

In Defense of Objectivism about Moral Obligation*

Peter A. Graham

This is a preprint of an article whose final and definitive form will be published in *Ethics* 2010; *Ethics* is available online at: <http://www.journals.uchicago.edu/toc/et/current>.

There is a debate in normative ethics about whether or not our moral obligations depend solely on either our evidence concerning, or our beliefs about, the world. *Subjectivists* maintain that they do and *objectivists* maintain that they do not.¹ I shall offer some arguments in support of objectivism and respond to the strongest argument for subjectivism. I shall also briefly consider the significance of my discussion to the debate over whether one's future voluntary actions are relevant to one's current moral obligations.

1. Preliminaries

First, a terminological stipulation, a conceptual presupposition, and a distinction. Throughout I treat 'X has a moral obligation to ϕ ' and 'X is morally obliged to ϕ ' as synonymous, meaning X has an all-things-considered moral obligation to ϕ . And I presuppose, as a conceptual truth, that it is morally wrong for X to ϕ if and only if X is morally obliged not to ϕ . I also distinguish between a person's *objective situation* and her *evidential situation*. A person's objective situation at a time consists of all the facts about her at that time except those about what evidence she has about the world. A person's evidential situation at a time consists in the totality of her evidence about the world at that time.²

Now, some definitions:

A moral theory, *T*, is *evidence-subjective* =_{def.} according to *T*, necessarily, a person has the moral obligations that she has at a time solely in virtue of facts about her evidential situation at (or prior to) that time.³

A moral theory, *T*, is *objective* =_{def.} it is not the case that *T* is evidence-subjective.

Evidence-subjectivism is the thesis that the true moral theory is evidence-subjective and objectivism is the thesis that the true moral theory is objective.⁴ Note that, so defined, objectivism is merely the denial of evidence-subjectivism.⁵ In what follows, I shall only be concerned with evidence-subjectivism and objectivism, so defined.

Before proceeding, I must address a complication with evidence-subjectivism. Most who call themselves ‘subjectivists’ can’t accept evidence-subjectivism as I’ve defined it. They can’t because, as I’ve defined it, evidence-subjectivism is inconsistent with the following principle that they also accept:

‘Ought’ Implies ‘Can’ (OIC): Necessarily, *S* is morally obliged to ϕ only if *S* can ϕ .⁶

As a person’s evidence may indicate that she can ϕ even when she can’t, if **OIC** is true, evidence-subjectivism is false. Take a being that lacks the capacity to ϕ whose evidential situation matches that of an agent who is morally obliged to ϕ . If **OIC** is true, then, despite being in the same exact evidential situation as the agent obliged to ϕ , the incapacitated being, because of **OIC**, is not so obliged. We need another definition:

A moral theory, *T*, is **ability-constrained-evidence-subjective** =_{def.} according to *T*, a person has the moral obligations that she has at a time solely in virtue of both facts about her abilities and facts about her evidential situation at (or prior to) that time.

Most subjectivists hold that the true moral theory is ability-constrained-evidence-subjective.⁷

Ability-constrained-evidence-subjectivists, though shying away from full-blooded evidence-subjectivism, retain its spirit in cleaving to the idea that a person’s moral obligations are *almost* solely determined by her evidential situation. But, precisely because of that ‘*almost*’, a worry lingers. Ability-constrained-evidence-subjectivists are under considerable pressure to explain why a person’s moral obligations are determined by her evidence in this only-almost way. The motivation behind evidence dependence, after all, is that a person’s moral obligations must, in some way, be readily accessible to her.⁸ But, as what a person can do may not be readily accessible to her, ability-constrained-evidence-subjectivism outstrips its intuitive motivation. The ability-constrained-evidence-subjectivist has two inconsistent ideas—that a person can’t be morally obliged to do something if she can’t do it and that a person’s moral obligations must always be readily accessible to her—and so, she lets in a little slack with one in order to accommodate the other. This highly suspect and uncomfortably cobbled together motivation seems unstable.

I won't press this criticism, however. Henceforth, I'll focus on ability-constrained-evidence-subjectivism and refer to it simply as 'subjectivism'.

2. The Case for Objectivism

Now I'll very briefly present two arguments for objectivism.

The first argument is due originally to Sir David Ross⁹:

Many people would be inclined to say that the right act for me is...that which on all the evidence available to me I should think to be my duty. But suppose that from the state of partial knowledge in which I think act *A* to be my duty, I could pass to a state of perfect knowledge in which I saw act *B* to be my duty, should I not say 'act *B* was the right act for me to do'?¹⁰

When we change our view about our moral obligations as a result of acquiring information, we don't take it that our moral obligations have changed. Rather, it seems to us that we, at last, come to see what our moral obligations were all along. If this seeming is not delusory, then objectivism must be correct.

The second argument for objectivism goes as follows: the question I want answered when I ask myself what my moral obligations are is the same as that which I want answered when, in seeking your help, I ask you what they are; but, to adequately answer me you don't need to consider my evidence concerning my situation; therefore, my moral obligations don't depend on my evidence concerning my situation; so, objectivism is true.

Consider the following scenario:

PROMISE: I have promised Kenneth that I will turn on a certain light at 6pm. All my evidence indicates that it is not 6pm. I ask you what I'm morally obliged to do. You know of my promise and that it is 6pm, but you also know that I neither believe, nor have any evidence, that it is 6pm. You tell me that I'm morally obliged to turn on the light.

Here's an argument:

- (1) In PROMISE, you speak truly in saying I'm morally obliged to turn on the light.
- (2) According to no plausible subjectivist moral theory do I have a moral obligation to turn on the light in PROMISE.
- (3) Therefore, objectivism is true.

(1) is intuitive and (2) follows from the description of the case.¹¹

Zimmerman considers and rejects an argument of this sort.¹² The argument he considers involves a case in which Jill can do one of three things: A, B, and C. She knows that A will help John significantly, though not completely, but she does not know of B and C which will help him completely and which will kill him. Zimmerman grants that if Jill asks a third party, Jack, what she is morally obliged to do and Jack knows that C will completely help John, Jack acts rightly only if he says that she is morally obliged to do C. But, Zimmerman claims, though Jack acts rightly only if he does this, that's not because he speaks truly, but because he, like Jill, is morally obliged to help John as best he can, and speaking falsely, thereby getting Jill to do C, is best for John. Thus, Zimmerman holds that cases like these can be treated as instances of morally obligatory lying.

This isn't an adequate response to Zimmerman's own case; it just doesn't seem to be a case of morally obligatory lying. More important, however, this kind of reply has no force against my argument. In PROMISE, you're morally obliged neither to make sure I keep my promise nor to shield Kenneth from disappointment. So, even if the intuitive thought weren't that in PROMISE you speak truly, but merely that you can only act rightly by telling me I'm obliged to turn on the light, that can't be explained by your having some moral obligation to speak falsely to me. Why you're required to tell me I'm morally obliged to turn on the light must be because you're morally obliged not to knowingly say false things to me. And so, in PROMISE you speak truly when you say I'm morally obliged to turn on the light. Zimmerman's response, as applied to my argument, then, fails.

These arguments seem to show that objectivism is true. In the absence of compelling reason to think otherwise, then, one should accept objectivism.

3. The Case for Subjectivism

In this section I'll discuss two arguments for subjectivism. Whereas the first confuses blameworthiness facts for obligation facts, the second does not. I'll present the first argument and expose its confusion. Then I'll consider and dismantle an attempt to dissolve the objectivism/subjectivism debate by way of a postulated ambiguity in 'morally obliged'. Finally, I'll consider the second argument. This argument is the best out there for subjectivism. I'll devote the rest of this section and the next to explaining how to respond to it.¹³

The first argument begins with the following case:

SYRINGES: Luke's patient, Nathan, has a painful non-fatal disease. In syringe C is the cure for the disease and in syringe P is a lethal poison. Unfortunately, the labels on the syringes have, unbeknownst to Luke, been switched.

