Moral Uncertainty for Deontologists

Christian Tarsney

(Forthcoming in *Ethical Theory and Moral Practice*.)

Abstract

Defenders of deontological constraints in normative ethics face a challenge: how should an agent decide what to do when she is uncertain whether some course of action would violate a constraint? The most common response to this challenge has been to defend a *threshold principle* on which it is subjectively permissible to act iff the agent's credence that her action would be constraint-violating is below some threshold t. But the threshold approach seems arbitrary and unmotivated: what would possibly determine where the threshold should be set, and why should there be any precise threshold at all? Threshold views also seem to violate "ought" agglomeration, since a pair of actions each of which is below the threshold for acceptable moral risk can, in combination, exceed that threshold. In this paper, I argue that *stochastic dominance reasoning* can vindicate and lend rigor to the threshold approach: given characteristically deontological assumptions about the moral value of acts, it turns out that morally safe options will stochastically dominate morally risky alternatives when and only when the likelihood that the risky option violates a moral constraint is greater than some precisely definable threshold (in the simplest case, .5). I also show how, in combination with the observation that deontological moral evaluation is relativized to particular choice situations, this approach can overcome the agglomeration problem. This allows the deontologist to give a precise and well-motivated response to the problem of uncertainty.

1 Introduction

We are often uncertain about what we morally ought to do. Such uncertainty can arise in two importantly distinct ways. On the one hand, some moral uncertainty is grounded in empirical uncertainty: for instance, is this substance that I am about to put in my friend's coffee sweetener, or is it arsenic (Weatherson, 2014)? On the other hand, there is what we might call *purely* moral uncertainty, uncertainty about the basic principles of morality, which no empirical information could directly or easily resolve: for instance, is it permissible to tell my friend a white lie about his new haircut?

Moral consequentialists have typically had the most to say about uncertainties of both kinds. With respect to empirically based moral uncertainty, consequentialists typically claim that agents should choose practical options that have maximal *expected value* (the expected value of an option O being found by summing, over all possible outcomes of that action, the value of that outcome times its probability conditional on one's choosing O). This expectational formula can be adjusted in a variety of ways, e.g. to allow for risk aversion or risk seeking.

The study of *purely* moral uncertainty is still in its infancy, but to date the dominant positive approaches in this literature have shared an almost exclusively consequentialist, expectational flavor: Lockhart (2000), Ross (2006), Sepielli (2009), and MacAskill (2014), for instance, all defend expected value approaches to decision-making under purely moral uncertainty. The critics of this approach, on the other hand, have predominantly been those who deny the need for a theory of rational choice under purely moral uncertainty, because they deny that what an agent ought to do depends on her purely moral beliefs or degrees of belief (e.g. Weatherson (2014), Harman (2015), Hedden (2016)).[1]

Expectational approaches to uncertainty, however, seem ill-suited for if not actively incompatible with non-consequentialist moral theories like Kantian deontology, most obviously because (unlike, say, classical utilitarianism) there is no natural way of interpreting these non-consequentialist theories as assigning degrees of rightness and wrongness that are amenable to being multiplied by probabilities and summed to yield quantitative expectations. Even setting aside any thought that theories of choice under uncertainty must take an expectational form, we will see in the next section that there are substantial reasons to doubt whether Kantians et al can provide any plausible answer to certain inescapable questions about moral choice under uncertainty.

The purpose of this paper, however, is to suggest one way in which deontologists *can* say something definite and precise about decision-making under moral uncertainty, of both the empirically based and purely moral varieties. I will propose that *stochastic dominance reasoning* provides powerful motivation for a *threshold principle* of choice under uncertainty, which implies *inter alia* that when an agent A faces a choice between practical options O and P, where P is known for certain to be morally permissible and the status of O is uncertain, A subjectively ought not choose O if her credence that O is objectively impermissible is greater than or equal to .5.

In the next section I describe in greater detail the challenge to deontology posed by both kinds of moral uncertainty, which some in the recent literature (e.g. Jackson and Smith (2006, 2016) and Huemer (2010)) have characterized as a decisive objection at least to absolutist forms of deontology. In §3 I introduce the idea of stochastic dominance reasoning. In §4 I show that stochastic dominance reasoning yields substantive and plausible conclusions for how an agent who accepts the general framework of deontological, agent-centered constraints (whether

[1]It is standard in the literature to reserve the term "moral uncertainty" for what I am calling *purely* moral uncertainty. I use the term more broadly in this paper both to avoid repeated inelegant references to "morally relevant empirical uncertainty" and because one of my aims will be to emphasize the continuity between empirically-based and purely moral uncertainty, so it is helpful to have an umbrella term that refers to both at once.

absolute or defeasible) ought to act under both kinds of moral uncertainty, which in at least the simplest cases will take the form of a straightforward threshold principle. Finally, §5 addresses a powerful objection to threshold principles given by Jackson and Smith, that such principles seem to violate "ought" agglomeration, since a pair of actions each of which is below the threshold for acceptable moral risk can, in combination, exceed that threshold. I argue that, in combination with the observation that deontological moral evaluation is relativized to the particular choice situations in which acts arise, the stochastic dominance approach can overcome this objection and preserve "ought" agglomeration.

2 Absolutist Deontology and Uncertainty

Following Portmore (2016, p. 7), a deontological ethical theory may be characterized as one that denies what all consequentialist theories affirm, namely, that the objective deontic status and degrees of objective rightness/wrongness of any act are determined by how the outcome or prospect of that act compares to the outcomes or prospects of its alternatives.[2] More precisely (to avoid accidentally counting as deontological views that deny that acts have deontic status or degrees of rightness), a deontological theory is one that at least sometimes attributes objective deontic status and/or degrees of objective rightness to acts for reasons that do not depend on how the outcomes of those acts compare to the outcomes of alternatives. An imperfect litmus test for whether a theory is deontological is whether it endorses the existence of *agent-centered constraints*, act types that are treated as *prima facie* objectively wrong or prohibited and such that a token of that act type is at least sometimes objectively wrong or prohibited even when its performance would prevent multiple future tokens of the same act type that are similar in all ethically relevant respects to the prohibited token.[3] Thus a deontological theory might hold, for instance, that it is wrong for me to lie now even though my lie will prevent three future acts of lying that are similar to mine in all ethically relevant respects (except that they will not themselves prevent future instances of lying).[4]

An *absolutist* deontological theory holds that, for some act types characterized non-relationally

[2] A prospect is an objective probability distribution over outcomes; hereafter I will use "outcome" to mean "outcome or prospect." I will likewise speak of "degrees of rightness" when I mean "degrees of rightness or wrongness."

