

Duty and Ignorance*

Chad Lee-Stronach[†]

School of Philosophy
Research School of the Social Sciences
Australian National University

Forthcoming in *Ethics*

Abstract

Holly Smith (2014) contends that subjective deontological theories – those that hold that our moral duties are sensitive to our beliefs about our situation – cannot correctly determine whether one ought to gather more information before acting. Against this contention, I argue that deontological theories can use a decision-theoretic approach to evaluating the moral importance of information. I then argue that this approach compares favourably with an alternative approach proposed by Philip Swenson (2016).

I Introduction

Subjective deontological theories hold that what we morally ought to do is sensitive to our beliefs about our situation. Holly Smith (2014) argues that these theories cannot correctly determine whether we ought to gather more information before acting.¹ I will call this: the Problem of Ignorance. To illustrate the Problem, Smith presents following scenario:

*For especially helpful discussion and comments on this paper, I wish to thank John Cusbert, Nicholas DiBella, Edward Elliott, Alan Hájek, Frank Jackson, Joshua Knobe, Seth Lazar, Kirsten Mann, Philip Pettit, Kai Spiekermann, Sergio Tenenbaum, Shang Long Yeo, and an audience of the Moral, Social, and Political Theory Graduate Workshop, at the Australian National University.

[†]Electronic Correspondance: chad.lee-stronach@anu.edu.au

1. Holly M. Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” *Ethics* 125, no. 1 (2014): 11–38.

Layoff: You are a human resources manager who must lay off either Max or Mina due to a recent economic downturn. You know that your decision should be morally just, and that justice is best served by laying off the least productive employee. As it stands, however, you do not have any beliefs about the relative merits of laying off one employee rather than the other. You can spend time today gathering accurate information about them, such that if you did so, you would have enough information to make a good decision tomorrow. Alternatively, you can choose to remain ignorant of these facts, and simply layoff one of them now (say, by adopting the second-best solution of firing the most recently-hired employee, who you know happens to be Max).

Let us stipulate that, morally speaking, you ought to gather more information. The challenge for subjective deontological theories is to give a principled and non-problematic explanation for why they agree with this verdict. Smith contends that no such explanation is available to them.

In response, I shall argue that deontological theories can avoid the Problem of Ignorance by using a standard decision-theoretic approach to evaluating the moral importance of information, even in Smith’s target cases in which we have no or mistaken beliefs about our situation.²

Before proceeding, it is worth noting that this decision-theoretic approach is not new to ethics.³ However, neither Smith nor, in his recent response, Philip Swenson (2016), engage with it in their respective explorations of the Problem of Ignorance.⁴ Instead, they draw conclusions on the basis of their own piecemeal extensions of deontological theory. Below, I outline their respective arguments and show that they lead to mistaken conclusions about whether deontologists face the Problem of Ignorance. This gives rise to a methodological conclusion: to safely navigate the morality of decision-making under incomplete information, moral theorists do well to consult decision theory before proceeding.

Part II presents Smith’s argument against deontology. **Part III** explains how deontologists can avoid the Problem of Ignorance by adopting a decision-theoretic approach to

2. Like Smith, I will focus primarily on cases involving no beliefs. The approach I give here also applies to cases involving mistaken beliefs.

3. See: Frank Jackson, “Decision-theoretic consequentialism and the nearest and dearest objection,” *Ethics* 101, no. 3 (1991): 461–482; Frank Jackson, “How Decision Theory Illuminates Assignments of Moral Responsibility,” chap. 2 in *Intention in Law and Philosophy*, ed. Ngaire Naffine, Rosemary J Owens, and John Williams (Aldershot, UK: Ashgate, 2001), 19–36. The locus classicus of value of information calculations is found in: I. J. Good, “On the Principle of Total Evidence,” *The British Journal for the Philosophy of Science* 17, no. 4 (1967): 319–321. For an accessible introduction, see: Michael D. Resnik, *Choices: An Introduction to Decision Theory* (London: University of Minnesota Press, 1987), pp. 57–59.

4. Philip Swenson, “Subjective Deontology and the Duty to Gather Information,” *Ethics* 125, no. October (2016): 257–271.

information gathering. **Part IV** offers a critical evaluation of Swenson’s (2016) proposed solution. **Conclusion** follows.

II Smith’s Argument

Smith contends that if you are a subjective deontologist then you can never, based on your understanding of your situation, have a moral duty to gather more information before acting. Her argumentative strategy is as follows: she distinguishes between two apparently jointly exhaustive types of duties that deontologists might posit, and then proceeds to argue that neither type allows deontology to avoid the Problem of Ignorance.