To many subjectivists it's just obvious that Luke is morally obliged to inject P. They also find it equally obvious that had the labels not been switched, Luke would then have been morally obliged to inject C. Subjectivism, they conclude, must be true.¹⁴

This argument is not persuasive. What's obvious is not that Luke is morally obliged to inject P, but rather, that he would be morally blameworthy were he to inject C. But it doesn't follow from this that he's not morally obliged to inject C. Whether one is morally obliged to do something depends on how one's doing it impacts, and fits with, the world; whether one is morally blameworthy for doing something depends (roughly) on what one takes oneself to be doing when one does it. Clearly what you take yourself to be doing when you do something depends on your beliefs about, and evidence concerning, your situation. But how what you do impacts, and fits with, the world depends (at least in part) on your objective situation.

In addition, there are independent intuitive grounds for thinking that moral obligation and blameworthiness can come apart in both directions. First, there is the phenomenon of blameless wrongdoing. If when leaving a party I take your jacket, mistaking it for mine, it's intuitive that I do something wrong even though I may be completely blameless. Second, a person may be morally blameworthy even when she does what she knows she's morally obliged to do. For example, X might be obliged to chop off Y's leg—suppose X and Y both know that it's gangrenous and Y begs X to chop it off—but, if X does so not to help Y, but, rather, to cause Y excruciating pain, X will be blameworthy even though she does what she knows she is obliged to do.¹⁵ Blameworthiness is one thing, obligation another. One should not slide, then, as the subjectivist does, from the thought that Luke is blameless in injecting P, and would be blameworthy for injecting C, in SYRINGES to the thought that he is not morally obliged to inject C in that scenario.

At this point, one might try to dissolve the dispute between objectivists and subjectivists by suggesting that 'morally obliged' is ambiguous. Perhaps there is one sense of 'morally obliged' corresponding to the subjectivist view of moral obligation and one corresponding to the objectivist view of it. If that's right, then perhaps objectivists and subjectivists are both right:

though Luke has a *subjective* moral obligation to inject P in SYRINGES, he does not have an *objective* moral obligation to do so. Beyond this, one might think, there's just nothing more to be said.

This apparently tidy resolution of the debate, on closer examination, however, is neither tidy nor acceptable. First, it's not clear that 'morally obliged' is ambiguous; subjectivists and objectivists don't seem to be arguing past each other.¹⁶ Second, and more important, this sense-splitting maneuver simply can't dissolve the dispute between objectivists and subjectivists. Even if there are two types of moral obligation, the question remains: in which are we interested when doing moral theory about action? That is, in which sense of 'moral obligation' are we interested when we debate whether Consequentialism or some other theory provides the correct account of our moral obligations with respect to action? Even if there are multiple senses of 'moral obligation' we still need an answer to this question. Now, one might think that we're at a complete loss when trying to answer this question for it might seem that we have no way of isolating the sense of 'moral obligation' we're interested in when we do moral theory about action. On the contrary, however, there is quite a natural way of doing this: the sense of 'moral obligation' we're interested in when doing moral theory about action is that with which the morally conscientious person is concerned when she deliberates about what to do.¹⁷

If this is right, then we can make some progress. We have intuitions both about moral conscientiousness and about how a morally conscientious person would behave in various circumstances, and we can deploy these intuitions to focus the objectivism/subjectivism debate. Even if there are multiple senses of 'moral obligation', the question at issue between objectivists and subjectivists is: with which sense of 'moral obligation' is the morally conscientious person concerned in her deliberations about what to do? Objectivists will claim it's the objective sense and subjectivists will claim it's the subjective sense. No amount of sense-splitting is going to resolve this dispute. And it surely won't do to suggest that the morally conscientious person is concerned with both senses of 'moral obligation'. In her deliberations about what to do, the morally conscientious person's concern is so single-minded as to preclude her being concerned with more than one type of moral obligation.

The SYRINGES-based argument fails to win the case for the subjectivist because it implicitly relies on the following premise:

- (I) Necessarily, a morally conscientious person never does something she is morally obliged not to do.

As I have argued, this premise derives its plausibility from a mistaken tethering of moral obligation and blameworthiness. Once the two phenomena are untangled, however, the surface allure of (I) evaporates. The SYRINGES-based argument thus provides the subjectivist no leverage against the objectivist. If, on the other hand, the subjectivist can fashion an argument against objectivism that does not rely on such a premise, then, perhaps, she can gain the upper hand.

Recently, a new kind of subjectivist argument has emerged that seems to do just that. The version of it I will discuss is based on the following case originally offered by Frank Jackson:

DOCTOR: Jill's patient, John, has a very painful, non-fatal disease. Jill has three drugs: A, B, and C. A would completely cure him; B would only partially cure him; and C would kill him. But, though Jill knows that B would only partially cure John, her evidence leaves it open which of A and C would cure him and which would kill him.^{18,19}

How is this case meant to establish that objectivism is false? According to Jackson, it's just "obvious" that Jill is morally obliged to prescribe B in DOCTOR.²⁰ Contra Jackson, this is not obvious. Again, what's obvious is that Jill would be morally blameworthy were she not to prescribe B. But distinguishing between moral obligation and blameworthiness will only go so far here. For though it is intuitive that one may be blameless in failing to do what one is morally obliged to do, it seems plausible that that should only happen when one doesn't know that one is failing to do what one is morally obliged to do when one does what one does. But the objectivist can't maintain that this is true in DOCTOR.

What's special about DOCTOR is not just that it is a situation in which Jill, insofar as she is morally conscientious, would do what she has an objective moral obligation not to do, but, rather, that it is a situation in which *she knows when she does what she does* that she has an objective moral obligation not to do it. If she is an objectivist, Jill will recognize that one of A and C, and not B, is the one she is morally obliged to prescribe, but, this notwithstanding, if she is morally conscientious, she will prescribe B nonetheless. It is this that gives the DOCTOR-based argument the dialectical oomph the SYRINGES-based argument lacks. It implicitly relies not on (I), but on the seemingly uncontestable:

- (II) Necessarily, a morally conscientious person never does something such that when she does it she knows she is morally obliged not to do it.

So, if being morally conscientious requires prescribing B, Jill's being morally conscientious precludes her believing objectivism. But then she can't be both an objectivist and a morally conscientious person. That, so goes the argument, shows that objectivism must be false. How could objectivism be correct and yet be such that anyone who wholeheartedly endorses it, necessarily, is not morally conscientious?

These reflections alone don't yet constitute a valid argument against objectivism. But we can easily work them up into one by adding the following premise:

- (III) It is possible for a person to both believe the true moral theory and be morally conscientious.

If (III) is true, then, as being morally conscientious requires being such that one would prescribe B in DOCTOR, and one could not do that if (a) (II) were true, (b) the true moral theory were objective, and (c) one believed it, it follows that objectivism is false. And, as (II) and (III) look plausible, the objectivist's goose may seem pretty well cooked.

4. Against the Argument for Subjectivism

The argument for subjectivism has four crucial premises:

- (II) Necessarily, a morally conscientious person never does something such that when she does it she knows she is morally obliged not to do it.
- (III) It is possible for a person to both believe the true moral theory and be morally conscientious.
- (IV) According to any plausible objectivist moral theory, Jill is morally obliged not to prescribe B in DOCTOR.
- (V) In DOCTOR, if Jill is morally conscientious, she will prescribe B.

Are they plausible? (V) seems beyond question. (IV) is also hard to dispute. It seems clear that (barring special mitigating circumstances, ones clearly absent in DOCTOR) a doctor, with full information, when faced with a choice between completely curing, partially curing, or killing a patient, is morally obliged to completely cure him. And given that according to objectivism what moral obligations a person has when faced with a choice under uncertainty in situations like these just are those she would have were she faced with that very same choice under full information, (IV) must be true. (III) looks quite plausible as well. It would be a very queer thing indeed if the correct moral theory were such that merely believing it entailed that one was not

morally conscientious. Perhaps (III) is false, but I can't come up with a reason for thinking that it is.²¹ To reply to the DOCTOR-based argument, then, one needs to tackle (II).