[3] The litmus test is imperfect because ethical theories that incorporate these *prima facie* deontological features can be "consequentialized" (Portmore, 2011). As we will see below, a desideratum for deontological approaches to uncertainty is to avoid a drift in the direction of expectational reasoning about the risk of constraint violations, which, if fully embraced, amounts to evaluating actions by a comparison of their prospects and hence to consequentialism, despite the presence of agent-centered constraints.

[4] Agent-centered constraints may also be positive, e.g., I am required to keep a promise even if my doing so will result in several fewer future instances of promise keeping, or more instances of promise violation.

(i.e., whether an act is an instance of the act type does not depend on any fact about how it relates or compares to its available alternatives), not only is their deontic status not fully determined by how their outcomes compare to those of alternatives, but all acts of that act type have the same deontic status regardless of *any* facts about their outcomes or those of alternatives. For instance, an absolutist theory might hold that all acts that have the non-relational property of *being a lie* (asserting something that the agent believes to be false with the intent that another agent believe it, in a circumstance where the second agent has not consented to being so deceived) is wrong even when telling the lie would prevent the deaths of a thousand innocent people who will be killed if the agent does anything other than lie.

Deontological theories in general and absolutist theories in particular face the following challenge: How should an agent decide what to do when she is uncertain whether a given act is an instance of a deontologically prohibited (or required) act type? For instance, consider the following case.

Possible Promise I have unexpectedly found myself in possession of tickets to the Big Game this afternoon. But as I am celebrating my good fortune, it occurs to me that I may have promised my friend Petunia that I would help her repaint her house later today, at just the time when the game is being played. I have a vague recollection of making such a promise but can't remember with any confidence. I have no immediate way of contacting my friend and must decide, here and now, whether to go to the game at the risk of breaking a promise, or go to her house, avoiding the moral risk but missing the game.

Possible Promise is an instance of empirically based moral uncertainty, but we can describe a closely related case of "pure" moral uncertainty.

Dubious Promise A week ago, Petunia sent me a text message asking if I would help paint her house today. I replied, saying that I would. Unbeknownst to Petunia, however, I was in the hospital at the time, recovering from a minor operation and under the influence of a fairly strong narcotic painkiller. By the time the influence of the painkiller subsided, I had completely forgotten my conversation with Petunia, and only just remembered it a moment ago, while planning my trip to the Big Game.

I take it that, under these circumstances, I may reasonably be uncertain whether I am morally required to skip the game and help my friend, even if I am certain of a background deontological conception of morality on which an ordinary, fully capacitated promise would have been morally

binding.[5] Because my moral uncertainty does not trace back to any empirical uncertainty (e.g. about what I told Petunia, my mental state at the time, or her expectations of my future behavior), it is pure moral uncertainty, uncertainty about the content of a basic moral principle.

In either version of the case, the question is what I should do given my uncertainty about what morality requires of me. To frame this problem rightly, we must introduce a familiar distinction between *objective* and *subjective* normative properties. The objective normative properties of an act depend on facts about the world, but not (in general) on facts about an agent's beliefs or evidence. The subjective normative properties of an act, conversely, depend on facts about the agent's beliefs and/or evidence but not on facts about the world independent of the agent's beliefs/evidence. For instance (borrowing a helpful illustration from Portmore (2016)), suppose you are tasked with defusing a bomb and must decide whether to cut the green wire or the red wire. You believe, and your evidence overwhelmingly indicates, that cutting the green wire will defuse the bomb while cutting the red wire will cause the bomb to explode. But in fact, cutting the red wire will defuse the bomb, while cutting the green wire will cause the bomb to explode. In this case, it is objectively right but subjectively wrong to cut the red wire, and objectively wrong but subjectively right to cut the green wire.[6]

This distinction in hand, let's return to the case of the Possible Promise. There are various things a deontologist can say about this case. However, it is generally agreed that she should *not* say any of the following: (1) It is subjectively wrong for me to go to the game iff I have *any* positive credence that I promised to help Petunia paint her house and hence that it would be *objectively* wrong for me to go to the game. (2) It is subjectively wrong for me to go to the game iff I am *certain* (i.e., credence 1) that I promised to help Petunia paint her house and hence that it would be objectively wrong for me to go to the game. (3) It is subjectively wrong for me to go to the game iff I *did in fact* promise to help Petunia paint her house and hence it *would in fact* be objectively wrong for me to go to the game.

[5] Adjust the strength of the imagined painkillers as needed until this case strikes you as one of substantial moral uncertainty.

[6] Among the possible objective normative properties of an act are being objectively right or wrong; being objectively obligatory, permissible, or forbidden; being what the agent objectively ought or ought not do; and being required, permitted, or forbidden in virtue of an agent's all-things-considered objective reasons. For each of these objective properties, we may define a subjective counterpart. In general, whether an action has a given subjective normative property will depend on the agent's beliefs and/or evidence regarding the corresponding objective property. For our purposes, the distinctions between the various normative properties in each category will not matter very much, so I will sometimes switch between talking e.g. of objective wrongness, objective prohibition, and objective "ought not" as convenience dictates. Our central focus, however, will be on the interaction between an agent's beliefs about the objective properties of being (deontologically) *morally required* or *morally prohibited* and the subjective property of being *rationally required* or *rationally prohibited*, that is, being required or prohibited in virtue of an agent's all-things-considered subjective reasons. These concepts are related by way of objective reasons: a rational agent who believes, e.g., that she is deontologically morally required to choose some practical option O will believe that she has decisive or at least very strong objective reason to choose O, and (at least in general) it is an agent's beliefs about her objective reasons that give rise to her subjective reasons.

Answer (1) is unsatisfactory because I can rarely if ever be *certain* that a course of action does not violate a promise or some other source of deontological obligation. Hence, the principle behind (1) implies that all or nearly all acts are subjectively wrong. Answer (2) is unsatisfactory for a symmetrical reason: I can rarely if ever be certain that an act *would* violate a deontological obligation, and hence if those obligations are only ever relevant to the subjective normative properties of my actions under conditions of certainty, they are effectively inert. Answer (3) is unsatisfactory because, by collapsing subjective deontic status into objective deontic status, it fails to serve the purpose for which subjective normative concepts are introduced, namely to provide epistemically imperfect agents with useful action guidance in the face of uncertainty.

Absent a better answer than these, deontology is in trouble. It will seem either to paralyze action entirely (as per (1)), to collapse into consequentialism insofar as its agent-centered constraints are never practically relevant to epistemically imperfect agents (as per (2)), or to be simply useless as guides to action for such agents (as per (3)). Thus Jackson and Smith (2006, 2016) and Huemer (2010), *inter alia*, have suggested that the problem of uncertainty poses a fatal objection at least to absolutist forms of deontology.

