https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

DECISIONAL NONCONSEQUENTIALISM AND THE RISK SENSITIVITY OF OBLIGATION*

By Horacio Spector

Abstract: *A good deal of contemporary moral nonconsequentialism assumes that agents have perfect knowledge about the various features and consequences of their options. This assumption is unrealistic. More often than not, moral agents can only assess with a certain degree of probability the factual circumstances that are morally relevant for their decision making. My aim in this essay is to discuss the problem of moral decisions under risk from the point of view of nonconsequentialism. Basically, I analyze how objective moral principles can be transformed into subjective, decisional prescriptions, and argue that the standard nonconsequentialist approach to moral decision making, which focuses on probability thresholds, is wrong. In accordance with the fundamental postulates of nonconsequentialism, I seek to solve the problem of risk in moral choice by proposing a theory about the marginal moral value of various options. Actions can vary along various dimensions, and each of these dimensions can offer a different moral value function. Nonconsequentialist marginalism can level the playing field with consequentialism. Whereas consequentialism can simply borrow the notion of expected utility from economics, nonconsequentialism must introduce the notion of expectational obligation to formulate a general principle of moral choice under risk. I finally suggest that further empirical work is needed to delineate the shape of various moral value functions that are critical for applying the general principle of moral decision making under risk to well-known cases.*

My aim in this essay is to show that there is a version of nonconsequentialism that can successfully manage the problem of moral choice under risk. The central idea is that, just as consequentialism has a decision-theoretic counterpart, so too must nonconsequentialism develop an associated theory that we can call *decision-theoretic nonconsequentialism*, or, for short, *decisional nonconsequentialism*. This theory must include: (1) an account of moral inference that formalizes the derivation of subjective ought statements applicable to concrete practical contexts from objective moral propositions, and (2) a coherent explanation of how decisional nonconsequentialism can supply conclusions about the moral status of risk-imposing actions that (a) avoids merely probabilistic boundaries between rightness and wrongness, and (b) is consonant with the nonconsequentialist tenet that rightness criteria must be grounded in the moral value of various kinds of action.

* I am grateful for helpful comments by Eleonora Cresto, Martin Farrell, Frances Myrna Kamm, Benjamin Kiesewetter, Hartmut Kliemt, Eduardo Rivera Lopez, Thomas Schmidt, Ted Sichelman, and other members of the audience at the Institute of Philosophy at the Humboldt University of Berlin. I am also indebted for excellent points made by an anonymous referee of *Social Philosophy and Policy*.

doi:10.1017/S0265052516000121
© 2016 Social Philosophy & Policy Foundation. Printed in the USA.

I. Moral Theory With Epistemic Constraints

Consider your right to defend yourself against aggression. Moral rights such as self-defense can be construed in an objectivist or subjectivist way. For *objectivism*, whether you are allowed to use violence to counteract an attack depends on the objective features of your situation; so your violent response is right if it is needed and effective to neutralize the attack (and its extent is proportional to the gravity of the attack). For *subjectivism*, your beliefs matter as well to determine the permissibility of the self-defensive action. Most forms of self-defense only admit of a subjectivist interpretation; so if you have to respond to a threat of attack, your defensive behavior is permissible only if you justifiably believe that it is necessary, proportional and effective to prevent a very likely attack.

In nonconsequentialist theory, moral principles are formulated as if they belonged to an *epistemically ideal morality*, that is, a moral system that is meant to bear on omniscient moral agents. In a world of perfect factual knowledge, the dictates of morality apply objectively to the relevant features of decisional contexts, and agents' beliefs need not be examined (because they are assumed to be complete and true).[1] However, full information is seldom available. Incomplete information can adopt two forms. First, we might have a gap of information in the sense that we do not have evidence about the truth of a proposition.[2] Second, we might have a body of evidence about the truth of a proposition that only warrants a certain degree of epistemic confidence, often called *credence*. Credences are expressed in terms of subjective probabilities that satisfy the axioms of probability theory. We can call the former situation, "uncertainty," and the latter, "risk."[3] Both uncertainty and risk are subjective notions because they represent our relative degrees of information and ignorance. However, in this essay I will only focus on risk and its impact on moral choice.