Why does (II) seem plausible? Here's why. The morally conscientious person, as we've identified her, is single-mindedly focused on satisfying her moral obligations. The only way she would ever ϕ while at the same time knowing that she was morally obliged not to ϕ , it can seem, is if she also knew that what she was morally obliged to do was ϕ . Many think, however, that it's impossible to be both morally obliged to ϕ and morally obliged not to ϕ . If all of this is right, then it is impossible for a morally conscientious person ever to do something she knows, when she does it, that she is morally obliged not to do. And so, if we think about moral conscientiousness in this way, (II) seems unquestionable.

We needn't think about moral conscientiousness in this way, however.²² We should agree that, in her deliberations, the morally conscientious person is solely concerned with her moral obligations. But, and here's the point, we needn't affirm that the morally conscientious person's sole concern with her moral obligations is as single-faceted as the above line of reasoning has it that it is. The morally conscientious person's concern with her moral obligations is not limited only to being concerned to do what she is morally obliged to do. Rather, the morally conscientious person has (at least) two concerns with respect to what she is morally obliged to do: to do what she is morally obliged to do and to *not do* what she is morally obliged *not* to do. The focus of the one concern is doing what morality commands while the focus of the other is avoiding doing what morality forbids.²³

But how does this construal of moral conscientiousness help the objectivist deny (II)? Another observation is in order: though, in a particular situation, there are no degrees of moral "obliged-ness", there most certainly can be varying degrees of moral "obliged-not-ness". Though there can be only one thing that a person is morally obliged to do in a particular situation, there may very well be more than one thing she is morally obliged not to do.²⁴ (I'm assuming here, and henceforth, that moral obligation dilemmas—situations in which all of a person's options are ones she's obliged not to take—are impossible.) And, what's more, in such a situation, it makes perfect sense to ask of the various different things that she is morally obliged not to do which she is most morally obliged not to do in that situation. Now, if one of the morally conscientious person's concerns is to not act morally wrongly and moral wrongness comes in degrees, then it's quite natural to think that a morally conscientious person, though certainly concerned not to do

either, will be more concerned not to do what she is very seriously morally obliged not to do than she will be not to do what she is much less seriously morally obliged not to do. We should think, then, that the strength of the morally conscientious person's concern not to act wrongly will vary in direct proportion to the wrongness of the actions she is concerned to avoid performing.

Now, if the morally conscientious person has not one goal, but two, then she might find herself in a situation in which she has to make certain tradeoffs between them. The point here is very general. Whenever a person has multiple distinct goals, there is at least the conceptual possibility of her confronting a situation in which pursuing one of them is incompatible with, or makes less likely, successful pursuit of others. Teachers, for example, not only have the goal of accuracy, i.e., teaching their students only truths, but also that of breadth of coverage, i.e., covering a sufficiently wide spectrum of topics. It often happens that these goals come into conflict. Sometimes a teacher must sacrifice accuracy for breadth of coverage; in order to not get bogged down in certain details, she presents as fact to her students information which, though close to the truth, is not, strictly speaking, true. Other times a teacher must sacrifice breadth of coverage for accuracy; in order not to make a hash of all the topics she covers by going over them too quickly, she omits discussing certain topics which deserve to be covered.

Likewise, if the morally conscientious person has two distinct goals, then her goals might similarly come into conflict. Pointing to this possibility, however, only helps the objectivist if it makes it plausible that (II) is false. As I'll now explain, it does.

Moral philosophy is dominated by "two-option cases", i.e., cases in which a person faces a choice simply between performing and not performing some action. Most real life cases, however, are not two-option cases. Most often, when faced with a choice, people are confronted with an array of morally relevant options. The unrealistic focus of moral philosophy on two-option cases is not to be criticized, however. It is important when doing theory to confront simple cases first. Only after having crafted principles covering simple cases should one try to extend those principles to more complicated cases. This is a sound methodological strategy in any theoretical discipline and it should no more be abandoned in ethics than it should in any other such discipline.

What this simplifying focus on two-option cases can do, however, is blind one to the fact that there are many cases in which the morally conscientious person's goals of doing what she is morally obliged to do and of not doing what she is morally obliged not to do come into conflict

with each other. This is because two-option cases have the following interesting feature: if a person knows she has but two options and she knows that one of her options is one she is morally obliged not to take, then she straightaway knows that her only other option is the one that she is morally obliged to take. Thus, in any such situation, if a morally conscientious person knows that, of X and Y, X is one that she is morally obliged not to take, she thereby straightaway knows that taking her other option, Y, will best satisfy both of her goals: in taking it she will both do what she is morally obliged to do and avoid doing anything she's morally obliged not to do.

This feature of two-option cases, however, is unique to two-option cases. In more-than-two-option cases, a morally conscientious person's knowing of a certain option, X, that she's morally obliged not to take it provides her with knowledge neither of which of her options best satisfies her goals, nor, more importantly, of which *worst* satisfies them. This is because not only is being morally obliged not to take X clearly consistent both with taking Y's being obligatory and with taking Z's being obligatory, but it is also consistent both with Y's being, and with Z's being, the option that worst satisfies both of the morally conscientious person's goals. If Y is the option that she is morally obliged to take and Z is an option that she is much more seriously obliged not to take than she is not to take X, then it is Z and not X that worst satisfies both her goals. (This is because, though Z fails to satisfy her goal of doing what she is morally obliged to do just as much as X does, it much more seriously fails to satisfy her goal of not doing what she is morally obliged not to do than does X.)

How does all of this show that (II) is false? Well, given her two goals, it seems quite plausible that if a morally conscientious person finds herself in a situation in which

- (1) she doesn't know which of her options will best satisfy both of her goals, and
- (2) she knows of a certain option, X, that it does not best satisfy both of her goals but that taking it much better satisfies them than does another option she would have to risk taking were she to attempt taking the option that best satisfies both of them

she might very well take option X—and if she knows that the morally worst option is very substantially morally worse than X, she most certainly *would not* do otherwise than take X—even though she knows that she's morally obliged not to take it: taking any other option in such a situation would demonstrate so insufficient a concern for not doing what she is morally obliged not to do as to conceptually preclude her being truly morally conscientious. So, whenever (1) and (2) hold a morally conscientious person might, and, in some instances, would be *conceptually*

required to, do something she knows she is morally obliged not to do.²⁵ If this is right, then (II) is false and the DOCTOR-based argument fails.

If my account of moral conscientiousness is correct, then not only is (II) false, but, also, it's clear why every version of this kind of argument for subjectivism (the kind that implicitly relies on (II)) has a certain conspicuous feature, viz., they each crucially involve a more-than-two-option case. For, though on my account of the morally conscientious person (II) is false, the following principle is certainly true:

(II') Necessarily, a morally conscientious person never, in a two-option case, does something such that when she does it she knows she is morally obliged not to do it.

Why is (II') true on my account of moral conscientiousness? As indicated above, in two-option cases, knowing that one is morally obliged not to take one option suffices for knowing that one is morally obliged to take the other option. And moral conscientiousness is certainly inconsistent with not taking an option while knowing of it that you're morally obliged to take it. That is,

(VI) Necessarily, a morally conscientious person never refrains from doing something such that when she refrains she knows she, nonetheless, is morally obliged to do it

is incontrovertibly true.²⁶ As (II') and (VI) are both true (and explicable) on my account of moral conscientiousness, if my account of moral conscientiousness is correct, then one can't construct a two-option case in which a morally conscientious person would knowingly do what she was objectively morally obliged not to do. That there are no two-option-case versions of this kind of subjectivist argument and my account of moral conscientiousness predicts that there couldn't be further confirms (though perhaps only slightly) my account of moral conscientiousness.²⁷

If the morally conscientious person has two concerns with respect to what she is morally obliged to do—to do what she is morally obliged to do and to not do what she is morally obliged not to do—then (II) is false. But if (II) is false, then the DOCTOR-based argument against objectivism crumbles.

5. A Subjectivist Reply

At this point, the subjectivist might respond as follows. Suppose that in DOCTOR Jill suffers a temporary lapse in moral conscientiousness and prescribes A for her patient and sends him on his

way. Then, suppose that, before seeing the results of her decision, she regains her moral conscientiousness and reflects upon what she's done. Will she not chastise herself by saying, "I ought not to have done that"? If so, the subjectivist will press, does that not show that she was morally obliged not to prescribe A?