There may be many lines of response more promising than (1)-(3) that the deontologist can pursue in the face of this challenge. But the major focus of the recent literature has been on variants of the *threshold view*, according to which, for any objectively prohibited act type W there is some threshold t, $0 < t < 1$, such that it is always subjectively wrong for me to choose a practical option O if my credence that O would be an instance of W is greater than or equal to t.[7]

The threshold view faces an important and widely acknowledged technical objection, which we will address in §5. But it also faces much simpler worries about *motivation* that have yet to be seriously addressed. First, why should there be any credal threshold at which an act abruptly switches status, from subjectively permissible to subjectively impermissible? What plausibility can there be, say, in the idea that if I have .27 credence that going to the game would break a promise to my friend, there is nothing at all wrong with going, but if I have .28 credence, I am strictly prohibited from going? Second, and closely related, what could serve to make any *particular* credal threshold the right one? Why should the threshold be set *here* rather than

[7] Jackson and Smith (2006) treat the threshold view as the most natural response the absolutist deontologist can give to the problem of uncertainty, though as we will see in §5 they argue that it is subject to a conclusive objection. Hawley (2008) and Aboodi et al. (2008) both defend versions of the threshold view against this objection. Isaacs (2014) suggests that it is subjectively permissible for an agent to take a morally risky action only if she *knows* that that action would not be an instance of a deontologically prohibited type, which is a close cousin of the threshold view provided one accepts that belief above some credal threshold is a necessary condition for knowledge.

there? And how can we ever hope to *know* where it has been set?[8]

In the next two sections, I will propose an answer to these questions: namely, that *stochastic dominance reasoning* both explains why there should be sharp thresholds for subjective permissibility and fixes the value of those thresholds in a perfectly non-arbitrary manner, thereby vindicating a version of the threshold view.

3 Stochastic Dominance

The idea of dominance reasoning is familiar to philosophers from game-theoretic contexts like the prisoner's dilemma. If I am uncertain about the state of the world, but certain that, *given any possible state of the world*, option O is more choiceworthy than option P, then O is said to *strictly dominate* P. If I am certain (i) that, given any possible state of the world, O is *at least as choiceworthy* as P, and (ii) that given some state(s) of the world O is more choiceworthy, then O is said to *weakly dominate* P.

Stochastic dominance extends these familiar ideas as follows. A practical option O stochastically dominates an alternative P iff, relative to the agent's doxastic or epistemic state

1. For any degree of choiceworthiness d, the probability that O is (or will turn out to be) at least as choiceworthy as d is equal to or greater than the probability that P is ("") at least as choiceworthy as d, and

2. For some degree of choiceworthiness d, the probability that O is ("") at least as choiceworthy as d is strictly greater than the probability that P is ("") at least as choiceworthy as d.

An illustration: Suppose that I am going to flip a fair coin, and I offer you a choice of two tickets. The Heads ticket will pay you \$1 for heads and nothing for tails, while the Tails ticket will pay you \$2 for tails and nothing for heads. The Tails ticket neither strictly nor weakly dominates the Heads ticket because, if the coin lands Heads, the Heads ticket will yield a more desirable outcome. But the Tails ticket *does* stochastically dominate the Heads ticket. There are three possible outcomes of the game, which in ascending order of desirability are: winning \$0, winning \$1, and winning \$2. The two tickets offer the same probability of an outcome *at least as good as* \$0, namely 100%. Likewise, they offer the same probability of an outcome at

[8]This objection from arbitrariness is pressed by both (Portmore, 2016, pp. 10-11) and (Jackson and Smith, 2016, p. 284).

least as good as \$1, namely 50%. But the Tails ticket offers a better chance of an outcome at least as good as \$2, namely 50%, versus the 0% probability offered by the Heads ticket.

The principle that it is never rational to choose a stochastically dominated option is extremely compelling. First, note that unlike the stronger principle that a rational agent should always maximize expected utility or expected value, the stochastic dominance principle places no unwelcome constraints on an agent's attitude toward risk. For instance, if I am offered a choice between a ticket that pays \$1 if a fair coin lands heads on a single flip and a ticket that pays \$4 if a fair coin lands tails on both of two flips, it may be rational for me to take the risk averse option and select the first ticket even though the expected payoff of the second ticket is twice as large. Unlike an expectational principle, stochastic dominance is silent in this sort of case: the second ticket offers a better chance of a payoff at least as good as \$4 (namely, 25% rather than 0%), but the first ticket offers a better chance of a payoff at least as good as \$1 (namely, 50% rather than 25%), so neither option is stochastically dominated.[9] Relatedly (as we will see at greater length in the next section), stochastic dominance reasoning does not entail the implausible results of expectational reasoning with respect to very small chances of very large or infinite payoffs (as in Pascal's Wager), since paying some definite cost for a very small chance of a very large reward will not stochastically dominate the null option of forgoing both the cost and the chance of reward.

Given these observations, it is hard to see how one can ever make the case for selecting a stochastically dominated option. It is, in general, *possible* that the stochastically dominated option will yield a more desirable payoff than the dominant option, but whatever level of payoff one is most concerned with getting, the dominant option offers an equally good or better chance of a payoff at least that desirable. The *way* in which that payout comes about is, by definition, a matter of indifference: If, say, I preferred winning \$0 with a Heads ticket to winning \$0 with a Tails ticket in the first case of stochastic dominance reasoning given above, then the desirability of the possible outcomes would depend on more than the monetary payout, and the Tails ticket would no longer stochastically dominate the Heads ticket. If I do *not* have such a preference, then the mere fact that the Heads ticket *might* turn out to yield a more desirable outcome is not a sufficient reason for choosing it over the Tails ticket.[10]

[9] Of course, expected utility theory will only requires that you choose the riskier ticket if the payoffs are given in utiles or some other unit of pure value, rather than in dollars. But the point can be made just as easily in these terms, since it is commonly believed that it is sometimes rationally permissible to be risk averse with respect to utility.

[10] In more formal terms, a stochastically dominant option will be preferred to a stochastically dominated alternative by any agent whose utility function is monotonically increasing, i.e., for whom a given chance of a higher-value outcome is preferred to the same chance of a lower-value outcome, all else being equal (Hadar and Russell, 1969, p. 28). The notion of stochastic dominance defined above, and which I employ throughout the paper, is sometimes called *first-order* stochastic dominance, distinguishing it from higher-order stochastic

In the next section, I will show that the principle of eliminating stochastically dominated options can justify a threshold principle for choice under deontological moral uncertainty.