The first type of duty is a *freestanding* one that holds that for *all* cases of uncertainty, you should gather information. Smith rejects any approach of this type:

Scrutiny of these proposed freestanding duties to seek information reveals that they have several flaws: either they provide poor advice, or they provide insufficiently detailed advice, or they provide advice for which no rationale is forthcoming.⁵

Simply put, any plausible moral theory should be able to distinguish between cases where we should gather information (and to what extent) from other cases where we should not or need not. Since a freestanding approach is (by definition, it seems) insensitive to context, it is difficult to see how it could correctly balance the various considerations of cost, reliability of information, and so on, which are relevant to determining whether we morally ought to gather more information. On this basis, Smith concludes that subjective deontologists cannot avoid the Problem of Ignorance by positing a freestanding duty to gather more information.⁶

If this is correct, then it seems that a deontological theory must instead adopt a *derivative* duty, one that is grounded on the fact that gathering information can help you to discharge other duties you have. Compared to freestanding duties, it is easier to see how derivative duties will be sensitive to the moral stakes of your situation: if gathering information will help you to better discharge your other duties, then you have duty to do so; if it won’t, then you don’t.

However, Smith argues that any such duty will also incorrectly determine whether to gather more information once it is combined with other putative deontological commitments: specifically, principles I shall call **Subjective Duty** and **No Deontic Value**.⁷ I will briefly

5. Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” p. 19.

6. For a cogent response to this charge, see: Sergio Tenenbaum, “Action, Deontology, and Risk: Against the Multiplicative Model,” *Ethics* 127 (2017): 1–36.

7. Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” pp. 32–34.

explain how these principles furnish the premises of Smith's argument against a derivative duty to gather information.

II.A Premise 1: Subjective Duty

Smith offers the following general principle that deontologists might accept when determining whether or not we ought to gather more information:

Subjective Duty: You have a duty to gather information if and only if you believe that doing so would lead you to produce the maximum amount of deontic value.⁸

I shall later argue that deontologists should reject **Subjective Duty**. For now, however, we will see how it plays an important role in Smith's argument. Applied to **Layoff**, **Subjective Duty** seems to imply the following moral code:

Code C

Clause 1. When you believe that you should lay off employees in light of financial retrenchment, you should:

- a. Lay off whoever you believe is the least productive employee;
- b. Lay off the most recently hired employee, if you do not believe of any employee that they are the least productive.

Clause 2. When you believe that you can gather more information relevant to satisfying **Clause 1**, then you should do so if and only if you believe that it would lead you to produce the maximum amount of deontic value.

Smith concludes, however, that **Code C** incorrectly leads you to avoid gathering information once it is coupled with the following principle: **No Deontic Value**.

8. This is a simplified version of Smith's principle. For the purposes of the argument given here, nothing of substance is lost by the simplification. Compare: "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 leave 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 upon him)." Smith, "The Subjective Moral Duty to Inform Oneself before Acting," p. 24.

II.B Premise 2: No Deontic Value

Smith argues that some actions can *create* a duty that would not otherwise exist. For example, making a promise can create a *pro tanto* duty to perform the promised action, where otherwise this duty would not exist. Likewise, she argues, in cases like **Layoff**, your action of investigating or not investigating will create a duty (by **Clause 1 of Code C**) to layoff a particular employee by leading you to believe whether a particular employee is the least productive.

Drawing on earlier work, Smith contends that deontologists must hold that there is no deontic value to creating a duty, or else we would be required to create increasingly onerous duties: rather than promising to look after someone's goldfish, we would be required to promise to look after their children; we may even be required to damage others' property in order to create the weighty obligation to pay compensation; and so on.⁹ Since we are not required to make such promises or, more generally, to create more burdensome duties, Smith concludes that deontologists must accept:

No Deontic Value: There is no positive deontic value to satisfying a created duty, but only a negative value to violating such a duty.¹⁰

Applied to the present case, Smith argues that since gathering information can create duties (once you become aware of a fact, you can come under a duty to respond to it appropriately), there will be no additional deontic value associated with gathering it as opposed to not doing so.

It is worth noting that deontologists have numerous resources for denying **No Deontic Value**, if they wish. For instance, they might argue that gathering information does not create duties in the relevant sense: rather, gathering information involves *discovering* what your duties actually are.¹¹ On top of this, deontologists can point out that they can take into account personal costs in determining whether one has an obligation to do what is otherwise best.¹² They can also deny that we should create obligations to compensate others (say, by harming them and becoming duty-bound to compensate), since we have reasons to refrain from harming others, even if the benefit we could give them would *overcompensate* them

9. Holly M. Smith, "A Paradox of Promising," *Philosophical Review* 106, no. 2 (1997): 153–196; Smith, "The Subjective Moral Duty to Inform Oneself before Acting."

10. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," p. 29.