Since long ago, criminal lawyers have realized that an innocent defendant's right not to be punished cannot be construed objectively if legal punishment is ever to be applied. Clearly, whatever evidentiary procedures are in place, courts must cope with various degrees of risk in condemning or acquitting suspects. For instance, Cesare Beccaria said in 1764:

> It may seem extraordinary that I speak of probability with regard to crimes, which to deserve a punishment, must be certain. But this paradox will vanish when it is considered, that, strictly speaking, moral

[1] I have discussed this idea in: Horacio Spector, "The Irrelevance of Ideal Morality to International Law," *Analisi e Diritto* (2013): 209–224.

[2] Carlos Alchourron and Eugenio Bulygin, *Normative Systems* (Vienna: Springer-Verlag, 1971), 33.

[3] Frank Knight, *Risk, Uncertainty and Profit* (New York: Augustus M. Kelley, 1964), 46. Note that "certainty" is opposed to both "uncertainty" and "risk" (just as "certain" is opposed both to "uncertain" and "risky").

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

> certainty is only probability, but which is called a certainty, because every man in his senses assents to it from an habit produced by the necessity of acting, and which is anterior to all speculation. That certainty which is necessary to decide that the accused is guilty is the very same which determines every man in the most important transactions of his life.[4]

Beccaria's point is that "legal certainty" (or "moral certainty" by the same token) is just a very high degree of probability that is taken to be good enough for practical purposes; in this case, the purpose is punishing a defendant. Probabilistic notions are not only necessary to analyze an innocent's right not to be punished. The problem of how obligations are sensitive to risk is relevant to all moral obligations, rights, and prohibitions. Any moral principle can be objectively applied if the *minimally sufficient set of factual conditions* for its applicability is true. But in practice this set can seldom, if ever, be known for sure, and so decisions under every principle must often be taken under risk. For instance, in international law the justification of self-defense, in terms of its necessity, proportionality, and effectiveness, relies on factual assumptions that can be established, if at all, only with various levels of epistemic confidence short of certainty.

Every moral theory must contemplate situations in which the moral agent lacks perfect knowledge about the relevant features of her options and potential outcomes. Consequentialists often feel reassured in their views by assuming that nonconsequentialist theories do not offer a clear solution to the problem of risk-taking in moral choice. In support of that assumption they mention various difficulties that result from trying to integrate probabilities with nonconsequentialist postulates. But the main problem is rarely made explicit. Standard systems of deontic logic do not allow attaching numerical values to deontic operators. Since nonconsequentialist theories are expressed in deontic terms, their normative modalities lack numerical values. In contrast, consequentialist theories often incorporate numerical measures of epistemic confidence as weights for utility that comply with the axioms of probability.[5] I think that the difficulties for nonconsequentialists of dealing with the problem of risk are caused by their failure to produce a decisional version of their theory. Decisional nonconsequentialism should be able to treat the problem of risk-taking at least as effectively as consequentialism does, and maintain at the same time its theoretical distinctiveness. (Obviously, nonconsequentialism would slide into consequentialism if it just resorted to the trick of treating acts and omissions as consequences.)

[4] Cesare Beccaria, *An essay on crimes and punishments*, 2nd ed., trans. Edward D. Ingraham (Philadelphia: Philip H. Nicklin, 1819), available at http:/www.constitution.org/cb/crim_pun.htm

[5] This is true of consequentialist models of weighing goods. See, for instance, John Broome, *Weighing Goods: Equality, Uncertainty and Time* (Oxford: Blackwell, 1991).

