

Subjective Deontology and the Duty to Gather Information*

Philip Swenson

Holly Smith has recently argued that Subjective Deontological Moral Theories (SDM theories) cannot adequately account for agents' duties to gather information. I defend SDM theories against this charge and argue that they can account for agents' duties to inform themselves. Along the way, I develop some principles governing how SDM theories, and deontological moral theories in general, should assign 'deontic value' or 'deontic weight' to particular actions.

I. INTRODUCTION

Agents sometimes have a duty to gather information. Agents who will be faced with a morally significant choice are often morally required to seek out new evidence that will aid them in making their choice. If a moral theory cannot yield the result that agents have such obligations in cases where it seems clear that they do, we have strong reason to reject it.

Holly Smith has argued that Subjective Deontological Moral Theories (SDM theories) cannot account for agents' duties to gather information in cases in which it seems clear that agents have such a duty.¹ In what follows I defend SDM theories against this charge and argue that they can account for agents' duties to inform themselves in the sort of case Smith is concerned with. Along the way, I develop some principles governing how SDM theories, and deontological moral theories in general, should assign 'deontic value' or 'deontic weight' to particular actions.

SDM theories are deontological moral theories which *subjectivize* moral duties in such a way that what an agent has a duty to do depends on her beliefs about the nature of her potential actions rather than on the

* I would like to thank Ben Bronner and two anonymous referees for very beneficial comments. Also, thanks to Holly Smith for insightful suggestions and helpful discussion.

1. Holly Smith, "The Subjective Moral Duty to Inform Oneself before Acting," *Ethics* 125 (2014): 11–38.

mind-independent nature of her actions.² For example, where an objective moral theory might hold that:

A ought to keep her promises;
or,
A ought to refrain from killing.

An SDM theory will instead hold that:

A ought to do what she believes will amount to keeping her promises;
or,
A ought to do what she believes will amount to refraining from killing.

SDM theories are attractive because they appear to be more “user friendly” than objective moral theories. It is much easier to accidentally fail to keep a promise than it is to accidentally fail to do what you believe amounts to failing to keep a promise. Insofar as we prefer moral theories that provide a high degree of guidance for the decision making of epistemically limited agents, Subjective Deontologists may seem to have an edge. However, Smith argues that SDM theories have a fatal flaw.

II. SMITH'S OBJECTION

In order to lay out Smith's objection, I need to get the notion of “deontic value” on the table.³ Prima facie moral duties can conflict. When they do, the more weighty or stringent duty wins out. The notion of deontic value nicely captures this feature of morality. As Smith puts it:

The deontic value of an act expresses the weight, or stringency, of the duty to perform (or not to perform) that act (or, in other terminology, the force of the moral reason to perform or not to perform that action). Thus the deontic value of saving a person's life is greater than the deontic value of keeping a minor promise. In determining what one ought all things considered to do, one weighs the deontic value of the various (sometimes conflicting) duties involved. If one has to choose between saving a life versus keeping a minor promise, one ought all things considered to save the life, since the deontic

2. Proponents of Subjective Deontological Moral Theories include H. A. Prichard, “Duty and Ignorance of Fact,” in *Moral Obligation and Duty and Interest* (Oxford: Oxford University Press, 1968), 18–39, and W. D. Ross, *Foundations of Ethics* (Oxford: Clarendon, 1939), chap. 7.

3. In her book manuscript “Making Morality Work” Holly Smith instead refers to this notion as “deontic weight.”

value of this act is greater than the deontic value of keeping the promise.⁴

I also need to introduce Smith's notion of a derivative duty. A duty is derivative if "the fact that it is a duty derives from the subsequent duties it would lead the agent to satisfy."⁵ For example, if I have a duty to pick you up at the airport at 9:30, and I would only succeed in fulfilling this duty if I left my home before 8:45, then it is plausible that I have a derivative duty to leave my home before 8:45. I have the duty to leave home before 8:45 because if I fail to do so I would not satisfy my duty to pick you up at 9:30.