11. See: Peter A. Graham, "In Defense of Objectivism about Moral Obligation," *Ethics* 121, no. 1 (2010): p. 91.

12. See: Seth Lazar, "Deontological Decision Theory and Agent-Centered Options," *Ethics* 127, no. 3 (2017): 579–609.

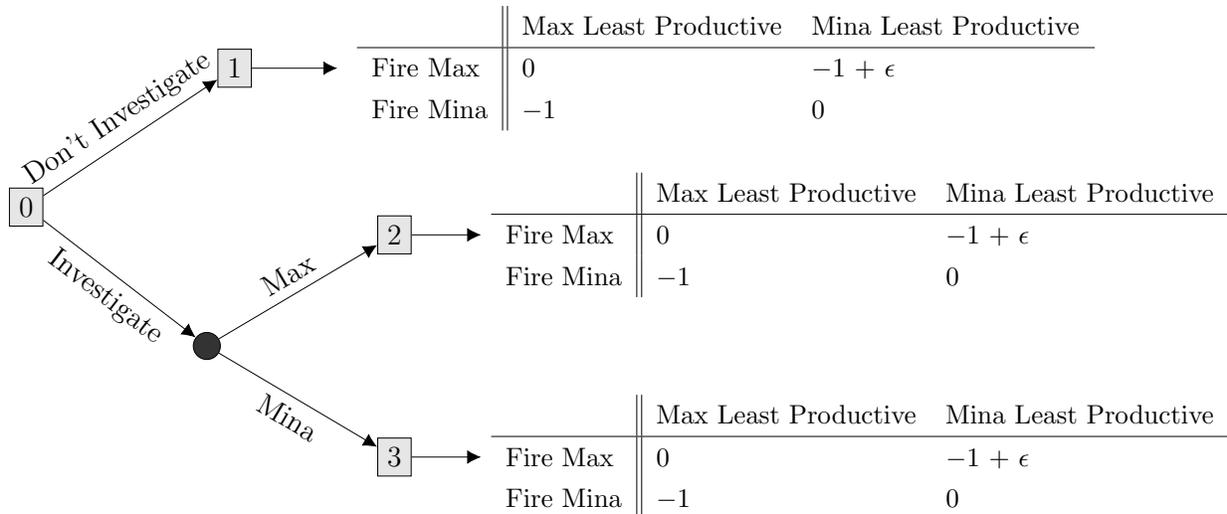


Figure 1: Smith's Argument, applied to **Layoff**

for the harm caused.¹³ More generally, there is no reason to think that deontologists hold that the more onerous a moral duty, the more valuable it is to conform to it. However, it turns out that these substantive lines of argument are unnecessary: as we shall see, **No Deontic Value** is a red-herring that does nothing to block deontology from adopting a decision-theoretic solution to the Problem of Ignorance.

As such, for the sake of argument, we will accept **No Deontic Value** and assign all outcomes in which you fulfil a created duty a value of zero.¹⁴ Without loss of generality, Figure 1 lists the values for the possible outcomes of your choices in **Layoff**, where ϵ is some small positive deontic value representing Smith's stipulation that, since you know that Max is the more recent employee, it is slightly more just that he be fired (other things equal).¹⁵

II.C Conclusion

As shown in Figure 1, whatever you choose, getting more information does not allow you to achieve more deontic value than not getting information. As such, from the point of view of deontic value maximisation, it is permissible for you to choose either option.

More generally, the argument against a derivative duty to gather information can sum-

13. See: Seana Valentine Shiffrin, "Wrongful Life, Procreative Responsibility, and the Significance of Harm," *Legal Theory* 5, no. 2 (1999): 117-148; Judith Jarvis Thomson, *The Realm of Rights* (Harvard University Press, 1990), pp. 228-48.

14. It is worth emphasising, however, that talk of 'zero-value' is misleading. As Smith herself notes (note 28, p. 29), given that we are using interval scales, the zero point is arbitrary, since the values can be rescaled up to positive affine transformation without distorting the ordering of the outcomes.

15. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," p. 25.

marised as follows:¹⁶

Premise 1: You ought to gather more information before acting if and only if you believe that doing so will produce uniquely maximal deontic value. (from **Subjective Duty**)

Premise 2: The deontic value of gathering more information is never greater than that of not gathering information. (from **No Deontic Value**)

Conclusion: Therefore, it is never the case that you ought to gather more information before acting.

Thus, having considered both freestanding and derivative approaches to the Problem, Smith concludes that there is no subjective deontological duty that can save deontology from the Problem of Ignorance.

III A Decision-Theoretic Approach

Crucially, **Subjective Duty** determines what you ought to do on the basis of your full beliefs, rather than on your degrees of belief. I will now argue that deontologists should reject **Subjective Duty** because – even in the specific cases that Smith is targeting – you ought to gather more information because of your degrees of belief, even if you lack all-out belief about whether doing so will maximise deontic value. By including information about degrees of belief, subjective deontologists gain access to the standard decision-theoretic approach to determining whether to gather more information before acting.

III.A Smith’s Argument for Subjective Duty

To motivate **Subjective Duty** and perhaps to preclude a decision-theoretic approach to the Problem, Smith makes the following caveat:

First, I will focus primarily on agents who have false beliefs or no relevant beliefs, rather than agents who are merely uncertain about their prospective action’s character. This restriction will not distort the inquiry or its outcome. If a moral theory flunks the test of appropriately handling the duty to gather information for agents having false beliefs or no relevant beliefs, then it is inadequate, even if it passes this test for agents laboring under uncertainty.