No, it does not. Upon regaining her moral conscientiousness Jill may well have the thought that she ought not to have prescribed A, but that does not show that moral obligation is not objective. The culprit here is the word 'ought'. 'Ought' is a notoriously slippery word. There is, to be sure, an 'ought' correlated with moral obligation, but that is not the 'ought' at play in Jill's thought. Beside the 'ought' associated with moral obligation there is, it is widely recognized, also a pragmatic 'ought' associated with means and ends. There are, in fact, at least two such pragmatic 'ought's: a subjective one and an objective one. According to the objective pragmatic 'ought' (ought_{pragmatic (objective)}), a person ought to do something just in case doing it will bring about the outcome, among the various outcomes from which she is choosing, she most prefers relative to her goals in acting. (With respect to the set of goals that a person has in acting she has a preference ranking at the top of which is complete satisfaction of all those goals and the order of the rest of which is determined by her preferences regarding the various different possibilities of partial fulfillment of those goals.) According to the subjective pragmatic 'ought' (ought_{pragmatic (subjective)}), a person ought to do something just in case, roughly, doing so is the output that results from inputting into whatever the correct decision theory happens to be the agent's preference, and subjective probability, functions.

When Jill, after recovering from her momentary lapse in moral conscientiousness thinks "I ought not to have done that", the 'ought' in her thought is not the 'ought' of moral obligation, but, rather, the subjective pragmatic 'ought'. She is quite rightly recognizing, that in prescribing A she did not opt for the course of action dictated by her evidence and her preference ranking over her goals as a morally conscientious person.²⁸ That she would have that thought, then, lends no support to the idea that she was not morally obliged to prescribe A in DOCTOR.

"But wait," the subjectivist might protest. "The 'ought' in Jill's thought is not a pragmatic 'ought'. It's clearly a *moral* 'ought'. It has a distinctly moral cast to it; when she entertains it she'll be riddled with guilt and self-blame." True enough. But, neither of these observations establishes that the 'ought' in her thought is the 'ought' of moral obligation.

Take the first observation first. The ‘ought’ in Jill’s thought indeed has a moral cast to it. But that is certainly consistent with its being a pragmatic ‘ought’. If the pragmatic ‘ought’ is, as I have indicated, an ‘ought’ relative to the goals of the agent in question, then it is natural that that ‘ought’ take on the cast of the goals to which it is relativized. The moral flavor of the ‘ought’ in Jill’s thought, then, is easily explained by the fact that the goals Jill has as she entertains it after recovering her moral conscientiousness are the two thoroughly moral goals of the morally conscientious person. The thought will also have a heightened urgency because Jill will, to the degree that she has recovered her moral conscientiousness, realize that there is a significant chance that she has done something horribly morally wrong and urgently want that not to be the case.

Now take the second observation. Though Jill will feel a significant sense of guilt and self-blame, that doesn’t establish that she was not morally obliged to prescribe A. As noted in section 3 in connection with the SYRINGES-based argument for subjectivism, blame and moral obligation are different things. So, attempting to undermine the objectivist’s response to the subjectivist’s DOCTOR-based argument against objectivism by appeal to these kinds of considerations about guilt and self-blame simply won’t get off the ground. Such an appeal is nothing more than the SYRINGES-based argument, once again, in new clothes.

Not only does this attempt to scuttle my reply to the subjectivist’s DOCTOR-based argument run aground, but it actually also provides grist for the objectivist’s mill. How so? Consider, once again, Jill, who has lapsed in her moral conscientiousness and prescribed A in DOCTOR. The subjectivist is correct that when Jill regains her moral conscientiousness she’ll have the thought that she ought not to have prescribed A; but he leaves out the whole story if he stops there. Not only will Jill think that she ought not to have prescribed A, she’ll also think that the reason why is that doing so was *too risky*. Now, the natural thing to say about what makes Jill’s prescribing A too risky, is that, in doing so, she has risked acting horribly wrongly. The subjectivist, however, cannot say this. If it was definitively morally wrong for Jill to have prescribed A, as the subjectivist contends, then the explanation of why prescribing A was morally wrong can’t be that in doing so Jill ran too high a risk of acting horribly morally wrongly.²⁹ So, if prescribing A is morally wrong, and wrong in virtue of its riskiness, then its riskiness must be a riskiness with respect to something other than horrible wrongdoing. But riskiness with respect to what?

It might seem that the subjectivist has a ready reply. There does seem to be a plausible factor—namely, harm—the risk of which a subjectivist might appeal to in order to explain the wrongness of Jill’s prescribing A in DOCTOR. He might say, then, that what makes Jill’s prescribing A morally wrong is that it runs too high a risk of causing terrible harm to her patient. The problem with this reply, however, is that it only has a shot of working in cases in which harm is intuitively the only morally relevant factor. But there are other similar kinds of cases in which an appeal to the risk of harm simply won’t do the trick. Consider

SURGEON: Jane is a transplant surgeon who has five dying organ failure patients. Bloggs and Slade are each both unconscious and such that by killing him and redistributing his organs, Jane can save all five of her organ failure patients. But though one, but only one, of Bloggs and Slade has rationally consented to being sacrificed to save the five, Jane’s evidence leaves it open which is which.

If Jane is morally conscientious, she’ll neither kill Bloggs nor Slade in SURGEON. But suppose she has a momentary lapse in moral conscientiousness and does chop one of them—Bloggs, say—up. Afterward, when she regains her moral conscientiousness, she, like our lately-lapsed Jill, will have the thought “I ought not to have done that; it was too risky”. But risky how? It can’t have been risky with respect to harm for the amount of harm that she did in killing Bloggs is the same as that which she would have done had she killed Slade instead.³⁰ The natural thing to say, and the thing subjectivists cannot say, is that Jane’s chopping up Bloggs was too risky because in doing so she ran the risk of horribly wrongfully killing someone.³¹

Perhaps the subjectivist will say that the risk Jane runs in chopping up Bloggs is the risk of killing someone who has not consented to being killed. But the following shows that that can’t be right:

SURGEON 2: Everything is as it is in SURGEON except that Jane knows that one of Bloggs and Slade has rationally consented to being sacrificed to save the five organ failure patients and the other is the one who villainously poisoned them, causing their organ failure. Jane’s evidence leaves it open which is which.

In this case Jane can retain her morally conscientiousness while killing Bloggs. (Intuitively, both Bloggs and Slade have waived their rights against being sacrificed to save the five: one has voluntarily relinquished it via his rational consent and the other has forfeited it by villainously poisoning the five organ failure patients.)³² Furthermore, if she does so, afterward she will neither have the thought that she ought not to have done so nor the thought that doing so was too risky. However, if she does kill Bloggs in SURGEON 2, she runs the very same risk of killing

someone who did not consent to being killed as she does in killing Bloggs in SURGEON. And it obviously won't do for the subjectivist to suggest that running the risk of killing-someone-who-hasn't-consented-to-being-killed-and-who-was-not-villainously-responsible-for-the-organ-failures-of-the-five-patients by killing Bloggs is what explains her thought that she ought not to have killed Bloggs in SURGEON: for one thing, it's simply implausible that a risk of something so hodgepoddgedly cobbled together could be the explanation, and, for another, yet other variants of SURGEON could be constructed to definitively show that it isn't.

Whatever maneuvers a subjectivist might try, in the end, she won't be able to shake the intuitive thought that the reason why lapsing Jill ought not to have prescribed A is that in doing so she took too high a risk of doing something seriously morally wrong. And if that's right, then, given that she interprets the 'ought' in Jill's thought as that identified with moral obligation, the subjectivist can't explain Jill's thought that prescribing A was too risky. The objectivist, on the other hand, can do this with ease. As, according to the objectivist, the 'ought' in Jill's thought that she ought not to have prescribed A is the subjective pragmatic 'ought', there's no problem in explaining why Jill ought_{pragmatic (subjective)} not to have prescribed A by appeal to the fact that her doing so ran too high a risk of her doing what she was very seriously morally obliged not to have done. Insofar as her prescribing A ran too high a risk of seriously falling short with respect to one of her goals as a morally conscientious person, viz., that of not doing what she is morally obliged not to do, doing so was not what she ought_{pragmatic (subjective)} to have done.