Smith makes use of both of these notions to develop a plausible principle for objective deontological theories to appeal to in accounting for the duty to gather information:

OD: An agent has an objective derivative prima facie duty to acquire information if and only if doing so would lead the agent subsequently to produce the maximum possible amount of deontic value (typically through his carrying out the various deontic duties that would later be incumbent on him).⁶

One might suspect that SDM theories could make use of some very similar principle, and indeed Smith provides one:

SD: An agent has a subjective derivative prima facie duty to do what he believes is acquiring information if and only if he believes that doing what he believes is gathering information would lead him subsequently to produce the maximum amount of deontic value (typically through his doing what he then believes to be carrying out the various deontic duties that would then be incumbent on him).⁷

However, Smith argues that SDM theories cannot, even by appealing to SD, account for certain cases in which it is clear that agents have a duty to gather information. Smith's discussion relies on the following case:

Claire's Choice: Suppose Claire, the human resources manager tasked with laying off an employee, is governed by [a subjective moral code] which includes a prima facie deontic duty to do what the agent believes to be laying off employees justly. . . . Justice will be best served by doing what the agent believes to be laying off the least productive employee, but in a case of employees who are all tied for minimum productivity, or a case in which the manager does not know which

4. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," 20.

5. *Ibid.*, 20.

6. *Ibid.*, 21.

7. *Ibid.*, 24.

employee is least productive, the second-best solution called for is to lay off the person the agent believes to be the most recently hired employee.

In this case, Claire currently doesn't have any beliefs about which employee is least productive, although she (correctly) believes that one of them is more productive than the other. She already believes that Max is the most recently hired employee. Claire truly believes that if she does what she believes is gathering information she will come to accurately believe of one of the employees that he or she is least productive, whereas if she does what she believes to be omitting to gather information, she will continue not to believe of any employee that he or she is least productive. If she gathers information, she will have a subsequent duty . . . to do what she believes to be laying off the least productive employee, whereas if she omits to gather information, she will have a subsequent duty . . . to lay off Max, whom she believes to be the most recently hired employee . . . Claire truly believes that she will carry out whichever of these subsequent duties is incumbent upon her.⁸

Since "Justice will be best served by doing what the agent believes to be laying off the least productive employee," and Claire will do that if she gathers information, intuitively this is a case in which Claire ought to do what she believes is acquiring further information. SD will yield this result if we assign deontic values such as the following:

*Claire does what she believes to be laying off the least productive employee (on the condition that she has gathered information). [+2 deontic value.]

**Claire does what she believes to be laying off Max (on the condition that she has not gathered information). [+1 deontic value.]

So far, all is well. But Smith argues that assigning a higher deontic value to * than to ** is untenable.

Smith takes it that Claire's duty to do * (on the condition that she has gathered information) would be a weightier duty than Claire's duty to do ** (on the condition that she has not gathered information).⁹ Furthermore, Smith points out that Claire's deciding to gather information would *create* her later duty to do *. Alternatively, Claire's deciding to

8. Ibid., 24–26. Smith builds into the case that Claire's beliefs regarding the results of gathering information are true. It is not clear that she needs to do so. A Subjective Deontologist who accepts SD would presumably want to say that Claire has a duty to investigate whether or not her beliefs are true.

9. This is because it would be worse to refrain from doing * (on the condition that she has gathered information) than it would be to refrain from doing ** (on the condition that she has not gathered information).

refrain from gathering information would *create* her later duty to do **. (This fact differentiates SDM theories from objective moral theories which would assign Claire a duty to lay off the least productive employee whether or not she has gathered information.) The fact that Claire creates her later duty underwrites Smith's objection to SDM theories. Smith asks:

But is it really true that one should act so as to create a weightier duty for oneself which one would then carry out, as opposed to creating a less weighty duty for oneself which one would then carry out?¹⁰

Smith's answer is that it is not really true. She provides a persuasive argument against the claim that *one has a prima facie duty to "act so as to create a weightier duty for oneself which one would then carry out, as opposed to creating a less weighty duty for oneself which one would then carry out."* (For Smith, a weightier duty is a duty that it would be worse to violate.) Consider the following case:

Devon's Promise: Suppose Devon receives two e-mail messages: one from Kate, who asks Devon to take care of Kate's cat next week while she is on vacation, and one from Fred, who asks Devon to take care of his goldfish next week while he is on vacation. Both friends say that if Devon can't do it, they can find another equally competent caretaker. On the other hand, if he promises, then the pet will only be fed if Devon himself does it.¹¹

Smith points out that (assuming that someone else will care for each creature if he declines) it is implausible to say that Devon ought to promise to take care of Kate's cat rather than Fred's fish. The promise to care for the cat would be weightier,¹² but this seems to provide Devon with no reason to prefer it. Devon can permissibly create the less weighty obligation.