16. See also: Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” pp. 31-32.

This restriction determines the content of **Subjective Duty**, which defines duties in terms of beliefs or lack of belief, rather than in terms of degrees of belief (what Smith refers to as ‘uncertainty’). Recall:

Subjective Duty: You have a duty to gather information if and only if you believe that doing so would lead you to subsequently produce the maximum amount of deontic value.

As we saw, this restriction plays a crucial role in Smith’s argument. In **Layoff**, since (by stipulation) you do not believe that investigating will lead you to produce the maximum amount of deontic value (due to the **No Deontic Value** assumption), it is not the case that you ought to do it.

Unfortunately, Smith never explains what she means by ‘belief’ or what it means to have ‘no relevant beliefs’. More to the point, she never explains why your having no beliefs or false beliefs about a proposition entails that you cannot be more, less or similarly confident that the proposition is true as opposed to its negation. As it stands, there is no apparent inconsistency between not believing that Max is the least productive employee, not believing that Mina is the least productive employee, while nevertheless being more, less, equally, or similarly confident that one of them is less productive than the other. This suggests that Smith’s target cases – those involving no or mistaken beliefs – and cases involving uncertainty are not mutually exclusive. Indeed, I shall argue that her target cases are a proper subset of cases of uncertainty. If that is correct, then – as we shall see below – deontologists can use standard decision-theoretic reasoning to explain why your degrees of confidence can make it the case that you should gather information *even if* you do not believe that doing so will maximise deontic value (thus rejecting the ‘only if’ clause in **Subjective Duty**).

III.B A Deontic Value of Information Approach

In presenting a positive proposal of how subjective deontology can avoid the Problem of Ignorance, I will introduce some terminology that helps to bridge deontological moral theory and decision theory.

For the moment, I will try (where possible) to avoid the term ‘duty’. This is because duty-speak can be unclear about how exactly duties are relevant to decision-making in risky situations. For instance, are duties the *output* of a decision process (e.g. your duties are just the set of permissible actions)? Or are duties the *inputs* of a decision process (e.g. you consider your *pro tanto* duties in your deliberations regarding which of your available actions are morally permissible)? Or are duties some *function* from the grounds of obligation

to some set of permissible actions (whereby your duty is to consider your reasons for action and to determine, from them, what you morally ought to do)?

In an attempt to side-step this ambiguity, I will adopt the following choice-theoretic framework.¹⁷ Let a *choice context* be a set of actions, a set of possible states of the world, and a belief function defined over those states that represents with probabilities how confident you are that a particular state is actually the way the world is. When an action is performed under a particular state of the world, it produces an *outcome*. This outcome encodes whether moral considerations have been upheld or violated: it may represent the nature of the action (for instance, it may represent the fact that you intentionally lied) and it may also represent the causal consequences of your action (for instance, it may represent whether or not your interlocutor believed your lie). An action's *prospect* is the set of its possible outcomes and their respective probabilities of obtaining. We will assume that these prospects can be ordered and represented by a *deontic value function*, which assigns a higher deontic value to outcomes, the more morally important they are. A *moral decision rule* is a function from a choice context to a set of morally permissible actions. It draws on some or all of the above information to determine what, if anything, is morally permissible.

Understood in this framework, the challenge for deontologists is to find an unproblematic moral decision rule that is consistent with their core tenets and which allows them to avoid the Problem of Ignorance. Smith's claim is that no such moral decision rule exists or – at the very least – coming up with such a rule would involve ‘dauntingly complex theoretical issues.’¹⁸

The good news for deontologists is that these theoretical issues are already well-recognised and catered for in the fields of decision theory and formal epistemology. For example, Smith raises the challenge of modelling cases involving ‘lack of belief’ using degrees of belief, arguing that these approaches must adopt problematic rules like the Principle of Indifference.¹⁹ However, deontologists should not accept this gambit: following orthodox Bayesianism, they can point out that even if the Principle of Indifference is false (which is debatable), they do not need to invoke it because even in these cases, we nevertheless have precise and coherent

17. This general framework is useful for modelling a wide range of non-consequentialist moral theories, since it allows for agent-relative perogatives, menu-dependence, and other structural features commonly associated with deontology, see: Franz Dietrich and Christian List, “What Matters and How it Matters: A Choice-Theoretic Representation of Moral Theories,” *The Philosophical Review* 126, no. 4 (2017): 421–479.

18. Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” p. 33.

19. Roughly speaking, the Principle of Indifference says that when we lack any evidence about our situation, we should assign equal probabilities to the various possibilities that might occur. The problem is that possibilities can be redescribed in seemingly equally eligible ways to yield different numbers of possibilities, leading to a conflicting assignment of probabilities. See: Bas van Fraassen, *Laws and Symmetries* (Oxford: Clarendon Press, 1989), ch. 12.

degrees of belief.²⁰

However, a more charitable understanding of Smith’s concern is that her target cases are those not only in which we lack all-out belief; they are cases where we lack all-out belief because we lack evidence. It is arguable – though still a matter of significant debate – that lacking evidence prevents us from having precise degrees of belief.