4 Stochastic Dominance and Deontological Uncertainty

Let's return to the pair of cases given in §2, in which I am unsure whether I am required to skip the Big Game and help my friend paint her house, either because I am empirically uncertain whether I told her I would help or because I am purely morally uncertain whether some past act constituted a morally binding promise.

Let's provisionally suppose the following about these cases: (1) The background deontological conception of morality that I accept is absolutist, and therefore treats the objective requirement to keep one's promises as lexically stronger than any non-moral reason (like the reasons stemming from my desire to see the game). (2) This means that (a) an act of promise keeping is *better* or *more choiceworthy* than any act that violates a moral obligation or is morally neutral (at least so long as the agent is appropriately motivated, e.g. by considerations of duty), and (b) an act of promise breaking is *worse* or *less choiceworthy* than any act that fulfills a moral obligation or is morally indifferent (at least so long as no special exculpatory condition applies).[11] (We will later consider what happens when we relax these assumptions.)

Accepting these assumptions allow a decision-theoretic representation of deontological moral considerations, along lines suggested by Colyvan, Cox, and Steele (2010). In their model, absolutist deontology is characterized by the following axiomatic extension of standard utility theory:

D1* If [outcome] O_{ij} is the result of an (absolutely) prohibited act, then any admissible utility function u must be such that $u(O_{ij}) = -\infty$.

D2* If [outcome] O_{ij} is the result of an (absolutely) obligatory act, then any admissible utility function u must be such that $u(O_{ij}) = +\infty$. (Colyvan et al., 2010, p. 512)

On this model, while actions from duty have "infinite" positive value and actions against duty have "infinite" negative value, actions that neither violate nor fulfill duties have finite value determined, presumably, by prudential or desire-based reasons and perhaps by other sorts of

dominance principles that place increasingly stringent constraints on an agent's risk attitudes.

[11] It may seem that we are some risk of describing deontological appraisal of actions in evaluative rather than normative terms, when part of what distinguishes deontological approaches to ethics is that they take normative notions to be primary, where consequentialists give primacy to the evaluative. But nothing turns on the choice of evaluative or normative language—we can, for instance, take the description of one action as "better" than another from the standpoint of a given agent as mere shorthand for the normative claim that it is more strongly supported by the agent's reasons, without raising any new difficulties for the argument given in this section.

moral reason (e.g. consequence-based reasons of benevolence) that do not generate absolute obligations or prohibitions.[12]

The mentions of infinite positive and negative value in Colyvan et al's axioms need not be taken too literally. As they point out (pp. 521ff), decision-theoretic models of non-consequentialist ethical theories may be descriptively adequate without being explanatory: A Kantian does not avoid lying, for instance, *because* she regards acts of lying as having infinite disvalue. Nevertheless, the fact that, for a Kantian, moral obligations and prohibitions are lexically stronger than prudential reasons can be accurately *represented* by treating the value of an action from duty as the upper bound on the scale of reason strength, and the disvalue of an action against duty as the lower bound.

This representation in hand, we can employ stochastic dominance reasoning to draw conclusions about how a committed deontologist ought to respond to moral uncertainty. Consider the case of the Possible Promise. Suppose that the prudential value of seeing the Big Game is +20, while the prudential value of helping Petunia paint is +5 (and that I know these prudential facts with certainty). It follows that the option of helping Petunia will stochastically dominate the option of going to the game if and only if the probability that I promised to help Petunia is greater than or equal to .5.

Suppose, first, that the probability is exactly .5, that is, I regard it as equally likely that I did as that I did not make a promise to Petunia. Then I may reason as follows. My decision has four possible outcomes (Table 1), which in order from best to worst are: (i) I did make a promise to Petunia, but I go to the game, violating that promise ($-\infty$). (ii) I did not make a promise to Petunia, but skip the game to help her paint anyway (+5). (iii) I did not make a promise to Petunia, and go to the game, violating no obligation (+20). (iv) I did make a promise to Petunia, and I skip the game to help her paint, thereby fulfilling that obligation ($+\infty$).

Given my credences, helping Petunia paint (option P) stochastically dominates going to the game (option G): P and G have the same chance of producing an outcome at least as good as $-\infty$ (100%); P has a better chance of producing an outcome at least as good as +5 (100% vs. 50%); P and G have the same chance of producing an outcome at least as good as +20 (50%); and P has a better chance of producing an outcome at least as good as $+\infty$ (50% vs. 0%) (Table 2). Thus, if I accept the principle that I should never choose a stochastically dominated option, I am compelled to choose P rather than G.[13]

[12]Talk of the "value" or "utility" of actions should be understood as simply a less cumbersome stand-in for talk of *choiceworthiness*—i.e., it should be understood as denoting a normative rather than an evaluative property of actions.

[13]The talk of "outcomes" should not be seen as illicitly consequentialist. Violating or fulfilling a deontological

	Promised (.5)	**Didn't Promise (.5)**
Painting	$+\infty$	$+5$
Game	$-\infty$	$+20$

Table 1: Possible Promise v1, payoff matrix

	$\mathbf{-\infty}$	$\mathbf{+5}$	$\mathbf{+20}$	$\mathbf{+\infty}$
Painting	1	1	.5	.5
Game	1	.5	.5	0

Table 2: Possible Promise v1, probability of payoff $\geq x$

	Promised (.49)	**Didn't Promise (.51)**
Painting	$+\infty$	$+5$
Game	$-\infty$	$+20$

Table 3: Possible Promise v2, payoff matrix

	$\mathbf{-\infty}$	$\mathbf{+5}$	$\mathbf{+20}$	$\mathbf{+\infty}$
Painting	1	1	.49	.49
Game	1	.51	.51	0

Table 4: Possible Promise v2, probability of payoff $\geq x$

On the other hand, suppose that my credence that I promised to help Petunia paint is only .49. In that case, option G has a better chance than option P of yielding an outcome at least as good as +20 (51% vs. 49%), so P does not stochastically dominate G (Tables 3-4).

Notice that this argument applies to Dubious Promise, the case of pure moral uncertainty, just as it applies to Possible Promise, the case of empirically based moral uncertainty. In either case, the option of helping my friend will stochastically dominate the option of going to the game iff my credence that I am objectively morally obligated to help my friend is greater than or equal to .5. Thus, it seems, stochastic dominance reasoning gives the deontologist something definite, precise, and well-motivated to say *both* about choices made under morally relevant empirical uncertainty *and* about choices made under pure moral uncertainty.[14]

An immediate worry: We have so far assumed that the background deontological conception is absolutist. But many philosophers who identify as deontologists and endorse characteristically deontological moral phenomena like agent-centered constraints do *not* think of deontological constraints as absolute.[15] Fortunately, however, the stochastic dominance argument does not

obligation is an "outcome" only in the formal sense of being a distinct act-state combination valued differently than other act-state combinations. Calling these "outcomes" does not imply, for instance, that the wrongness of breaking a promise has anything to do with its causal consequences.