So, far from aiding the subjectivist's cause, an appeal to Jill's thought upon recovering her moral conscientiousness that she ought not to have prescribed A actually bolsters the objectivist's case.

6. An Argument For Objectivism Reconsidered

Now consider, once again, the second of the two arguments for objectivism offered in section 2. That argument, recall, went something like this: the question I want answered when I ask myself what my moral obligations are is the same as the question I want answered when, seeking your help, I ask you what my moral obligations are; intuitively, your adequately answering me doesn't require your knowing anything about my evidence; so, objectivism is true. But now consider the following. In the midst of my deliberations, instead of asking you what my moral obligations are,

I might just simply ask you what I ought to do in that situation. What's more, you might think, that's the more natural question for me to ask. But if that's the question I would ask, does the argument for objectivism lose its bite? No, but to see why will take some explaining.

Suppose, then, that in the midst of my deliberations about what to do, I ask you not what my moral obligations are, but what I ought to do. How should we understand what I'm asking here? Here's how. When I ask you what I ought to do in this saliently moral situation, I am asking the same question as I am asking when I ask you what I ought to do in any patently non-moral situation. That is, I am asking you what I ought_{pragmatic (objective)} to do in my situation. That the situation is a saliently moral one may be relevant to the goals against which the objective pragmatic 'ought' in that question is to be relativized, but that 'ought' is not the moral 'ought'; it is, indeed, the objective pragmatic 'ought'.

That the 'ought' I use when I ask you, quite generally, what I ought to do is a pragmatic 'ought' is obvious from the fact that if you don't take my goals and preference structure over them as the basis for the determination of your answer, you won't count as having answered my question. For example, suppose it is common knowledge that my goal is to make it to the airport on time to catch my flight and I ask you "which road ought I to take?". If you say road A, a road leading away from the airport and toward a post office, because you think I am morally obliged to mail off to Oxfam my entire life savings, then you're not cooperating and haven't answered my question. Similarly, suppose it is common knowledge (1) that I want both to get to the airport and to buy a toy for my niece, and (2) that, though I want to do both, I vastly prefer making my flight to buying the toy, and I ask you "which road ought I to take?". If you say road C instead of D, which you know leads to the airport, because doing so will allow me to get the toy for my niece, though at the cost of missing my flight, and you think it's better that I buy the toy than make my flight, then you're not cooperating and haven't answered my question. So, the 'ought' in my question, when I ask you, quite generally, what I ought to do, is clearly a pragmatic one.

It is also the *objective* pragmatic 'ought'. Why? From my perspective, the optimal choice among the options from which I am choosing is that which achieves the combination of goal satisfaction highest in my preference ranking over my goals. It's analytic that what I ought_{pragmatic (objective)} to do in my situation is pick the option with that property, whichever it happens to be. And as that is the optimal choice, clearly that is what I am asking you to identify for me if you can. To see this, consider also that in answering my question, though you do need to adopt my

goals and my preference ordering over them, you do not need to take into account my evidence about my situation.

Now, in a context in which I am facing a difficult moral decision, when I ask you what I ought to do, I am asking you what I ought_{pragmatic (objective)} to do in my situation. But as the situation is saliently a moral situation, the presupposition of my question (a presupposition which can, of course, be cancelled) is that my goals in action are those of any morally conscientious person. You will only count as answering my question, then, if you take those goals and a morally conscientious person's preference ordering over them as the basis of your answer. So, you will correctly answer my question only if you indicate which option does, in fact, lead to the outcome I, as a morally conscientious person, prefer most. But since, as a morally conscientious person, I have but the two goals of doing what I am morally obliged to do and of not doing what I am morally obliged not to do, the only way to correctly answer my question will be to say that I ought to take the option that I am morally obliged to take for, as a morally conscientious person, I must prefer that outcome to any other.

The argument for objectivism can be recast without loss of force, then, even if we take it that the question I ask myself, and perhaps a third-party, in deliberating about what to do is not "what am I morally obliged to do?", but, instead, "what ought I to do?"

A question remains, however. What if the third party I ask what I ought to do in my situation both can't simply share his evidence with me and, like me, is such that his evidence leaves it open what the outcomes of my different options will be? If so, then the third party is not in a position to say what I ought_{pragmatic (objective)} to do. The best he can do is tell me what I ought_{pragmatic (subjective)} to do, where the evidence about my situation against which that is determined is his evidence, not mine. What's more, this is what I want him to do, *if, in fact, his evidence about my situation is better than mine is*. (The reason being that if your evidence about my situation is better than mine is, the course of action you advise is more likely to lead to an outcome higher up on my preference ordering over my goals than is the course of action I would choose.) It is clear, then, what cooperating with my request requires when the third party is unsure about the outcomes of my different options. He must first preface his answer by saying that he does not know what I ought_{pragmatic (objective)} to do and can, at best, only say what I ought_{pragmatic (subjective)} to do, given my goals and his evidence about my situation; he can do this by saying some such thing as: "I'm not really sure what you ought to do in your situation

because I'm not sure what the outcomes of your different choices will be...". If he is sure that his evidence about my situation is better than mine is, he should then go on to say what I ought_{pragmatic (subjective)} to do in my situation relative to my preference ordering over my goals and his evidence about my situation. If he knows that his evidence about my situation is worse than mine is, he must demur, saying something such as "you're in a better position to know what's going to happen if you do this or that, so I shouldn't offer an opinion". And, finally, if he is unsure whether his evidence is better or worse than mine is, he must either simply state this or, at least, qualify his answer with it. If the third party acts in any other manner, he is not cooperating and isn't responding in good faith to my request.

Third-party inquiry about what one ought to do in a situation is a complicated business. But if we think straight about it, it is clear that this argument for objectivism survives being recast in terms of 'ought' instead of 'moral obligation'.

7. The Isn't-This-Whole-Dispute-Just-Verbal? Reply

One might have the following reaction to everything I've argued so far: "I see both how (II) is false on your construal of moral conscientiousness and how adopting that construal thus provides the objectivist a way of responding to the subjectivist's argument. In the previous two sections, however, you've brought this other word, 'ought', into the conversation and you've also admitted that it has a number of distinct senses. And though you claim that the *moral* sense of 'ought' does not admit of a subjective sense, you have granted that there is at least a subjective sense of what you've called "the pragmatic 'ought'". What's more, you've even granted that there are subjective pragmatic 'ought's associated with the goals and preference ranking over them of the morally conscientious person. But what you call the 'subjective pragmatic 'ought' relativized to the goals and preference ranking over them of the morally conscientious person' *just is* the one and only 'ought' that subjectivists were ever claiming is subjective. You can call it a pragmatic 'ought' and claim that there is no sense of 'moral obligation' associated with it if you like. But I don't see why a subjectivist needs to go along with any of this. I don't see why she can't say that her sense of 'ought' is a moral 'ought' associated with a sense of 'moral obligation'. And so, given that you can't really prevent her from doing this, you will, in the end, have to accept that there is a subjective sense of 'moral obligation'."

Even if a subjectivist retrenches so far as my hypothetical interlocutor indicates she may, I don't think I do need to grant that there is a subjective sense of 'moral obligation' for the reasons I've previously laid out. But even if I do accede to the above line of reasoning and grant the retrenched subjectivist her sense of 'moral obligation', I will, nonetheless, have secured a significant victory for the objectivist. Why? Well, note first that any subjectivist who presses the above line against me herself grants the existence of an objective sense of 'moral obligation' and must, *ipso facto*, back down from her claim that her DOCTOR-based argument refutes objectivism. And not only must she grant that there is such an objective sense of 'moral obligation', but she must also grant that, given my story of moral conscientiousness—one that is quite intuitively plausible and one which this hypothetical reply in no way discredits or calls into question—this objective sense of 'moral obligation' is one which is, as I've explained, of special concern to the morally conscientious person when she deliberates about what to do.