Nevertheless, decision theory can help us navigate these cases, too. For example, deontologists can represent themselves as having an indeterminate credal state that reflects their equivocal, incomplete, or otherwise ambiguous evidence.²¹ Where a determinate degree of belief specifies a precise numerical value (indeed, precise up to infinite decimal places), an indeterminate degree of belief admits of a range of such values, represented by a set of probability functions. By positing sets of probability functions, deontologists can model the cases that Smith targets, without appealing to the Principle of Indifference.²²

Following this latter approach, in **Layoff** we will model your epistemic situation of ‘lacking beliefs’ in terms of indeterminate degrees of belief. We will suppose that you are roughly equally confident that Max or Mina is the least productive, such that the probabilities range within the interval $[0.4, 0.6]$ (or ≈ 0.5 , for ease of exposition).

With this prior, albeit indeterminate, degree of belief, investigating will allow you to update your degrees of belief and choose from a more informed standpoint. Let us suppose that your investigation will be 99% accurate. Given this information, being a rational agent you will update your degrees of belief to reflect the evidence. As we shall see in **Layoff**, this information will improve the prospects associated with laying off one employee rather than the other. Following this approach, deontologists should reject **Subjective Duty** in favour of the following moral decision rule:

Choose Undefeated Prospects: An action is morally permissible if and only if its prospect is undefeated by that of any alternative action.

An ‘undefeated prospect’ is one that is not ranked lower than any other prospect.²³ This can

20. See, for example: Adam Elga, “Subjective Probabilities should be Sharp,” *Philosophers’ Imprint* 10, no. 5 (2010): 1–11.

21. James M. Joyce, “How Probabilities Reflect Evidence,” *Philosophical Perspectives* 19 (2005); James M. Joyce, “A Defense of Imprecise Credences in Inference and Decision Making,” *Philosophical Perspectives* 24 (2010); Alan Hájek and Michael Smithson, “Rationality and Indeterminate Probabilities,” *Synthese* 187 (2012): 33–48; Isaac Levi, “Why Indeterminate Probability is Rational,” *Journal of Applied Logic* 7, no. 4 (December 2009): 364–376.

22. For defences of this approach, see: Isaac Levi, *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance* (Cambridge, Massachusetts; London, England: MIT Press, 1980); Richard C. Jeffrey, *The Logic of Decision*, Second (Chicago and London: The University of Chicago Press, 1983); Mark Kaplan, *Decision Theory as Philosophy* (Cambridge: Cambridge University Press, 1996); Joyce, “How Probabilities Reflect Evidence”; Joyce, “A Defense of Imprecise Credences in Inference and Decision Making.”

23. A defeat relation is logically weaker than a weighing relation. Unlike a weighing relation, a defeat

be conveniently represented as one action having a determinately higher probability-weighted average (or ‘expected’) deontic value (EDV) than another.

Choose Undefeated Prospects solves the Problem of Ignorance because it allows deontologists to use what is equivalent to a ‘value of information’ calculation to determine when it is morally required to gather more information. Roughly put, value of information calculations involve comparing your optimal future actions in an informed context (which results from having gathered more information) versus an uninformed one. If the informed context has an optimal action whose prospect defeats that of every action in the uninformed context, then you should choose to gather more information.²⁴

For example, in **Layoff**, you have a choice between acting in an informed choice context (by investigating) or in an uninformed choice context (by not investigating). The rule **Choose Undefeated Prospects** determines whether or not you should investigate by calculating whether the optimal action in the informed choice context has higher expected deontic value than the optimal action in the uninformed choice context. Figure 2 illustrates your choice situation based on your prior and posterior degrees of belief, giving the expected deontic values of your available courses of action (where *Max* denotes that Max is found to be the least productive and *Mina* denotes that Mina is found to be the least productive).

If you choose to remain uninformed, the morally optimal action of your later choice situation has an expected deontic value of between -0.4 and -0.6 (or ≈ -0.5 for short, as in the diagram above). Compare this with the context in which you investigate, where the optimal actions have an expected deontic value of approximately -0.01 , which is greater than that of the optimal actions in the uninformed context.²⁵ According to **Choose Undefeated Prospects**, you are morally required to investigate, even though you do not all-out believe that doing so will maximise actual (as opposed to expected) deontic value. This result shows that there is a moral decision rule (or ‘subjective duty’) that subjective deontologists can adopt to avoid the Problem of Ignorance, thus refuting Smith’s main contention. The question, then, is whether deontologists can avail themselves of this solution without giving up their core substantive commitments.

relation can allow for incommensurable considerations, strong moral dilemmas, and (arguably) moral options. See: Dietrich and List, “What Matters and How it Matters: A Choice-Theoretic Representation of Moral Theories,” p. 432.

24. For an accessible introduction, see: Resnik (1987). For details on how such calculations work with indeterminate probabilities (as in our current case), see: Bradley and Steele (2016). Resnik, *Choices: An Introduction to Decision Theory*; Seamus Bradley and Katie Steele, “Can Free Evidence Be Bad? Value of Information for the Imprecise Probabilist,” *Philosophy of Science* 83, no. 1 (2016): 1–28.