[14]Interestingly, the view that it is permissible to take a morally risky act only when the probability that the act is not objectively wrong is greater than .5 has a history in the Catholic moral theology literature, under the name "probabiliorism" (Sepielli, 2010, pp. 51-2). Stochastic dominance, in these terms, implies that a deontological absolutist must adopt a position *at least* as rigorous as probabiliorism in the cases of "asymmetrical" moral risk with which this literature is chiefly concerned.

[15]The literature on deontological moral uncertainty is divided in focus between absolute and non-absolute versions of deontology. Jackson and Smith (2006) direct their criticisms at absolutist theories. Aboodi et al.

turn on the assumption of absolutism. Following Colyvan et al, who conclude that the absolutist version of deontology is implausible, we may represent the rightness of fulfilling an obligation not as $+\infty$ but as merely a very large finite positive number, i.e., as a reason of finite strength though much stronger than ordinary prudential reasons. And likewise, we may represent the wrongness of violating an obligation not as $-\infty$ but as merely a very large finite negative number.[16] Suppose that keeping a promise to my friend has a value of $+9001$ while breaking that a promise has a value of -9001. It is easy to see that the stochastic dominance argument for helping Petunia paint will go through *mutatis mutandis*.

Note, however, that a reliance on stochastic dominance rather than expectational reasoning lets us avoid one of the chief pitfalls for absolutist approaches to uncertainty, namely, the risk of "Pascalian paralysis": If violating a moral obligation is treated as an outcome with the value of $-\infty$, then actions that are almost certainly morally permissible but carry even a vanishingly small risk of violating a moral obligation will carry an expected value of $-\infty$—the same expected value, in fact, as acts that are *certain* to violate a moral obligation.[17] Worse still, if (as seems plausible to me) *every* possible action has some non-zero chance of fulfilling a moral obligation and some non-zero chance of violating a moral obligation, then the expected value of every possible action is undefined ($\infty + (-\infty)$).

One way to avoid Pascalian difficulties is to hold that a rational agent may never have non-

(2008) suggest that this is a mistake, since "hardly any (secular) contemporary deontologist is an absolutist" (p. 261, n5). But Huemer (2010, p. 348, n3) marshals a credible array of apparent (contemporary or near-contemporary) deontological absolutists. Note that the difficulties of absolutism arise for anyone who takes one class of reasons to be lexically stronger than another: assuming that there is always, in any choice situation, a non-zero risk that the lexically stronger reasons are in play (e.g. that a given course of action *might* cause some innocent person torturous suffering), one may worry that the lexically weaker reasons (e.g. headache prevention) will always be preempted by that risk and hence never rise to the level of practical relevance.

But in any case, the challenge posed by uncertainty is no less acute for non-absolute deontologists. The non-absolutist must still answer the question of how an agent should decide what to do when she is unsure whether a given course of action would violate a deontological constraint. As with the absolutist, the most natural answer she can give is some version of the threshold view, but to defend this view she will have to overcome all the same obstacles as the absolutist (Jackson and Smith, 2016, pp. 287-8). The only new option that the non-absolutist position seems to open up is an expectational view on which the subjective reason-giving force of a deontological constraint for an uncertain agent is the product of the inherent stringency of that constraint (which is finite, since the constraint is non-absolute) times the probability that a given course of action would violate the constraint. But if she goes this expectational route, it is no longer clear that her view should be counted as a form of deontology, rather than simply a version of consequentialism that incorporates agent-centered constraints. (Cf. discussion in Portmore (2016). Portmore allows that a deontologist may give an expectational account of subjective rightness without becoming a consequentialist, so long as the underlying account of *objective* rightness remains deontological. But he argues convincingly that the challenges of uncertainty arise not just in cases where we are ignorant of facts (e.g., what the consequences of some action will be or would have been) but also in cases where there is *no fact of the matter* about whether some course of action would have constituted a constraint violation (e.g., due to physical indeterminism or counterfactual underdetermination), and that therefore the deontologist who wishes to go the expectational route must give a partially expectational account of *objective* as well as subjective rightness, which he takes it *would* amount to a form of consequentialism. For a different line of argument against expectational approaches to deontological moral uncertainty, see Tenenbaum (2017).)

[16] See Colyvan et al (2010, pp. 515-8) for an axiomatic characterization of such a non-absolutist deontological theory.

[17] In the literature on "pure" moral uncertainty, this is often described as the worry that "fanatical" moral theories will hijack expectational reasoning—i.e., moral theories that attribute infinite value and disvalue to options, perhaps in very strange or counterintuitive ways, will take precedence over all finitary moral theories, so long as one assigns them even the most vanishingly small degree of positive credence (Ross, 2006, pp. 765-7).

zero credence in any outcome with infinite positive or negative value. But this seems implausible and has only *ad hoc* motivation. The better response is to modify our decision theory, either weakening or amending expected utility theory in a way that allows an agent with modest credence in infinite values and disvalues to nevertheless remain responsive to finitary considerations. Stochastic dominance is one such weakening of expected utility theory: As far as stochastic dominance principles are concerned, it is rationally permissible for me to go to the game, despite the risk of infinite moral turpitude, so long as that risk has a probability less than .5.[18]

The .5 credal threshold for permissibility follows from stochastic dominance reasoning when an agent is certain of all her non-moral reasons. But when she is uncertain which option her non-moral reasons favor, stochastic dominance may become more demanding. Suppose, for instance, that in the Possible Promise case I am uncertain whether I would have a better time at the game or painting with my friend. Perhaps I believe that there is a one-in-three chance that my team will lose, and I know that while seeing my team win would give me a utility of $+20$, seeing them lose would give me a utility of -10. If I simply ignore the game and help my friend paint, on the other hand, I am guaranteed a utility of $+5$. In this case, the degree of belief that I promised my afternoon to Petunia at which stochastic dominance will require that I do so is reduced—specifically, to .4 rather than .5.[19] But in simple cases of deontological moral-prudential conflict, where I am certain that one option O is morally permissible but prudentially worse than another option P, and certain that P is prudentially better than O but uncertain whether it is morally permissible, we may say that stochastic dominance requires me to choose option O iff the probability that P is objectively impermissible is greater than or equal to .5.[20]

Of course, even allowing for the possibility of a higher threshold in cases of prudential uncertainty, the requirements of stochastic dominance will not rise to the level of practical stringency that we intuitively expect at least some deontological constraints to possess. The kind of rea-

[18]Of course, stochastic dominance need not be understood as the *only* or the *strongest* principle of rational requirement under uncertainty (cf. the last three paragraphs of this section). Various principles are possible that occupy an intermediate position between stochastic dominance and expected utility maximization and that might impose a more rigorous form of moral caution on agents acting under uncertainty while still avoiding the problem of Pascalian paralysis.