Second, and more important, though the subjectivist may not call the 'ought' she claims corresponds to her subjective moral obligation a pragmatic 'ought', I've shown how it behaves very much like all other pragmatic 'ought's. That, I take it, is good grounds for viewing it as a species of pragmatic 'ought's more generally. And given that pragmatic 'ought's more generally—'ought's associated with means and ends—are not plausibly viewed as corresponding to any general notion of obligation (except perhaps a notion of an obligation of rationality, though here talk of obligation is, in all truth, rather strained), insofar as this 'ought' is a species of pragmatic 'ought's more generally it is not intuitively a good candidate for being a moral 'ought' associated with a notion of moral obligation.

Third, and most important, if a subjectivist wants to take this line, she will have to grant that the objective sense of 'moral obligation' has, in a very important sense, primacy over her subjective notion. This is because her subjective notion of 'moral obligation' is explicable in terms of the objective notion: it is explicable in terms of the goals and preference ranking over them of the morally conscientious person, where the morally conscientious person's goals are understood to be both doing what she is objectively morally obliged to do and not doing what she is objectively morally obliged not to do. It is not at all clear, however, how the notion of objective moral obligation and wrongness here articulated and defended can be explicated in terms of this subjectivist notion of moral obligation.³³

One thing that must be conceded, however, is that my reply to the subjectivist's argument rests heavily on the plausibility of my account of moral conscientiousness. That account does seem intuitively quite plausible, but, it must be granted, if it is somehow defective, then my reply to the subjectivist's argument collapses. Its intuitive plausibility aside, however, there is more to be said in its favor. Not only does my account of moral conscientiousness offer the objectivist a response to the subjectivist's strongest argument for subjectivism, but also, an appreciation of it can help to make some headway in another contentious debate in normative ethics, namely that between possibilists and actualists. This further feature of my account of moral conscientiousness, which I discuss in the next section, provides additional reason to think that that account is true.³⁴

8. A Defense of Possibilism

In addition to the objectivism/subjectivism debate, there is also a debate in normative ethics about whether a person's own future voluntary actions are relevant to her moral obligations at a particular time. Actualists maintain that they are and possibilists maintain that they are not.

Consider:

HEADACHE: Aidan has a splitting headache. Mike, can either do nothing or give Aidan drug D. If Mike does nothing, after five excruciating hours, the headache will go away on its own. If Mike gives Aidan drug D, the headache will go away immediately. If, after administering drug D, however, Mike does not also administer drug E to Aidan later in the evening, Aidan will die. Mike can administer drug D now, but he knows that if he does so, even though he will be able to administer drug E later, because of his own laziness, he won't.³⁵

Actualists claim that Mike is morally obliged to do nothing. Because he won't in fact administer E later if he administers D now and his administering D now and not administering E later would lead to Aidan's death, they reason, shows that he has a moral obligation not to administer D now. Possibilists, on the other hand, think that Mike is morally obliged to administer D now. That he won't, because of his own future laziness, administer E later, doesn't make it the case that he isn't morally obliged to administer D now.³⁶

How, you might ask, can a possibilist maintain that Mike is morally obliged to administer D in HEADACHE if his doing so will lead, and he knows it will lead, to Aidan's death? The first thing a possibilist will say here is that it is slightly misleading to say that if Mike administers D

now, his doing so will lead to Aidan's death, for if Mike does so, Aidan will die not simply because Mike administers D, but because he administers it *and then fails to administer E later*. What's more, it is only because of Mike's own future laziness, and not any inability on his part, that if he administers D now he won't later administer E. Here, the possibilist contends, lies the fatal flaw of actualism.³⁷ According to actualism, a person's own voluntary and avoidable future moral failings can get her out of having certain present moral obligations. But it just doesn't seem that morality works that way. One can't get off the moral hook so easily and in virtue of one's own potential future moral badness. As actualism entails that one can, so much the worse for actualism.

The actualist might reply that, though he is not morally obliged to administer D now, Mike is morally obliged to perform the temporally diachronic complex action *administer D now and later administer E*. If so, actualism may be able to dodge the charge that it allows people to get off the moral hook in virtue of their own future moral failings. But if the actualist goes this route to salvage her theory, she must deny the following seemingly-platitudinous principle:

'And' Elimination for Moral Obligation (AEMO): Necessarily, if *S* is morally obliged to both ϕ and ψ and *S* can ϕ , then *S* is morally obliged to ϕ .

As Mike can perform the temporally diachronic complex action *administer D now and later administer E* and he is morally obliged to do the best that he can for Aidan, it seems that everyone who admits of the existence of temporally diachronic complex actions and moral obligations with respect to them, the actualist included, must grant that Mike is morally obliged to administer D now and later administer E in HEADACHE. And as Mike can also administer D now in HEADACHE, the actualist must grant that the antecedent of **AEMO** is true in HEADACHE. But, as the actualist denies that Mike is morally obliged to administer D now, she must deny that the consequent of **AEMO** is true in HEADACHE. So, because of the possibility of situations like HEADACHE, actualists must deny **AEMO**.³⁸ But that, contends the possibilist, is wildly implausible.

Actualism seems untenable. Nonetheless, it persists. It persists because possibilism seems untenable in its own right. The possibilist says that Mike is morally obliged to prescribe D in HEADACHE; but how could Mike in good conscience do that given that he knows that he won't administer E later? If he had that moral obligation, he would be morally obliged to do something that he knew would, in the end, lead to Aidan's death. That, according to the actualist, is simply

too much. Because possibilism has it that morality requires people to do what they know will end up having disastrous consequences, the actualist contends, it must be false.

At this point in the dialectic, it may seem that we have arrived at stalemate. Actualism and possibilism both seem untenable, but one of them must be true.

Here, however, the possibilist can blunt the actualist's criticism by appeal to the account of moral conscientiousness defended above. The possibilist can maintain that though, in some circumstances, morality might require people to do things they know will lead to disastrous consequences, it might very well be that no morally conscientious person would do as morality requires in those circumstances. Recall that, on my construal of moral conscientiousness, a morally conscientious person sometimes will do what she knows she's obliged not to do if not doing it would run a substantial risk of doing something she's much more seriously morally obliged not to do.

So, if she adopts my construal of moral conscientiousness, the possibilist can maintain both that morality requires Mike to administer D, and that if he is morally conscientious he won't. The reason Mike will not, insofar as he is morally conscientious, administer D is that, by stipulation, he knows that doing so would result, ultimately, in his doing something horribly morally wrong. True, if he does not administer D, he acts wrongly by the possibilist's lights, but if he does not administer D he will act *much less* wrongly, overall, than he will if he does administer D, for if he administers D now he will not administer E later, and not administering D now and doing anything else later, though wrong, is much less seriously morally wrong than is administering D now and not administering E later. Thus, as administering D now, given that he knows that if he does so he won't later administer E, would demonstrate a level of regard for the goal of not doing what you are morally obliged not to do inconsistent with moral conscientiousness, if Mike is morally conscientious, he will not administer D in HEADACHE.³⁹ And this can be true even if he nonetheless is, and knows he is, morally obliged to administer it. So, the possibilist can agree with the actualist that if Mike is morally conscientious, he won't administer D, while still affirming that he is morally obliged to do so.

There are significant problems for actualism and this independently compelling account of moral conscientiousness offers the possibilist a plausible reply to the actualist's strongest argument against possibilism. Possibilism, therefore, is preferable to actualism.

9. Conclusion

I claim that moral obligation is objective. I've offered two arguments for this view. First, the fact that when we change our view about what we are morally obliged to do after acquiring more evidence about our situation we view our information-deprived past selves as having been mistaken about what our moral obligations were seems to indicate that moral obligation is objective. And second, reflection upon third-party inquiry about what an agent is morally obliged, or morally ought, to do also seems to indicate that moral obligation is objective.

I've also defanged what I take to be the strongest argument for subjectivism. I've shown that that argument crucially depends on a mistaken view of moral conscientiousness according to which it is impossible for a morally conscientious person ever to do what she knows at the time of her doing it that she is morally obliged not to do. Once we get a clearer grip on what the goals of a morally conscientious person are, we see that under certain special circumstances—the very kinds of circumstances at play in the cases around which proponents of subjectivism build their arguments for subjectivism—a morally conscientious person may, and, sometimes, indeed must, insofar as she is morally conscientious, do something she knows she's morally obliged not to do. When the mistaken view of moral conscientiousness is swept aside, the argument for subjectivism falls away.