25. Strictly speaking, this will be an interval of values around the value 0.1.

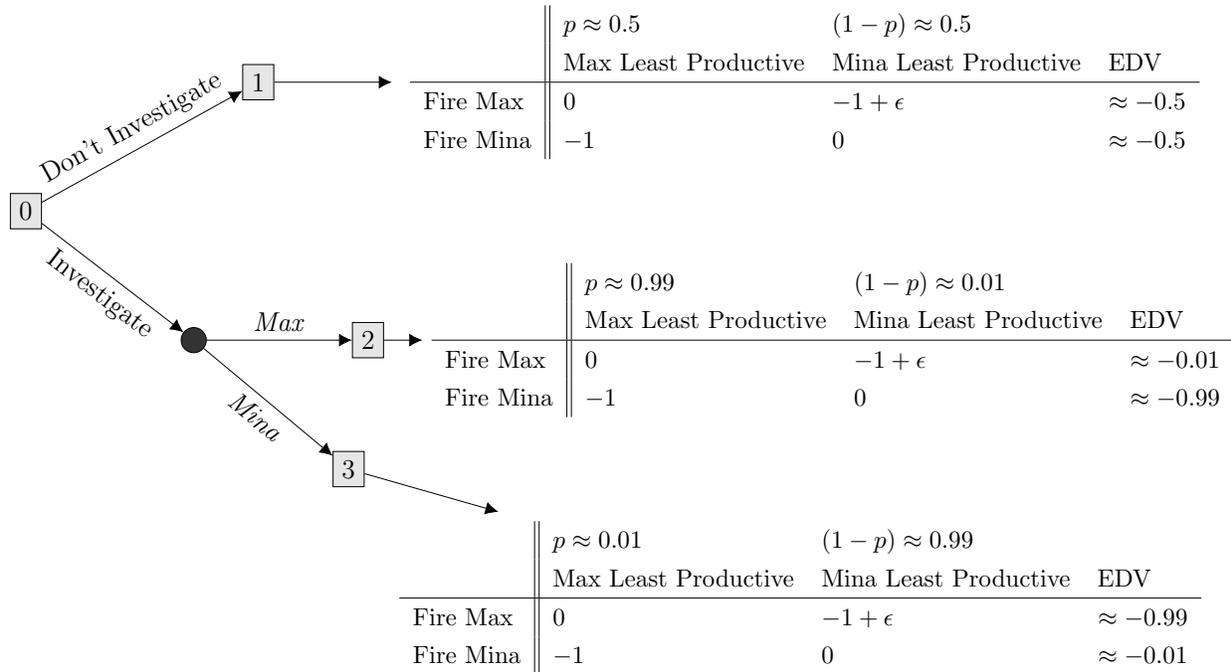


Figure 2: Value of Information Calculation

III.C Objections

Perhaps anticipating this approach, Smith offers some lines of resistance to any decision-theoretic model of deontology. For instance, she raises the following concerns:²⁶

Claire could be uncertain whether laying off Mina would be laying off the least-productive employee. This suggests that we need a new code incorporating multiple clauses in place of the earlier clause 1. Each clause would define a distinct duty, such as a duty to “lay off a person who you think has a .9 chance of being least productive,” a duty to “lay off a person who you think has a .8 chance of being least productive,” and so forth. What is the deontic value of fulfilling these various duties? Perhaps, if laying off someone you think has a 1.0 chance of being least productive has a deontic value of 2, then laying off someone you think has a .9 chance of being least productive has a lesser deontic value of 1.8 ($2 \times .9$). But what is the deontic value of laying off someone you think has merely a .4 chance of being least productive? Should this be a small positive value, or a negative value? Moreover, if laying off someone you think has a 1.0 chance of being least

26. Smith, “The Subjective Moral Duty to Inform Oneself before Acting,” p. 33. In fact, some deontologists have raised similar concerns. See: Ron Aboodi, Adi Borer, and David Enoch, “Deontology, Individualism, and Uncertainty: A Reply to Jackson and Smith,” *Journal of Philosophy* 105, no. 2 (2008): p. 272; Tenenbaum, “Action, Deontology, and Risk: Against the Multiplicative Model,” p. 9.

productive has a value of 0 rather than 2 (as I have argued), then what is the deontic value of laying off someone you think has a .9 chance of least productive? Should this value also be 0, or should it be less than 0, or more than 0?

There are a few ways of understanding this passage. On one interpretation, the concern is that deontological moral considerations are not ordered in the right way to allow for a well-defined measure of deontic value. Bracketing the fact that Smith's argument presumes that deontological considerations *can* be represented by a deontic value function, there are reasons to think that deontology should arguably reject at least some of the axioms of orthodox expected value theory.²⁷ However, this requires further investigation, not least because decision theorists have made many extensions and generalisations of orthodox decision theory that allow it to accommodate risk attitudes, incommensurability, and lexical priority – all while maintaining a systematic and coherent approach to uncertainty.²⁸ It is therefore not enough to show that deontology cannot be represented by *some* particular version of expected value theory; those pressing this challenge need to show that deontology cannot be represented by *any* version of expected value theory. We are yet to see an impossibility result to this effect.