[19]Suppose my credence that I promised to help Petunia paint is exactly .4. There are five possible outcomes, with values of $-\infty$, -10, $+5$, $+20$, and $+\infty$. Options P and G have the same chance of producing an outcome at least as good as $-\infty$ (100%); P has a better chance of producing at outcome at least as good as -10 (100% vs. 60%); P has a better chance of producing an outcome at least as good as $+5$ (100% vs. 40%); P and G have the same chance of producing an outcome at least as good as $+20$ (40% — the chance that G produces an outcome at least this good is a product of the .6 chance that I do not violate an obligation in going to the game and the $.\overline{6}$ chance that my team wins the game); and P has a better chance than G of producing an outcome at least as good as $+\infty$ (40% vs. 0%) (Tables 5-6). If, however, my credence that I made the promise were any lower, then G would have a better chance than P of producing an outcome at least as good as $+20$ and P would no longer stochastically dominate G.

[20]Cases of this sort have been a major focus of the recent literature on "pure" moral uncertainty, in particular the cases of abortion and vegetarianism, both of which seem to present agents (at least in many cases) with a conflict of prudential reasons on one side and uncertain moral reasons on the other. See for instance Guerrero (2007), Moller (2011), and Weatherson (2014) for discussion of these cases.

	Prms (.4)	**¬Prms & Win (.4)**	**¬Prms & Lose (.2)**
Painting	$+\infty$	$+5$	$+5$
Game	$-\infty$	$+20$	-10

Table 5: Possible Promise v3, payoff matrix

	$-\infty$	-10	$+5$	$+20$	$+\infty$
Painting	1	1	1	.4	.4
Game	1	.6	.4	.4	0

Table 6: Possible Promise v3, probability of payoff $\geq x$

soning I have described makes no distinction between "weaker" constraints (like the constraint against telling white lies) and "stronger" constraints (like the constraint against killing the innocent).[21] And while a threshold of .5 may seem plausible for the weaker constraints, it seems much less plausible when we think about the stronger constraints—surely a deontologist should not conclude, for instance, that a judge need not worry about the constraint against punishing the innocent so long as the probability that her sentencing decision violates that constraint is a mere .49 rather than .51.

Importantly, however, stochastic dominance is only a *sufficient* condition for rational requirement, not a *necessary* condition. Everything I have said so far leaves it open to the deontologist to argue for a higher threshold, at least for some categories of constraint, on grounds other than stochastic dominance. And even if one takes the laxity of the stochastic dominance threshold to be intuitively unacceptable with respect to any sort of deontological constraint, progress has still been made insofar as the deontologist can now say with confidence that you must *at least* believe it to be more likely than not that your morally risky action violates no constraint, even a relatively weak constraint, before non-deontological (e.g. prudential or consequentialist) reasoning may take over.

Nevertheless, it seems to me that this objection from the variable stringency of constraints represents a residual difficulty that it will be hard for deontologists to fully overcome. It is intuitive to hold that the practical force of deontological constraints must be sensitive to both (i) the *probability* of a constraint being violated by a given course of action and also (ii) the *seriousness*, *importance*, or *stringency* of the particular constraint in question. But taken together, these intuitions push us in the direction of an expectational view that, while it may still incorporate agent-centered normative considerations, has lost the rest of its distinctively deontological character (see note 15 *supra*). The deontologist may resist this pressure, but to do

[21] Except, on a non-absolutist view, when there is some chance that the constraint in question is outweighed by other kinds of considerations—since the chance of a stronger constraint being outweighed will presumably be smaller, *ceteris paribus*, the threshold of moral safety at which the option that risks violating that constraint ceases to be stochastically dominated will be higher.

so she must be prepared at some point to deny at least one of the above intuitions.

5 Option Individuation and Ought Agglomeration

So far I have argued that stochastic dominance reasoning offers a foundation for a threshold view of deontological moral choice under both empirically based and pure moral uncertainty. But such threshold views face another important difficulty, which we have yet to confront.

Jackson and Smith (2006) argue that threshold views violate the principle of "ought" agglomeration, that $ought(A)$ & $ought(B) \vdash ought(A \,\&\, B)$, in cases where an agent is faced with the possibility of performing two acts, each of which individually is below the threshold for permissible moral risk, but which in combination exceed that threshold. They illustrate this difficulty by way of the following thought experiment.

Two Skiers Two skiers are headed down a mountain slope, along different paths. Each, if allowed to continue, will trigger an avalanche that will kill ten innocent people. (If both are allowed to continue, they will trigger two such avalanches and kill both groups of ten.) The only way to save each group is to shoot the corresponding skier dead with your sniper rifle. You can shoot either skier individually or, being an extraordinarily crack shot, you can shoot both with a single bullet. The moral theory you accept (with certainty) tells you that you ought to kill culpable aggressors in other-defense, but are absolutely prohibited from killing innocent threats. Unfortunately, you are uncertain of the intentions of the skiers, assigning them each equal and independent probabilities of acting obliviously, and thus as innocent threats, rather than as ill-intentioned aggressors.

Suppose you subscribe to a threshold view on which you ought to kill a potential aggressor in other-defense iff your credence that he is acting innocently is less than t. And suppose that, with respect to each of the two skiers, your credence that he is innocent is *just* less than t. Thus, it seems that you ought to shoot Skier 1 and you ought to shoot Skier 2. But, if you shoot *both* Skier 1 and Skier 2, the chance that *one* of them is innocent and hence that you will have violated a deontological constraint is *greater* than t and hence, it seems, you ought not perform the compound action: *shoot Skier 1 and shoot Skier 2*.

It is immediately tempting to point out that *shooting Skier 1* and *shooting Skier 2* are two separate actions, and that all that should matter about the case from a moral standpoint is whether the chance of either action individually violating a deontological constraint is greater than t. And I will shortly argue that something very much like this is the right response for

the deontologist to give. But Jackson and Smith seem to have headed off this response: While shooting each skier separately looks like two separate actions, shooting both skiers *with the same bullet*, as they have imagined you have the option of doing, is a single action, with a chance greater than t of violating a deontological constraint. But surely it is implausible that you are permitted (indeed, required) to shoot both skiers with separate bullets, but prohibited from shooting them both with the same bullet.