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

It should be possible, thus, to level the playing field between nonconsequentialist and consequentialist accounts of risk-taking once we introduce adequate notational, conceptual, and logical resources. My goal is to write a condensed version of the "second part" of nonconsequentialism — that part of nonconsequentialism that merges epistemic considerations with moral propositions in order to provide concrete answers to moral questions in real conditions of partial ignorance. It is not that nonconsequentialists have always overlooked the moral problem of risk-imposition. For instance, Judith Jarvis Thomson says: "The central problem which risk-imposition raises for moral theory is this: which instances of risk-imposition are wrongful, and which are not."[6] But proposals to solve "the central problem" have been unsystematic. As is well known, nonconsequentialists find it difficult to uphold a side constraint when abiding by it causes an immense loss of welfare. The typical nonconsequentialist response to the problem of "loss-inflation" is to moderate the stringency of the side constraint by defining a threshold of avoidable bad above which the infringement becomes permissible.[7] Yet there is a parallel problem of "probability-deflation." That is, one can bring the nonconsequentialist to a dilemma either by adjusting up the severity of the welfare loss, or by adjusting down the risk of actually violating the side constraint, or by changing both parameters. Thus, it is natural to think that a loss that is below the *welfare threshold* nonetheless justifies incurring a very low probability of infringing a side constraint.[8] A standard nonconsequentialist response to the problem of "probability-deflation" is what I propose to call the *Probabilitarian View*; this is the idea of picking out a *probability threshold* of "moral certainty" below which the risk of infringing a side constraint does not count as a moral reason. I think, however, that dividing up the moral space into the right and wrong by relying on a probabilistic threshold stands in the way of formulating a principled moral position. One obvious difficulty is that every chosen threshold can be challenged by asking how a little variation of probability attached to a relevant feature can weigh so much as to turn over the moral status of a choice.

The essay is organized as follows. In Section II, I will show that nonconsequentialism can be interpreted as a three-tier moral theory. The first tier contains abstract or pro tanto principles about the kinds of actions that

[6] Judith Jarvis Thomson, "Imposing Risks," in *Rights, Restitution, and Risk: Essays in Moral Theory*, ed. William Parent (Cambridge, MA: Harvard University Press, 1986), 176. For readability reasons, I omitted the word "pure" in the quotation.

[7] Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), 30 (footnote); Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), 153; Samantha Brennan, "Thresholds for Rights," *Southern Journal of Philosophy* 33, no. 2 (1995); Larry Alexander, "Deontology at the Threshold," *San Diego Law Review* 37, no. 4 (2000); Eyal Zamir and Barak Medina, *Law, Economics, and Morality* (New York: Oxford University Press, 2010).

[8] Larry Alexander, "Scalar Properties, Binary Judgments," *Journal of Applied Philosophy* 25, no. 2 (2008): 97–99.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

are associated with fundamental moral duties and rights. These are *objective moral propositions*; they do not have direct decisional force. The other two tiers are subjectivist or decisional. One includes what I call *moral requirements*: wide-scope moral norms that impose constraints on sets of beliefs and actions. The third tier contains unconditional norms that can be detached from moral requirements by a special rule of inference. I call these norms *moral prescriptions*. In Section III, I will develop decisional nonconsequentialism for a real world of competing moral obligations and imperfect knowledge by relying on the three-tier structure discussed in the previous section. I will offer a set of notational and conceptual resources that enables nonconsequentialism to gain a foothold in quantification that is similar to that claimed by consequentialists. I will go beyond Cicero's and W. D. Ross's traditional accounts of competing moral duties by developing a theory of expectational obligations that integrates degrees of epistemic confidence with the relative weights of various duties. In Section IV, I will differentiate two modes of speech in terms of which nonconsequentialists can formulate various approaches to risk, and I will concentrate myself on the Probabilitarian View. In Section V, I will argue that the risk sensitivity of any kind of moral obligation is an epiphenomenon that depends on the function of moral value associated with the factual feature(s) that is (are) relevant to that kind of obligation. I will also argue that nonconsequentialists are committed to interpreting some moral value functions as nonlinear. This is a generally unnoticed difference between nonconsequentialism and consequentialism.

II. The Objective Ought and the Subjective Ought

Robert Nozick suggests that the exploration of the ways in which risk affects moral obligations and rights "would be a fertile area for investigation, if one could avoid drowning in the morass of considerations about 'subjective-ought' and 'objective-ought'."[9] I think, however, that the most fruitful way of discussing the risk sensitivity of obligations and rights cannot avoid considerations about "subjective ought" and "objective ought." Indeed, risk must be defined by reference to probability and, according to the Bayesian subjectivist conception of probability that I use here, probabilities are subjective and simply provide a mathematical representation of our ignorance (with the exception of quantum mechanics).[10] In de Finetti's words, "only subjective probabilities exist — i.e., the *degree of belief* in the occurrence of an event attributed by a given person at

[9] Nozick, *Anarchy, State, and Utopia*, 107.