Drawing on her earlier work,¹³ Smith claims that the best way to avoid the implication that Devon ought to create the weightier duty is to take the view that

while breaking a promise has negative value, keeping a promise as such has no positive value—no value above and beyond what the promised act would have had if it had not been promised.¹⁴

10. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," 27.

11. *Ibid.*

12. This is because failing to keep the promise to care for the cat would be worse.

13. Holly Smith, "A Paradox of Promising," *Philosophical Review* 106 (1997): 153–96.

14. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," 28.

On this approach we would assign a value of zero (assuming no other relevant factors are in play) to both of the following:

KC: Devon's *keeping the promise* to take care of Kate's cat (on the condition that he made this promise) [0 deontic value.]

FF: Devon's *keeping the promise* to take care of Fred's fish (on the condition that he made this promise) [0 deontic value.]

Given this, Devon gains no additional deontic value by selecting KC. Thus, he has no obligation to do so. And this seems like the right result.

Smith thinks we should treat Claire's duty-creating choice the same way we treat Devon's. Since we assign a value of zero to both KC and FF, we should also assign a value of zero to both * and **. And if we assign a value of zero to both * and **, then gathering information and then laying off the least productive employee no longer uniquely maximizes deontic value. So SD will no longer license the conclusion that Claire ought to gather information. And SDM theories can no longer account for the fact that agents like Claire ought to gather information.

I take something like the following to be Smith's crucial claim:

Crucial Claim: Just as we should assign equal (more specifically zero) deontic value to fulfilling the more weighty potential obligation and the less weighty potential obligation in cases like Devon's Promise, so too we should assign equal (more specifically zero) deontic value to fulfilling the more weighty potential obligation and the less weighty potential obligation in cases like Claire's Choice.¹⁵

This is where I part ways with Smith. I think Subjective Deontologists can plausibly reject Crucial Claim. One could offer a flatfooted response to Smith and merely insist that Smith hasn't definitively shown that Crucial Claim is true. But it would be much better if one could make a positive case that Crucial Claim is false.

One might attempt to attack Crucial Claim by rejecting Smith's analysis of Devon's Promise. I suspect that some philosophers will not want to go along with Smith's claim that keeping a promise has no positive

15. The following remarks make it fairly clear that Smith endorses Crucial Claim: "Just as making a promise . . . create[s] new duties, so, within a subjective deontological code, acquiring morally relevant information creates a new duty. . . . These created duties must be handled in the same way that the duties created by making promises are: fulfilling such prima facie duties must be accorded no positive value, even though violating these duties is accorded negative value. . . . Thus Claire's doing what she believes is gathering information and her doing what she believes to be not gathering information would have the same value (0)" (ibid., 29).

value. But I find Smith's claims about promising quite plausible. Thus I will provide a response which is consistent with Smith's insights regarding Devon's Promise. I will argue that Crucial Claim is false because there is an important asymmetry between Claire's Choice and Devon's Promise.

III. SETTING UP MY RESPONSE

Suppose the following claims are both true as a matter of distributive justice:

- (a) Employers ought to do what they believe to be paying their employees a living wage if they can do so consistent with their other moral obligations.

And,

- (b) Employers who cannot do what they believe to be paying their employees a living wage (consistent with their other moral obligations) ought to do what they believe to be paying their employees as close to a living wage as possible.

Now consider two business owners:

Business owner 1: does what she believes to be paying her employees a living wage.

Business owner 2: cannot do what she believes to be paying her employees a living wage but does do what she believes to be paying her employees as close to a living wage as possible.

