where the expectation is calculated by taking into account the agent's credence and actual intrinsic value of outcomes. Any of these views seems less counter-intuitive than the Prospective View.

But, as Zimmerman points out, for each of the theories above we can construct a Jackson-case where the theory in question has to say that the guaranteed second-best option is morally wrong. But then we need to remember that we can still say that it would be *rational* for a morally conscientious agent to perform the second-best option. So, some sting of this objection seems to be taken out.

5 Concluding Remarks

Zimmerman is insistent on finding an acceptable moral theory that is not vulnerable to Jackson-cases. But perhaps this is just a wild goose chase (to use a metaphor Zimmerman adopts in another context). His own theory can avoid the original Jackson-case, as well as a Jackson-case with evaluative uncertainty. But, as Holly Smith points out, in order to avoid Jackson-cases where one has not actually availed oneself of all of one's available evidence, he has to make moral obligations dependent only on availed evidence.¹² But this makes the

12 Perhaps this is not enough, for could one not construct a plausible Jackson-style case where what is best, given one's *manifest* beliefs about available evidence, is guaranteed not

theory extremely counter-intuitive. Furthermore, even if we go this far, Zimmerman's account is still vulnerable to Jackson-cases of fundamental moral uncertainty. The two culprits here seem to be constraint #2, that one never conscientiously does what one believes to be wrong, and the idea that moral obligation is sensitive to evaluative uncertainty, two claims I suggest we should therefore drop.

To end on a more conciliatory note, I agree with Zimmerman that it is often a good moral ideal to act like the morally conscientious person. But whereas he thinks that, in cases of empirical certainty, we morally ought to do what would be most reasonable to do, given our *actually availed evaluative evidence*, if we were morally conscientious, I think instead that we morally ought to do what would be most reasonable for us to do, if we were morally conscientious and *knew* all the relevant evaluative facts. When we do not know the relevant evaluative or moral facts, but still care about value or moral wrongness, there is still a *rational* choice to be made, given our moral preferences and beliefs about empirical, evaluative, and moral matters.

Acknowledgements

I would like to thank Michael Zimmerman, Holly Smith, Benjamin Kiesewetter, and the audience at the APA, in Seattle, 2017, for very useful comments.

Author Biography

Krister Bykvist is Professor of Practical Philosophy at the Department of Philosophy, Stockholm University, and Research Fellow at the Institute for Futures Studies, Stockholm.

References

Krister Bykvist, 'Objective versus Subjective Moral Oughts,' *Uppsala Philosophical Studies* 57: 39–55, (2009).

to be best, given one's *dispositional* beliefs about available evidence? Would Zimmerman be willing to go even further and define moral obligation in terms of only manifest beliefs about available evidence?

- Krister Bykvist, 'How to do wrong knowingly and get away with it,' *Uppsala Philosophical Studies* 58, (2011).
- Krister Bykvist, 'Evaluative uncertainty and consequentialist environmental ethics,' in L. Kahn and A. Hillel (eds.), *Environmental ethics and consequentialism*, (London: Routledge, 2014).
- Krister Bykvist, William MacAskill, Toby Ord, *Moral uncertainty*, (Oxford: Oxford University Press, forthcoming).
- Peter Graham, 'In Defence of Objectivism about Moral Obligations,' *Ethics* 121 (2010), pp. 88–115.
- Gideon Rosen, 'Review of Michael J. Zimmerman: Ignorance and Moral Obligation,' in *Notre Dame Philosophical Reviews*, 2015-03-09.
- Holly Smith, 'The "Prospective View" of Obligation,' *Journal of Ethics and Social Philosophy*, Discussion Note (2011), pp. 1–8.
- Michael Zimmerman, *Ignorance and moral obligation*, (Oxford: Oxford University Press, 2014).