Last, I've demonstrated how the account of moral conscientiousness I favor offers proponents of possibilism a plausible way of replying to the actualist's strongest argument against it. In light of all of this, one should think that the only moral theories that have a shot at being true are the objectivist possibilist ones.

NOTES

* I am grateful to the participants in Liz Harman's Spring 2009 graduate seminar at Princeton University and an audience at Dartmouth College in Fall 2009 for comments on and discussion of earlier drafts of this article. I am also very grateful to David Enoch, Fred Feldman, Liz Harman, Chris Meacham, Michael Otsuka, David Owens, Brad Skow, Holly Smith, and Michael Zimmerman for very helpful comments on earlier drafts of this article. And, finally, special thanks to two reviewers and several editors at *Ethics* for a number of comments and suggestions that improved the article significantly.

¹ There is an entirely different debate in metaethics between camps that sometimes go by the names 'subjectivists' and 'objectivists'. I will have nothing to say about this completely separate debate.

² By 'the totality of a person's evidence at a time', here, I mean what internalists in epistemology would consider to be the totality of a person's evidence at that time. On this conception of evidence a person and her brain-in-a-vat-duplicate have exactly the same totality of evidence.

³ The '(or prior to)' qualification is added in order to accommodate subjectivists who deny the possibility of one's getting out of a future moral obligation by purposefully acting in ways so as to impoverish one's future evidential situation. Some subjectivists don't deny this possibility. The definition in the text is meant to cover both types of subjectivists.

⁴ With the notion of a person's *doxastic situation* (and a corresponding notion of her objective situation), it would be easy to define a property of being belief-subjective (and a corresponding property of being objective) exactly parallel to the definition of that of being evidence-subjective. Everything I say in the sequel about evidence-subjectivism holds *mutatis mutandis* for belief-subjectivism as well. Henry Sidgwick, for example, seems to endorse belief-subjectivism in H. Sidgwick, *The Methods of Ethics* (London: Macmillan and Company, 1907). Examples of evidence-subjectivism include H. A. Prichard, "Duty and Ignorance of Fact," in *Moral Writings*, ed. J. MacAidan (Oxford: Oxford University Press, 2002): 84-101; W. D. Ross, *Foundations of Ethics* (Oxford, Oxford University Press, 1939); F. Jackson, "Decision-theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101 (1991): 461-482; T. M. Scanlon, "Thomson on Self-defense," in *Fact and Value: Essays on Ethics and Metaphysics for Judith Jarvis Thomson*, ed. A. Byrne, R. Stalnaker, and R. Wedgwood (Cambridge: MIT Press, 2001): 199-213; M. Zimmerman, "Is Moral Obligation Objective or Subjective?" *Utilitas* 18 (2006): 329-361; M. Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008). Examples of objectivism include G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903); G. E. Moore, *Ethics* (Oxford: Oxford University Press, 1912); W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 2002); F. Feldman, *Doing the Best We Can* (Dordrecht, Holland: D. Reidel Publishing Company, 1986); J. Thomson, "Imposing Risks," in *Rights, Restitution, and Risk*, ed. W. Parent (Cambridge: Harvard University Press, 1986): 173-191; J. Thomson, *The Realm of Rights* (Cambridge: Harvard University Press, 1990).

⁵ According to this definition, of course, objectivism is consistent with belief-subjectivism. However, as I note in the text, everything I shall go on to say both against evidence-subjectivism and in favor of objectivism, understood as the denial of evidence-subjectivism, will hold *mutatis mutandis* with respect to belief-subjectivism and a corresponding notion of objectivism, understood as the denial of belief-subjectivism, as well. In the end, then, my arguments, taken all together, will constitute a defense of a narrower notion of objectivism, one understood as the

denial of the disjunction of evidence-subjectivism and belief-subjectivism. I grant that this narrower notion of objectivism is, perhaps, the most natural one. But, as I say, my arguments, taken all together, do constitute a defense of this narrower notion. I choose, in the text, to construe objectivism as the denial of evidence-subjectivism merely for ease of exposition.

⁶ Though many subjectivists are committed to **OIC**, few discuss this commitment in connection with their subjectivism. Sidgwick, Ross (in *Fundamentals of Ethics*), and Zimmerman all commit themselves both to **OIC** and to subjectivism in their writings.

⁷ Even this is perhaps not quite right. Some subjectivists also probably hold that certain agents, such as children, don't have moral obligations at all, and so such subjectivists would probably prefer some further restricted version of ability-constrained-evidence-subjectivism. I thank an anonymous referee for pointing this out to me. Concerns similar to those I raise below in the text about the consistency of an ability-restriction with the standard intuitive motivation for subjectivism may well also bedevil any such further restrictions.

⁸ In *Living with Uncertainty*, however, Zimmerman explicitly denies that this is his motivation for the kind of subjectivism he favors.

⁹ Ross originally advocated objectivism in *The Right and the Good*, but later came to endorse subjectivism in *Foundations of Ethics*.

¹⁰ Ross, *The Right and the Good*, 32.

¹¹ This argument cannot be side-stepped by pointing out that your telling me that I'm morally obliged to turn on the light changes my evidential situation in such a way that 2 is no longer plausible. If we added to PROMISE that I have overwhelming, though, unbeknownst to me, misleading, evidence that you are an inveterate liar, both premise 1 and premise 2 would, nonetheless, still seem true.

¹² Both in "Is Moral Obligation Objective or Subjective?" 344-345 and in *Living with Uncertainty*, 32.

¹³ Just a small point about the dialectic before proceeding. In this and the following section, I will not be arguing for objectivism. Rather, I'll only be offering what I take to be a plausible reply to the subjectivist's arguments for subjectivism. So, in responding to the subjectivist's first argument, I'll appeal to what seem to be intuitive grounds for rejecting the kind of straightforward connection between moral obligation and blameworthiness upon which that argument implicitly relies. A subjectivist may deny these intuitive grounds. That's fine. My aim is not to convert the staunch subjectivist—that would be too much to hope for; my aim, rather, is merely to provide a coherent reply to the subjectivist's arguments that is intuitively plausible, the elements of which can be motivated independent of the need to respond to those arguments.

¹⁴ It is true, strictly speaking, that the purported data to which subjectivists appeal in this SYRINGES-based argument (and also in the DOCTOR-based argument I discuss later) are consistent with objectivism as I've defined it. No matter; all plausible versions of objectivism deny these data, and most, if not all, objectivists would grant that were these purported data true, subjectivism would win the day.

¹⁵ It is, to be sure, not widely recognized that it is possible for there to be cases in which a person is blameworthy even though (she knows) she does not act wrongly. Zimmerman is one who does actually recognize it; see his M. Zimmerman, *An Essay on Moral Responsibility* (Totowa, NJ: Rowman and Littlefield, 1988).

¹⁶ Or, more cautiously, *not all* subjectivists and objectivists seem to be arguing past each other.

¹⁷ Zimmerman was the first, as far as I know, to appeal in this way to the notion of the morally conscientious person in her deliberations about what to do as a way of pinning down the sense of ‘moral obligation’ about which subjectivists and objectivists disagree. See “Is Moral Obligation Objective or Subjective?” 335 and *Living with Uncertainty*, 15.

¹⁸ This is my presentation of the case, but, in detail, it does not differ from that offered in “Decision-theoretic Consequentialism,” 462-464.

¹⁹ Donald Regan, as far as I can tell, was the first to employ a case with this kind of structure. See D. Regan, *Utilitarianism and Cooperation* (Oxford: Oxford University Press, 1980): 265 n.1.

²⁰ I have cast the objectivism/subjectivism debate as one over the nature of moral obligation, but, strictly speaking, all that Jackson claims is obvious is that Jill “ought” to prescribe B in DOCTOR. For the purposes of my discussion here, I interpret Jackson’s ‘ought’ as the moral ‘ought’ associated with moral obligation. See section 5 for a discussion of ‘ought’ and its relation to moral obligation.