Both of the business owners are acting as they ought. However, there is an important sense in which Business owner 1 is doing better (with regard to the deontic norm of acting justly) than Business owner 2. Intuitively there is something *regrettable* about Business owner 2's actions. How should we make sense of this?

I suggest that deontic norms (or prima facie duties) are associated with *deontic ideals* which tell us what justice, or benevolence, or loyalty, and so forth would require in the absence of competing considerations or limitations in ability. And they are also associated with *approximation principles* which tell us how to best approximate deontic ideals if we cannot achieve them. So, for example, we might have:

Deontic ideal for distributive justice in employment: when operating a business do what you believe to be paying your employees a living wage.

Approximation principle for distributive justice in employment: when operating a business do what you believe to be paying your employees as close to a living wage as possible.

In my view we need to make use of something along the lines of deontic ideals to account for the intuitive reaction to Business owners 1 and 2. Business owner 2's behavior is regrettable (although permissible) because she does not achieve the relevant deontic ideal, while Business owner 1's behavior is not regrettable. Appealing to deontic ideals allows us to explain this asymmetry. Thus we have reason to add deontic ideals to our theoretical toolbox. I now want to suggest that there are important connections between deontic ideals and deontic value.

First Deontic Value Principle (FDVP): All things being equal, option *X* has higher deontic value than option *Y* if *X* approximates some *deontic ideal* to a greater degree than *Y* (where achieving a deontic ideal counts as maximally approximating it).

The FDVP seems to have a lot going for it. It is intuitively plausible in its own right. And it also nicely explains both why Business owner 1 is obligated to pay her employees a living wage and why Business owner 2 ought to pay her employees as close to a living wage as possible. Thus Subjective Deontologists should be quite pleased if the FDVP underwrites a response to Smith's objection.

It is an interesting question whether deontic ideals or deontic value should be thought of as explanatorily and/or conceptually prior. The FDVP is neutral on this question. And I think it is plausible whichever answer is given. I will speak of deontic ideals 'explaining' or 'accounting for' duties and deontic values. But it may be that deontic ideals depend on deontic values and thus reveal rather than explain them.¹⁶ I now want to examine how the FDVP might be applied to a more complex case. My examination of this case will illuminate the approach that I suggest Subjective Deontologists take toward Claire's Choice.

Consider the following case:

Bob's Business: Bob is deciding between starting business A or starting business B. If he starts business A he will be able to (and will) pay his employees a living wage. If he instead starts business B he will be unable to pay his employees a living wage (though he will pay them as close to a living wage as possible).

16. If this is so, then I will have to be satisfied with identifying principles which reveal deontic values vindicating Subjective Deontology. I will not have explained why the deontic values obtain. I regard this weaker result as satisfactory.

I'm inclined to view Bob's Business as a case in which (all other things being equal) Bob ought to create a duty to pay his employees a living wage.¹⁷ He ought to start business A. If we assume that what Bob ought to do is maximize deontic value, then a plausible explanation of the source of his obligation is that:

BA: paying his employees a living wage (having started business A)

has a higher deontic value than:

BB: paying his employees as close to a living wage as possible (having started business B).

Since on BA Bob achieves the deontic ideal for distributive justice in employment, the FDVP underwrites the claim that (all other things being equal) BA has higher deontic value than BB. On this approach Bob has a *derivative duty* to start Business A because doing so would allow him to fulfill his subsequent obligation to pay his employees a living wage (thus maximizing deontic value).

Although I find it plausible, I am not strongly committed to the claim that this is the correct way of making sense of Bob's Business. One could resist my analysis of Bob's Business by claiming that agents have a non-derivative duty to start businesses that can pay a living wage. (Thus rendering my appeal to a derivative duty superfluous.) I have presented this approach to Bob's Business because it is structurally parallel to the approach that I suggest Subjective Deontologists take in cases like Claire's Choice.