A different interpretation of the concern attends to the question of whether deontic value is linear with probabilities. This is related to the question of whether deontologists should be *risk-neutral*: that is, whether they should consider two prospects that have same expected deontic value to be equally choice-worthy, even if one prospect has much greater variance (or greater 'spread') in possible deontic value than the other. As it turns out, however, non-neutral attitudes to risk are consistent with orthodox decision theory.²⁹ Although deontologists are yet to settle such questions as whether non-neutral risk attitudes are morally appropriate, the solution given above is compatible with various answers to the question.³⁰

On either interpretation, these speculative concerns are not sufficient to block deontologists from adopting the tools of decision theory. While more work needs to be done to con-

27. For arguments to this effect, see: Lazar, "Deontological Decision Theory and Agent-Centered Options"; Seth Lazar, "Deontological Decision Theory and the Grounds of Subjective Permissibility" (2018); Tenenbaum, "Action, Deontology, and Risk: Against the Multiplicative Model."

28. John W. Pratt, "Risk aversion in the small and in the large," chap. 19 in *Handbook of the Fundamentals of Financial Decision Making* (1964), 317–331; Robert J. Aumann, "Utility Theory without the Completeness Axiom," *Econometrica* 30, no. 3 (1962): 445–462; Melvin Hausner, "Multidimensional Utilities," chap. 12 in *Decision Processes*, ed. R. M. Thrall, C. H. Coombs, and R. L. Davis (New York: John Wiley / Sons, Inc., 1954). See also: Dietrich and List, "What Matters and How it Matters: A Choice-Theoretic Representation of Moral Theories."

29. Pratt, "Risk aversion in the small and in the large"; Kenneth J. Arrow, "The Theory of Risk Aversion," chap. 2 in *Aspects of the Theory of Risk Bearing* (Helsinki: Yrjo Jahnsolin Saatio, 1965); John Broome, *Weighing Goods* (Cambridge, Massachusetts: Blackwell, 1995), Ch. 4.

30. For example, they may adopt a risk-weighted utility approach, as argued by: Lara Buchak, "Taking Risks behind the Veil of Ignorance," *Ethics* 127 (2017): 1–35.

struct a fully-fledged deontological decision theory, it appears that subjective deontologists can adopt a highly general moral decision rule, such as **Choose Undefeated Prospects**, to avoid the Problem of Ignorance.

IV Swenson's Account

Let us now compare the solution given above with that proposed by Philip Swenson (2016).³¹ In response to Smith's argument, Swenson accepts **No Deontic Value**, but gives a positive proposal for why it does not apply to actions that create duties by gathering more information, thereby allowing subjective deontologists to assign higher deontic value to gathering more information as opposed to not doing so. I will briefly show that Swenson's proposal is underspecified in many crucial respects, such that it does not provide a satisfactory solution to the Problem of Ignorance. I will argue that a more complete version of his proposal will be viable only insofar as it conforms to the decision-theoretic approach presented earlier.

Swenson's proposal is that deontologists should assign higher deontic value to actions that better *approximate* (in an undefined sense) deontic ideals. In **Layoff**, he argues that you should investigate because, intuitively, this better approximates the deontic ideal of achieving *justice in layoffs*, and hence that action has higher deontic value than not investigating. How do we know that it would better approximate this ideal? Because we would view it to be morally "regrettable" if you were to instead choose to stay ignorant. More generally, Swenson argues that:

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.

Swenson is confident that deontologists will be able to use these 'regretability' and 'approximation' heuristics on a case-by-case basis to derive a well-defined measure of the deontic value of actions. However, deontologists should not be so sanguine, for the following reasons.

31. Swenson, "Subjective Deontology and the Duty to Gather Information."

IV.A Approximation and the Second-Best

Swenson’s solution for deontologists puts forward the idea that, other things equal, an action has a higher deontic value than another if it better approximates a deontic ideal. As mentioned, however, no account is given of what ‘approximation’ means, nor why deontologists can rely on it to induce a coherent assignment of deontic value.

In particular, one immediate concern with any approximation-based approach to defining the deontic value of actions is that it not commit the Fallacy of the Second-Best: cases where it is mistakenly assumed that the closer an action approximates an ideal, the better that action is.³² These cases arise most clearly when approximating an ideal is valuable when and only when a number of interdependent conditions are met. In the absence of a sufficient combination of these conditions being met, closer approximations of an ideal are potentially worse than more distant ones.

For example, in **Layoff**, suppose that the deontic ideal of *justice in layoffs* is much better approximated by gathering information than by simply remaining ignorant and firing the most recent employee. However, suppose also that you can only gather information by violating the privacy of Max and Mina. On a straightforward application of an approximation principle, it would seem that gathering information via a privacy violation better approximates the deontic ideals, since it at least involves some weighty component of the ideal (taking the effort to gather information to determine the least productive employee). However, at the same time, it would seem mistaken to simply conclude that this act of gathering information therefore has greater deontic value from the point of view of *justice in layoffs*, since it is arguable that the value of acting on meritocratic concerns regarding productivity in this case also depends on having respected the privacy of others.