A plausible solution to the agglomeration problem must therefore find some appropriate criterion for "individuating" the shooting of Skier 1 and the shooting of Skier 2 that allows us to distinguish the two sources of moral risk, even given the option of shooting both skiers with a single bullet. Aboodi, Borer, and Enoch (2008) offer one such proposal. They suggest that, if all deontological constraints can be thought of as *rights* possessed by particular rightsholders, a deontological agent acting under uncertainty ought to ensure that she does not run a risk greater than the threshold t of violating the risks of any *particular* rightsholder, though she may permissibly run a risk greater than t of violating the rights of *some* rightsholder. They suggest this might be justified, *inter alia*, by the broadly contractualist thought that what determines the permissibility of my conduct is whether anyone can reasonably or legitimately object to it: A rightsholder cannot reasonably object to my conduct merely because it creates *some* risk that her rights will be violated, nor because it creates a large aggregate risk that *someone's* rights will be violated. But she can reasonably object if I run an excessive risk of violating *her rights in particular*. Thus, Aboodi et al conclude, in the Two Skiers case you ought to shoot both skiers, because in so doing there is no rightsholder that you run a risk greater than t of wronging. (Skier 1 and Skier 2 are each rightsholders, but there is no compound entity, Skier 1 + Skier 2, that is a rightsholder or can lodge objections against your conduct.)

This patient-based approach, however, has seriously counterintuitive consequences, as Huemer (2010) points out. Suppose, as I have argued holds true in the simplest sort of case, that the threshold for permissible risk of violating a deontological constraint is .5, and consider the following case (a slightly simplified retelling of Huemer's "War Options").

The Weapon Minerva, a military officer, faces a situation in which she knows that, if she does not act, a powerful weapon will fall into the hands of a ruthless enemy, who will use it to kill 100,000 innocent people. Fortunately, there are two surefire ways to prevent the weapon from falling into enemy hands: (1) Minerva has received intelligence indicating that a certain scientist will shortly reveal the plans for constructing the weapon to the enemy. It is unclear whether the scientist is willingly helping the enemy or acting under exculpating

duress (say, threats on the life of his family). Minerva thinks it slightly more likely than not (.55) that he is acting under duress and that it would therefore be objectively wrong to kill him even to save a greater number. (2) In addition to the weapon plans, the enemy needs certain materials to construct the weapon, which are being kept in a small town of 25,000 people. Minerva knows that if she orders the carpet bombing of this town, the weapon materials will be destroyed, but 45% of the town's population will be killed. Since she has no specific information about which members of the town's population will be at risk, the epistemic probability that any individual townsperson will be killed is .45.

On Aboodi et al's patient-centered approach, it looks as though Minerva is permitted to carpet bomb the town (since there is no rightsholder whom she runs a risk greater than .5 of violating) but is not permitted to assassinate the scientist (since in doing so she would run a risk greater than .5 of wronging him).[22]

Fortunately, there is another way of resolving the agglomeration problem that both avoids these counterintuitive consequences and follows naturally from a central feature of the deontological approach to ethics: Because deontological moral assessment of options is relativized to the particular *choice situations* in which those options arise, it seems to me, deontological moral theories simply cannot and do not assess *compound* options, i.e., combinations of options that arise in different choice situations. Thus, while *shoot Skier 1* and *shoot Skier 2* are each individually subject to deontological moral assessment, the compound option *shoot Skier 1 & shoot Skier 2* is not. But since it is a theory's moral assessment of options that provides the basis for stochastic dominance reasoning, which in turn grounds the threshold principle for deontological moral risk, this means that the threshold principle is simply inapplicable to compound options, and hence does not generate a prohibition against shooting both skiers. Rather, we are left free to make the natural inference from the fact that you ought to shoot Skier 1 and ought to shoot Skier 2 to the conclusion that you ought to shoot both skiers, preserving "ought" agglomeration.

Let's spell out this line of reasoning in a bit more detail. Deontological moral assessment is relativized to choice situations in the following sense: Just as deontological theories (in contrast, for instance, to agent-neutral consequentialist theories like classical utilitarianism) direct me to distinguish between the rights and wrongs of my own actions and those of others, and hence, for instance, not to commit one rights violation to prevent someone else from committing five, so likewise they direct me to distinguish between the rights and wrongs *of the choice immediately before me* and the rights and wrongs of my other choices, past, present, or future. Thus for

[22] Of course, the particular value of the threshold plays no role in this argument—the case can be easily modified to accommodate any threshold greater than or less than .5.

instance a deontological theory does not direct me to lie now even if it is the only way of extricating myself from a situation in which I can predict with certainty that I will tell five lies. To put this in terms of the quantitative representations we employed in the last section, a deontological theory will assign the option of lying now the large negative value (either finite or infinite) associated with violating a constraint, but will not assign this same degree of disvalue to the option of not lying, even though this option carries with it the certainty of future constraint-violating lies, which will themselves be assigned the disvalue of constraint violations in the choice situations where they arise.[23]

Colyvan et al. (2010) put this point by suggesting that deontological moral duties are both agent-relative and *time*-relative (p. 513), but this is not quite right, since one can make multiple choices at the same time, and in such cases the relativization of deontological considerations remains to the individual choice situation, rather than to the time at which both choices are made.[24] For instance, consider the following case:

The Buttons of Wrongness A mad ethicist has rigged a contraption that will test your judgments about two ethical dilemmas at once. You may press one of four buttons, colored red, green, yellow, and blue. If you press the red button, a message will be sent from your phone to one of your friends, telling him a white lie about his recent haircut. If you press the green button, $1000 will be stolen from the accounts of a large corporation and donated to GiveDirectly. If you press the yellow button, *both* these things will happen. If you press the blue button, *neither* will happen. Pressing any of the buttons will deactivate the rest, and you will have no other opportunity to reverse the effects of your selection once it has been made.

There is a sense, clearly, in which you make only one choice in this scenario, namely *which button to press*. But in a more ethically relevant sense, you make two choices, namely *whether to tell a white lie* and *whether to steal money from the large corporation for GiveDirectly*—even though you will put both of these choices into effect by means of a single action, namely, pressing a button. And to the deontologist, what you will do with respect to one of these choices simply has no bearing on what you are permitted to do with respect to the other—not because they are potential violations of two different rightsholders (it would make no difference if your friend with the new haircut was also the majority shareholder of the large corporation) but simply because

[23] A deontological theory should likely assign *some* disvalue to options that I know will lead me to future constraint violations, but the point is that it does not assign the same kind or degree of disvalue that it assigns to options that directly violate a constraint.