IV. ACCOUNTING FOR CLAIRE'S DUTY TO GATHER INFORMATION

Recall that on Smith's setup for Claire's Choice we are supposed to assume that:

Justice will be best served by doing what the agent believes to be laying off the least productive employee, but in a case of employees who are all tied for minimum productivity, or a case in which the manager does not know which employee is least productive, the second-best solution called for is to lay off the person the agent believes to be the most recently hired employee.¹⁸

17. Note that this may provide a counterexample to the strong claim that one is never obligated (absent reasons independent of deontic value) to create a more weighty duty rather than a less weighty duty.

18. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," 24.

Given this, the following deontic ideal and approximation principle seem plausible:

Subjective deontic ideal for justice in layoffs: do what you believe to be laying off the least productive employee.

Subjective approximation principle for justice in layoffs: (a) do what you believe to be laying off the most recently hired employee, (b) if (a) is not an option . . .

Now recall:

*Claire does what she believes to be laying off the least productive employee (on the condition that she has gathered information). [+2 deontic value.]

**Claire does what she believes to be laying off Max (on the condition that she has not gathered information). [+1 deontic value.]

Since on * Claire achieves the deontic ideal and on ** she does not, the FDVP tells us that we were correct to assign a higher deontic value to * than to **. And if that is right then we can use SD to account for Claire's derivative duty to gather information. Claire's gathering information will lead her to maximize deontic value. Thus, Subjective Deontologists can account for the duty to gather information (at least in cases like Claire's).

Now of course you are wondering why the same approach does not apply to Devon's Promise. Recall:

KC: Devon's *keeping the promise* to take care of Kate's cat (on the condition that he made this promise) [0 deontic value.]

and,

FF: Devon's *keeping the promise* to take care of Fred's fish (on the condition that he made this promise) [0 deontic value.]

If the FDVP required us to assign a higher deontic value to KC than to FF, something would have gone wrong. Devon would then (implausibly) be obligated to make the weightier promise. However I believe that we can avoid this result.

In my view the fact that Devon is not obligated to select KC tells us something about the nature of the deontic ideal for promise keeping. An agent that refrains from making a weighty promise is in no worse standing, with regard to the deontic ideal for promise keeping, than an agent who makes and keeps a weighty promise.

Recall that we can identify deontic ideals by examining whether certain acts are regrettable. We identified the deontic ideal for distributive justice in employment by noting that Business owner 2's actions are regrettable while Business owner 1's are not. Similarly, it seems that Claire's selecting ** would be regrettable, thus it is plausible that only * achieves the deontic ideal for justice in layoffs. But Devon's opting for FF does not appear regrettable at all. (This is why it is clear that he is not obligated to choose KC.) Thus the following deontic ideal looks plausible:

Subjective deontic ideal for promise keeping: do what you believe to be keeping all of the promises you have made.¹⁹

This ideal explains why Devon is not obligated to create the more weighty obligation (*keeping the promise to care for Kate's cat*). Creating and then fulfilling the less weighty obligation (*keeping the promise to care for Fred's fish*) also achieves the relevant deontic ideal. Thus the FDVP allows us to assign them equal deontic value.²⁰

Some may find my proposed asymmetry between the deontic ideal for promise keeping and the deontic ideal for justice in employment surprising. I want to emphasize that the asymmetry is motivated by a similar asymmetry in what we judge to be regrettable. Recall that Claire's selecting ** is regrettable while Devon's choosing FF is not. Thus I think the asymmetry is well supported.

Of course it does sometimes matter how weighty a promise is. I suggest that the relevant approximation principle will account for this:

Subjective approximation principle for promise keeping: when you cannot do what you believe to be keeping all of the promises you have made, prioritize doing what you believe to be keeping more weighty promises over doing what you believe to be keeping less weighty promises.

This explains why Devon ought to care for Kate's cat if he has promised to care for both Kate's cat and Fred's fish but cannot care for both. So

19. A referee has suggested that achieving the deontic ideal for promise keeping might require making and keeping a promise, rather than declining to make the promise, when some further good will result from making and keeping the promise. I am inclined to explain the badness of not making such a promise by appealing to the deontic ideal associated with achieving the further good, rather than by modifying the deontic ideal for promise keeping. But endorsing such a modification would not substantially alter my response to Smith. In Devon's Promise the same goods will come about whether Devon selects KC or FF, so the modification would not come into play.