²¹ (III) is consistent with its being the case, as Sidgwick and other utilitarians have argued, that there may very well be utilitarian grounds for suppressing the truth of utilitarianism. It doesn’t follow (at least not obviously) from the fact that people would perform actions more in accord with the principle of utility if they didn’t believe it than if they did that if one did believe it one would thereby not be morally conscientious.

²² Though I find the objectivist reply to the subjectivist’s argument that I go on to offer below quite plausible, I do not mean to suggest that it is the only plausible reply to the subjectivist’s argument that an objectivist might give. There may be other (and, perhaps, even better) objectivist replies. I aim only to present, and tout the particular merits of, my own reply.

²³ This conception of the morally conscientious person’s concern with her moral obligations being two-faced, I take to be similar, and somewhat analogous, to William James’s contention that, as would-be knowers, our concern with truth is two-faced: “*We must know the truth; and we must avoid error*—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws.” W. James, “The Will to Believe,” in his *Essays in Pragmatism* (New York: Haffner, 1969): 113.

²⁴ Mightn’t someone be morally obliged to do more than one thing in a situation? Well, in one sense, of course, yes; but, in the most important sense, no. You might, for instance, be morally obliged to open the door and also to say ‘Hello’. But, in such a case, given that temporally synchronic complex actions are actions (which they most certainly are), there is only one action, namely the complex action of opening the door and saying ‘Hello’, such that the performance of it is sufficient for your not acting morally wrongly.

²⁵ Here I say that the morally conscientious person may be required in certain situations to make certain tradeoffs among her two goals. One might think that in claiming that the morally conscientious person is *required* to do this I am here implicitly invoking another notion of obligation or requirement that a subjectivist might insist is the very notion she was ever claiming was subjective. This is not correct. The notion of requirement I am here invoking is merely that of conceptual necessity. I am claiming that, as a conceptual matter, one simply won’t count as being morally conscientious if she does not make these kinds of tradeoffs in these kinds of situations. This notion of requirement is in no way like that of a moral requirement. On this notion of requirement a rapist is conceptually required to commit rape—a person, as a matter of conceptual necessity, won’t count as a rapist unless she commits rape, and so, a rapist, insofar as

she is a rapist, is conceptually required to commit rape—but there is no notion of obligation, let alone moral obligation, according to which a rapist is obliged to commit rape.

²⁶ (II') and (VI), though incontrovertibly true in cases in which the morally conscientious person is choosing between options on which her own future voluntary actions have no bearing, are subject to qualification with respect to cases in which the morally conscientious person has knowledge of her own potential future moral wrongdoing. I discuss such cases in section 8.

²⁷ Throughout my discussion I have been presupposing that in DOCTOR Jill has but three options available to her: *prescribe A*, *prescribe B*, and *prescribe C*. I have presupposed, that is, that Jill does not, in addition, have yet other disjunctive options, such as *prescribe-A-or-prescribe-C*, *prescribe-B-or-prescribe-C*, etc. My claim that (VI) is true, however, may seem false on the supposition that Jill has not only the three non-disjunctive options I've assumed she has, but also, in addition, a host of other disjunctive options. This is because, one might argue, in DOCTOR, even though Jill is morally obliged to *prescribe-A-or-prescribe-C*—which, it might seem, follows from its being the case that she either has a moral obligation to prescribe A or has a moral obligation to prescribe C—and is in a position to know this, insofar as she is morally conscientious, she will nonetheless prescribe B. For the purposes of my discussion, however, I assume that we needn't worry about any complexities that may arise due to the postulation of disjunctive options. That is, I assume that the only options that the principles (II') and (VI) quantify over are those of the normal non-disjunctive kind. I thank two anonymous referees for impressing upon me the need to restrict these principles so as to exclude disjunctive options.

²⁸ Being morally conscientious is, of course, not only a matter of having the two goals that a morally conscientious person has, but also a matter of having the kind of preference ranking over them that a morally conscientious person has. It is also quite plausible, nor do I deny, that there is a range of different preference rankings over the goals of the morally conscientious person consistent with moral conscientiousness. All I claim is that there are some such preference rankings that are definitively inconsistent with moral conscientiousness. There being such a range does, of course, entail there being many cases in which, even though there is an option that is morally obligatory, there simply is no fact of the matter as to what a morally conscientious person would do in that case. But that is overwhelmingly plausible; there *are* many such cases.

²⁹ Zimmerman explicitly acknowledges that subjectivists are committed to its being conceptually incoherent for a person to be morally obliged not to do something in virtue of its riskiness with respect to wrongdoing (*Living with Uncertainty*, 58). He does not, however, see, or feel the force of, the problem for subjectivism I suggest that this poses.

³⁰ Nor can it be risky with respect to killing, obviously, for the amount of killing she does will be the same whether she kills Bloggs or Slade.

³¹ SURGEON also puts paid to the suggestion that the intuitive riskiness thought can be explained in terms of riskiness with respect to badness of outcomes (or loss of value). The state of affairs in which Jane chops up the one who has consented and the one in which she chops up the one who hasn't don't differ with respect to value and both are better in terms of value than that of her not chopping up either, but chopping up either of them is still intuitively risky. It is precisely the non-consequentialist's point, after all, that if given the choice between sacrificing one who has rationally consented to being sacrificed in order to save five others and sacrificing one who hasn't so consented, one is morally obliged to sacrifice the one who has consented, if one sacrifices anyone at all, even though sacrificing either would lead to an equally good state of affairs. What's more, even if a non-consequentialist were to hold that facts about consent were

not only relevant to moral permissibility but also to the values of states of affairs, still other versions of SURGEON, ones no non-consequentialist could view as involving a risk with respect to badness, would establish beyond a shadow of doubt that the intuitive riskiness is not a riskiness with respect to the badness of outcomes.

³² Here I am prescinding from questions of legal permissibility. Of course it might be legally impermissible to redistribute the organs of a villainous poisoner among his victims in order to save them. But doing so would surely not be morally impermissible.

³³ A straightforward counterfactual translation certainly won't work. It won't do to say that what a person is objectively morally obliged to do in a situation just is what she would be subjectively morally obliged to do in that situation were her evidence completely in line with the facts in that situation. This is because the changes necessary to make the situation one in which the agent's evidence was in line with the facts might be such as to make it the case that she was no longer objectively morally obliged to do just what she was in the absence of those changes.

³⁴ Special thanks to David Owens for encouraging me to think about cases like the one I discuss in the next section in connection with my account of moral conscientiousness.

³⁵ A case of this form was originally discussed in the context of the actualism/possibilism debate in H. Goldman, "Dated Rightness and Moral Imperfection," *Philosophical Review* 85 (1976): 449-487. Other cases with similar structures are discussed in H. Sobel, "Utilitarianism and Past and Future Mistakes," *Noûs* 10 (1976): 195-219; H. Goldman, "Doing the Best One Can," in *Values and Morals*, ed. A. Goldman and J. Kim (Dordrecht, Holland: D. Reidel Publishing Company, 1978): 185-214; F. Jackson and R. Pargetter, "Oughts, Options, and Actualism," *Philosophical Review* 95 (1986): 233-255; L. Goble, "The Logic of Obligation, 'Better', and 'Worse'," *Philosophical Studies* 70 (1993): 133-163; M. Zimmerman, *The Concept of Moral Obligation* (Cambridge: Cambridge University Press, 1996).

³⁶ Goldman (in "Dated Rightness and Moral Imperfection"), Sobel, Jackson and Pargetter, and Goble uphold actualism. Goldman (in "Doing the Best One Can"), Feldman, and Zimmerman all uphold possibilism.

³⁷ In the remainder of this and the following paragraph, I rehash fairly closely Zimmerman's discussion of some of the problems for actualism in *The Concept of Moral Obligation*, 191-192.

³⁸ Some actualists, e.g., Jackson and Pargetter, are quite forthright and sanguine about having to deny **AEMO**.

³⁹ This is the kind of case I had in mind when, in footnote 26, I claimed that (II') and (VI) are subject to qualification with respect to cases in which the morally conscientious person has knowledge of her own future potential wrongdoing.