[10] John Maynard Keynes, *A Treatise on Probability* (New York: Cosimo, 2007); F. P. Ramsey, "Truth and Probability," in F. P. Ramsey, *Philosophical Papers*, ed. D. H. Mellor (Cambridge: Cambridge University Press, 1990); Bruno de Finetti, *Theory of Probability*, Vol. 1 (New York: John Wiley and Sons, 1974). Prichard also embraced Keynes's epistemic approach: H. A. Prichard, *Moral Obligation, Essays and Lectures* (Oxford: Clarendon Press, 1949), 30.

a given instant and with a given set of information."[11] However, my adoption of the subjectivist conception is an expository decision.[12]

Prichard pioneered the rejection of deontic objectivism by using an example of decision under risk:

> *Forked Road*
> We are watching a car approaching along a road which forks, and of which one fork has just suffered a landslide, and we have no idea which road the driver is intending to take, or whether he knows about the landslide. We can stop the car knowing that we may be doing something that will not help the driver, or abstain from stopping the car knowing that we may be letting the car go to destruction.[13]

We might be tempted to conclude that we have a duty to stop the car only if it is going toward destruction, that is, only if we can assert that the driver of the car is going to take the road that leads to the landslide. But we don't know which road the driver is going to take. We have two options: to stop the car or to refrain from stopping the car. Prichard claims that "we in fact think without any doubt whatever that we are bound to stop the car, unless we have reason for being quite confident that the car is about to take the safe fork."[14] Indeed, he concludes that the objective view of moral obligation is contrary to our ordinary thought because it "is in direct conflict with all the numerous cases in which we think without question that we ought to do something which we are thinking of as of the nature of *an insurance* in the interest of someone else."[15]

In *Forked Road* we cannot ascertain which of the two options is right, but we know that stopping the car maximizes our chance to make the right choice because, in the worst scenario, the driver's trip is just delayed unhelpfully. The puzzle raised by objectivism can be addressed more clearly by means of an example proposed by Parfit:

> *Three Gates*
> We (the rescuers) know that a hundred miners are all trapped in shaft *A* or in shaft *B*, but we don't know which shaft is the one in which the miners are trapped. Flood waters on the surface are rising. We know that all the miners in a flooded shaft will be drowned. We could close three flood-gates (1, 2 and 3) in order to block the flood. We have three options: if we close gate 1, we save 100 lives if the miners are in shaft

[11] Bruno de Finetti, *Theory of Probability*, 3–4.

[12] In fact, decisional nonconsequentialism could also be understood in terms of the frequentist conception of probability by reinterpreting the doxastic operator *B* that I will use later on as a nonquantifiable operator that ranges over frequency-based probabilities.

[13] Prichard, *Moral Obligation*, 29.

[14] Ibid., 29–30. Prichard attributes this example to Collingwood.

[15] Ibid., 30; italics in the original.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

> *A* and no lives if the miners are in shaft *B*; if we close gate 2, we save no lives if the miners are in shaft *A* and we save 100 lives if the miners are in shaft *B*. Finally, if we close gate 3, we are certain we will save 90 miners regardless of which shaft they are trapped in.[16]

This example is troubling for objectivism because, for all we know, both closing gate 1 and closing gate 2 are wrong. (In *Forked Road* one of the options is permissible, namely, to stop the car.) We can appreciate the difficulty by focusing on the following dilemmatic reasoning:

1. If the miners are in shaft *A*, we ought to close gate 1
2. If the miners are in shaft *B*, we ought to close gate 2
3. Either the miners are in shaft *A* or they are in shaft *B*

Either we ought to close gate 1 or we ought to close gate 2[17]