20. This ideal also explains Smith's important insight that agents do not normally have a prima facie duty to create obligations via promising (see Smith, "A Paradox of Promising"). Agents can achieve the deontic ideal for promise keeping without making any promises.

the fact that the weightiness of a promise sometimes matters does not entail that 'the weightiness of promises' features in the deontic ideal for promise keeping. Weightiness matters only when the ideal will not be achieved.

I take myself to have made a plausible case that Smith's Crucial Claim is false. We can treat Claire's case differently from Devon's and account for her duty to gather information. I now want to consider whether my solution generalizes to another sort of case in which we judge that the agent has a duty to gather information.

V. DOES MY APPROACH GENERALIZE?

There is a worry that my approach will not generalize to the full range of cases in which agents have a duty to gather information. The FDVP gives us the right result in Claire's case because we were evaluating her options relative to only one deontic ideal (the deontic ideal for justice in layoffs). But there appear to be cases in which an agent will achieve one deontic ideal if she gathers information and another deontic ideal if she does not.²¹ Here is a case in which it is plausible that this occurs:

Daisy's Decision: Suppose Daisy is an agent governed by the *subjective deontic ideal for promise keeping* as well as the following deontic ideal

Subjective deontic ideal for parental obligation: when you believe your child is very ill, do what you believe to be acquiring the best available medical care for your child.

Suppose also that Daisy is currently in the dark about where to acquire the best care for her daughter who is suffering from a serious illness. She truly believes that if she (at t1) researches which medical clinics offer the best care, she will identify a clinic far from her town. In that case she truly believes that she would be obligated to do what she believes to be traveling to the clinic with her daughter at t2. But this travel would render her unable to keep a promise she has made to give a lecture at a local college at t2. Daisy also truly believes that if she did not (at t1) research which medical clinic is best then she would be unable to achieve the *subjective deontic ideal for parental obligation*, and thus permitted to seek treatment for her daughter at the local clinic. And she would then be obligated to do what she believes to be keeping her promise to give the lecture at t2. Furthermore, she truly believes that she will do whatever she is obligated to do at t2.

Intuitively, Daisy ought to gather information in this case. So we need to assign a higher deontic value to ST than to GL:

21. Thanks to Holly Smith for suggesting the possibility of such cases to me.

ST: do what she believes to be traveling to the clinic with her daughter (on the condition that she has gathered information)

and

GL: do what she believes to be keeping her promise to give the lecture (on the condition that she has not gathered information).

But since there are at least two deontic ideals in play here (the deontic ideal for parental obligation and the deontic ideal for promise keeping), the FDVP fails to yield the result that ST has a higher deontic value than GL. (This is because the ‘all things being equal’ clause is not met. GL’s achieving the deontic ideal for promise keeping renders ‘all other things’ unequal.)

However, I do not think this is a serious problem. Intuitively, in some cases where two deontic ideals give conflicting recommendations, we have more reason to achieve one ideal than we do to achieve the other. Daisy’s Decision is such a case. Thus the following principle seems plausible:

Second Deontic Value Principle: All things being equal option A has higher deontic value than option B if A achieves *some deontic ideal* X and X is (in the relevant context) a more significant ideal than any deontic ideal achieved by B.

Since (in the context of Daisy’s Decision) the deontic ideal for parental obligation is more significant than the deontic ideal for promise keeping, we can assign a higher deontic value to ST than to GL.²² The Second Deontic Value Principle seems intuitively plausible and nicely accounts for Daisy’s duty to gather information. (Note that the Second Deontic Value Principle allows for the possibility that in other contexts the deontic ideal for promise keeping is more significant than deontic ideal for parental obligation. We can identify which deontic ideal has more significance in a particular context by noting which ideal it would be more regrettable to fail to achieve.)