[24] In fairness, Colyvan et al may not have meant to assert otherwise, since later on the same page they put the point in terms of choice situations rather than times.

they are *two different choices*.[25]

What does this mean for Jackson and Smith's Two Skiers case and the agglomeration objection to threshold views? Simply this: Because deontological moral assessment is choice situation relative, deontological moral theories are just not in the business of assessing compound options like *shoot Skier 1 & shoot Skier 2*. Just like the options of telling your friend a white lie and stealing for GiveDirectly in the Buttons of Wrongness case, these options inhabit different choice situations, *even if* they are (or can be) effectuated by means of a single action (viz., shooting both skiers with the same bullet). For this reason, we cannot say for instance that the option of shooting neither skier stochastically dominates the option of shooting both skiers: Stochastic dominance reasoning makes reference to a probabilistic assignment of degrees of objective value or choiceworthiness to options, but since deontological theories don't assign degrees of value or choiceworthiness to compound options, they can be neither stochastically dominant nor dominated. Since the threshold principle of choice under deontological moral uncertainty is grounded in stochastic dominance reasoning, this principle in turn is only applicable to simple options. Thus in the Two Skiers case, the threshold principle does not imply that you ought not shoot both skiers. Rather, we can and should reason, by the principle of agglomeration, from the fact that you ought to shoot Skier 1 and ought to shoot Skier 2 to the conclusion that you ought to shoot both skiers.

6 Conclusion

Moral uncertainty, of both the empirically based and purely moral varieties, presents a challenge for deontologists. They must give an account of how agents should deliberate and act in the face of such uncertainties, ideally an account that is as precise and intuitively well-motivated as the expectation-maximizing account available to consequentialists. I have suggested that a version of the threshold account grounded in stochastic dominance reasoning may meet this need. If we treat morally right/obligatory actions as having very large positive moral value (whether finite

[25]This is not to say that deontologists are committed to metaethical possibilism, the view that what I ought to do with respect to a particular choice situation can *never* depend on what I believe I will do with respect to other choice situations. For instance, the deontologist might still allow that in the classic test case for actualism vs. possibilism (Jackson and Pargetter, 1986, p. 235), Professor Procrastinate should decline the invitation to review a recently published book, even though he knows himself to be the person best qualified to review it, because he also knows that given his inveterate tendency to procrastinate, he will not finish the review on time. The point is rather the narrower one that *the strength of a deontological moral prohibition or requirement* is not altered by one's expectations about how one will behave in other choice situations. (Thus, what allows the deontologist to take the actualist line with respect to Professor Procrastinate is that reviewing the book is, one assumes, morally optional. If it were the case, for whatever reason, that Professor Procrastinate were morally obligated to agree to review the book, then the deontologist should agree with the possibilist that he ought so to agree, whatever his expectations about any of his other choices or actions.)

or infinite) and morally wrong/prohibited actions as having very large negative moral value (“”), then it will often turn out that morally risky options are stochastically dominated. In the simplest sort of case, where only one of two options carries moral risk, and that option is certain to be preferable *but for* the risk of violating a moral requirement, it turns out that that option is stochastically dominated if and only if the probability that it is objectively prohibited is greater than or equal to .5, though this threshold may be less than .5 in more complex cases. The principle that one rationally ought not choose stochastically dominated options is extremely compelling, representing a weakening of expectational reasoning that both permits a variety of risk attitudes (risk aversion and risk seeking as well as risk neutrality) and avoids worries about moral fanaticism or Pascalian paralysis arising from tiny probabilities of infinite value and disvalue. It therefore provides an appealing basis for the threshold approach to deontological moral uncertainty, one that moreover (when combined with the right understanding of deontological moral evaluation) can avoid the most powerful objection to that approach. It seems, then, that deontologists are better equipped to respond to the challenges of moral uncertainty than many philosophers have thought.

References

Aboodi, R., A. Borer, and D. Enoch (2008). Deontology, individualism, and uncertainty: A reply to Jackson and Smith. *Journal of Philosophy 105*(5), 259–272.

Colyvan, M., D. Cox, and K. Steele (2010). Modelling the moral dimension of decisions. *Noûs 44*(3), 503–529.

Guerrero, A. A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies 136*(1), 59–97.

Hadar, J. and W. R. Russell (1969). Rules for ordering uncertain prospects. *The American Economic Review 59*(1), 25–34.

Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 10. Oxford University Press.

Hawley, P. (2008). Moral absolutism defended. *Journal of Philosophy 105*(5), 273–275.

Hedden, B. (2016). Does MITE make right? On decision-making under normative uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 11. Oxford University Press.

Huemer, M. (2010). Lexical priority and the problem of risk. *Pacific Philosophical Quarterly 91*(3), 332–351.

Isaacs, Y. (2014). Duty and knowledge. *Philosophical Perspectives 28*(1), 95–110.

Jackson, F. and R. Pargetter (1986). Oughts, options, and actualism. *Philosophical Review 95*(2), 233–255.

Jackson, F. and M. Smith (2006). Absolutist moral theories and uncertainty. *Journal of Philosophy 103*(6), 267–283.

Jackson, F. and M. Smith (2016). The implementation problem for deontology. In E. Lord and B. Maguire (Eds.), *Weighing Reasons*, pp. 279– 291. Oxford University Press.

Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. Oxford University Press.

MacAskill, W. (2014). *Normative Uncertainty*. Ph. D. thesis, University of Oxford.

Moller, D. (2011). Abortion and Moral Risk. *Philosophy 86*(03), 425–443.

Portmore, D. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press USA.

Portmore, D. W. (2016). Uncertainty, indeterminacy, and agent-centered constraints. *Australasian Journal of Philosophy*, 1–15.

Ross, J. (2006). Rejecting Ethical Deflationism. *Ethics 116*(4), 742–768.

Sepielli, A. (2009). What to Do When You Don't Know What to Do. *Oxford Studies in Metaethics 4*, 5–28.

Sepielli, A. (2010). *'Along an Imperfectly-Lighted Path': Practical Rationality and Normative Uncertainty*. Ph. D. thesis, Rutgers University Graduate School - New Brunswick.

Tenenbaum, S. (2017). Action, deontology, and risk: Against the multiplicative model. *Ethics 123*(3), 674–707.

Weatherson, B. (2014). Running Risks Morally. *Philosophical Studies 167*(1), 141–163.