Of course, given that Swenson’s notion of ‘approximation’ is undefined, his theory neither determinately commits (nor, for that matter, avoids) the Fallacy of the Second-Best in these types of cases. However, the main point is that deontologists who choose to develop his approach will need to do so with one eye on preserving a well-defined structure of deontic value, as defined by the kind of decision-theoretic model given earlier.

32. This fallacy is now well-recognised in moral and political philosophical discussions of feasibility and the ideal/non-ideal theory distinction. See, for example: David Enoch, “Against Utopianism: Noncompliance and Multiple Agents,” *Philosopher’s Imprint* 18, no. 16 (2018): 1–20; Juha Räikkä, “The Problem of the Second Best: Conceptual Issues,” *Utilitas* 12, no. 2 (2000): 204–218. Its origins can be found in economic theory: Kelvin Lancaster and R. G. Lipsey, “The General Theory of Second Best,” *The Review of Economic Studies* 24, no. 1 (1956): 11–32.

IV.B Regretability and Intransitivity

Swenson’s proposal also repeatedly appeals to an undefined notion of ‘regretability’ to determine an action’s deontic value.³³ As before, it is not possible to determinately attribute a particular problem to an underspecified account. Nevertheless, deontologists who wish to develop this account should note that decision-theorists have long identified a number of pitfalls for adopting this notion in decision-making. In particular, if ‘regretability’ departs from the decision-theoretic foundations sketched in **Part III**, it will lead to intransitive orderings of options, such that there is no well-defined measure of value available.³⁴

This point is also now familiar in moral theory.³⁵ To illustrate, suppose now that in **Lay-off** you must decide between laying off Max, Mina, or Manuel, taking into consideration the following factors: productivity (approximating the deontic ideal of *justice in layoffs*), their number of dependents (approximating *beneficence in layoffs*), and fairness (approximating *equality of opportunity*). Suppose that Mina is the most productive (followed by Max, then Manuel), Max has the most dependents (followed by Manuel, then Mina), and that fairness asks you to prioritise Manuel or Mina over Max. As others have pointed out, in cases like these, there is a very robust sense in which whatever you choose will be regrettable in some respect compared to some other action.³⁶ But if ‘regretability’ is your proxy for deontic value, as Swenson recommends, then there is no deontic value to speak of, since whatever you choose will have to have lower deontic value than some other alternative.

IV.C Deontic Value and Uncertainty

A final concern about Swenson’s proposal is that it does not give subjective deontologists any guidance in situations involving degrees of uncertainty. Recall that in our rendering of **Layoff**, you are 99% certain that your investigation will allow you to correctly determine who is the least productive. Should you investigate? Swenson’s account has no answer. Indeed, it falls silent for *any* degree of uncertainty. The question is: can it be extended to cover such cases, thus providing a general solution to the Problem of Ignorance for deontologists?

At this point, one might think to follow Swenson’s preferred approach of positing ad-

33. Swenson, “Subjective Deontology and the Duty to Gather Information,” pp. 263, 264, 267, 269, 270.

34. David E. Bell, “Regret in Decision Making under Uncertainty,” *Operations Research* 30, no. 5 (1982): 961–981; Han Bleichrodt and Peter P. Wakker, “Regret Theory: A Bold Alternative to the Alternatives,” *Economic Journal* 125, no. 583 (2015): 493–532; Sushil Bikhchandani and Uzi Segal, “Transitive Regret,” *Theoretical Economics* 6, no. 1 (2011): 95–108.

35. The following case is based on: Larry S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford; New York: Oxford University Press, 2012), ch. 7. See also: Frances Myrna Kamm, “Supererogation and obligation,” *Journal of Philosophy* 82, no. 3 (1985): 118–138.

36. See also: Caspar Hare, “Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had not Done What You Did,” *Noûs* 45, no. 1 (2011): 190–206.

ditional principles on a case-by-case basis. However, this would be to ignore the work of decision theorists who have already charted this logical space, determining which combinations of principles are (and are not) able to give us acceptable guidance in an uncertain world. Their work tells us that insofar as a more fully-developed version of Swenson's account will help subjective deontologists in cases of uncertainty at all, it will do so only by replicating some version of the decision-theoretic approach presented above.

V Conclusion

Before concluding whether particular moral theories, like deontology, can give correct recommendations about cases involving incomplete information, moral theorists do well to investigate the resources and insights of decision theory. By doing so, we have seen that subjective deontologists can avoid the Problem of Ignorance by using deontic value of information calculations. Given the many extensions and generalisations of decision theory that now exist, it seems increasingly clear that deontologists can avail themselves of a variety of useful, but often still overlooked, decision-theoretic approaches to determining what is right in an uncertain world.