I do not think that the endorsement of the First Deontic Value Principle and the Second Deontic Value Principle will account for all cases in which an agent ought to gather information. We could devise more complex cases, for example cases in which achieving one very significant deontic ideal conflicts with achieving many less significant de-

22. Some subjective deontologists may want to say that the significance of a deontic ideal depends on an agent’s beliefs about significance, while others may claim that significance is objectively determined. I stipulate that Daisy regards the deontic ideal for parental obligation as more significant and I think it is plausible that it is objectively more significant.

ontic ideals or cases in which various deontic ideals are approximated to various degrees. But I do think my approach to cases like Claire's Choice and Daisy's Decision is instructive.

When Subjective Deontologists are confronted with cases in which it is intuitive that the agent has a duty to gather information, they should look for principles that allow them to assign deontic value in a manner that accounts for this duty. Given that the agent has the duty, the failure to gather information will be regrettable. This reveals that the deontic ideal has not been achieved. *So the very fact that the agent has the duty is evidence that the deontic values should be assigned in a way that accounts for the duty.* Thus it is highly plausible that there will always be (perhaps very complex) principles of deontic value that can do the necessary work.

In order to make Smith's critique of SDM theories stick, one would have to present a case in which (a) the agent intuitively has a duty to gather information and (b) no plausible principles of deontic value could be utilized by Subjective Deontologists to account for this duty. It seems unlikely to me that such a case will be discovered.

I also think that my treatment of Devon's Promise is instructive. When we judge that an agent lacks a duty to create a particular obligation (e.g., the duty to care for Kate's cat), we should expect to find that the agent can approximate the relevant deontic ideals to the same degree whether or not she creates the obligation. And we need not generalize from the fact that there are some cases in which an agent has no reason to create an obligation to the strong claim that duties to create obligations are suspect in general. Different cases of created obligation will involve different deontic ideals and can thus be handled quite differently.

VI. TWO POINTS OF REMAINING (IN ONE CASE POTENTIAL) AGREEMENT

I want to note two important points at which I have not taken issue with Smith's view. First I entirely agree that cases like Devon's Promise show that the following principle is true:

D: The mere fact that option 1 creates a weightier duty than option 2 does not entail that there is a reason to select option 1. [A weightier duty is a duty that it would be worse to violate.]

Principle *D* is clearly right. My explanation of Claire's duty to gather information does not appeal to the fact that the duty she would create by gathering information is weightier. Rather it appeals to the fact that fulfilling this duty would achieve the relevant deontic ideal.

Second, recall Smith's view that actions should not be assigned a deontic value higher than zero in virtue of the fact that they fulfill cre-

ated duties. On my approach, deontic value is assigned based on how closely an act approximates the relevant deontic ideal, rather than based on whether the act fulfills a created duty. So, even if an act of keeping a promise has a value of +5, it does not have that value in virtue of the fact that it fulfills a created duty. Thus my approach trivially vindicates Smith's view. No act has any deontic value other than zero merely in virtue of being a created duty.

Furthermore, nothing in my approach requires that we assign positive deontic value to acts that fulfill created duties. It could be that actions which fulfill created duties and achieve the relevant deontic ideals are assigned a deontic value of zero. While acts that fulfill created duties but fail to achieve the relevant deontic ideals are assigned a negative deontic value. Whether an act counts as fulfilling a created duty depends on the agent's other available options (not just on how closely the act approximates the relevant deontic ideal(s)). So on this approach an act could have a negative deontic value, say -5, in virtue of failing to achieve the relevant deontic ideal(s). And it would then maintain that value whether or not it fulfills a created duty. The crucial claim for my purposes is that in some cases acts which fulfill some created duties are assigned higher deontic value (because they more closely approximate the relevant deontic ideals) than acts which fulfill others. But this is consistent with capping the deontic value of such acts at zero.

VII. CONCLUSION

I have argued that Subjective Deontologists can account for the duty to gather information in cases like Claire's Choice. I have also offered reasons for thinking that, in general, Subjective Deontologists will be able to account for agents' duties to gather information in cases where it seems clear that they have such duties. Along the way, I introduced the notion of deontic ideals and explored how deontologists can make use of this notion in assigning deontic value to particular actions. It seems likely that appealing to deontic ideals could prove useful in answering various further questions in moral theory. So I hope that my drawing attention to the notion proves beneficial.