The reasoning yields a disjunction of obligations that is paradoxical because both members of the disjunction are false in a practical, decisional sense. The reasoning appears to be valid, though this appearance might be deceptive. From a logical viewpoint, the paradox derives from the combination of two forms of inference. First, the reasoning appeals to a deontic variant of constructive dilemma (CD):

1. $p \supset \mathrm{O}x$
2. $q \supset \mathrm{O}y$

_____ by CD

$p \vee q \supset \mathrm{O}x \vee \mathrm{O}y$

Second, the inference draws on what I call the *factual detachment* model of moral inference. According to this model, objective moral rules guide moral choice by allowing an agent to detach a practical conclusion via modus ponens (MP). As applied to the dilemmatic conclusion, the model would validate the following schema:

1. $p \vee q \supset \mathrm{O}x \vee \mathrm{O}y$
2. $p \vee q$

_____ By MP

$\mathrm{O}x \vee \mathrm{O}y$

On this possible logical reconstruction, both the deontic variant of constructive dilemma and the deontic variant of modus ponens permit drawing

[16] Derek Parfit, On What Matters, Volume I (New York: Oxford University Press, 2011), 159.
[17] I took the structure of the reasoning from: Niko Kolodny and John Macfarlane, "Ifs and Oughts," *The Journal of Philosophy* 107, no. 3 (2010).

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

a disjunction of obligations from (i) a logical consequence of the objective propositions that establish our conditional obligation to close gate 1 and our conditional obligation to close gate 2, together with (ii) the disjunctive factual premise that says that either the miners are in shaft A or they are in shaft B.

If the deontic variant of constructive dilemma and the deontic variant of modus ponens are valid forms of inference, it seems that we should conclude that either we ought to close gate 1 or we ought to close gate 2. In fact, it might be claimed that one of such options is the "objectively" right action because, by performing it, we will save the one hundred miners. However, we are inclined to think that we ought to choose neither of these options. Since one option is right and the other is seriously wrong, and we do not know which one is the right option, it is too risky to choose one of them, randomly or in any other way. Clearly, given that we don't know which shaft the miners are trapped in, by closing gate 1 or 2 we would impose a 0.5 risk of death on all the miners. In contrast, if we choose to close gate 3, we are certain that we will save 90 miners, with a sure loss of only 10 miners. It seems that we ought to close gate 3, but this intuition is inconsistent with the dilemmatic conclusion of the reasoning.

The objectivist approach yields a counterintuitive result in *Three Gates*. The example shows that we need a different method for inferring practical conclusions from objective moral propositions. The paradox derives from the fact that the deontic premises of the dilemmatic reasoning — propositions (1) and (2)— establish abstract or pro tanto obligations, whereas the conclusion signals use of a notion of decisional or concrete ought. Insofar as we keep in mind that the reasoning can only draw a conclusion that establishes an abstract or pro tanto duty (since it is premised on abstract moral propositions), we might be able to dissolve the paradox by connecting objective and practical moral propositions in a different way. What is troubling in the application of deontic constructive dilemma and deontic modus ponens in *Three Gates* is that, in a concrete, decisional sense, it is not true that we ought to close gate 1(option x) or that we ought to close gate 2 (option y). It is evident that we ought to close gate 3 (option z).

Various authors have noted that moral theory must develop a bridge that leads from objective normative principles to practical conclusions that are binding in concrete situations, especially when our factual beliefs are probabilistic. For instance, Frank Jackson says with respect to consequentialism: "We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics, and having the best consequences is a story from the outside."[18] Jackson goes on to draw a distinction between ethics and physics: "It is fine for a theory in physics

[18] Frank Jackson, "Decision–Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101 (1991), 467.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

to tell us about its central notions in a way which leaves it obscure how to move from those notions to action, for that passage can be left to something which is not physics; but the passage to action is the very business of ethics."[19] Michael Smith asks: "How do we get from abstract knowledge of what it is about an action that makes it right — that it maximizes value — to a decision about what to do in a concrete situation?"[20] And Holly Smith argues that "principles of objective status" stating, for instance, that an action is wrong if it involves killing an innocent must be supplemented by "principles of subjective right": "We need principles of subjective right precisely because agents must often make decisions despite their lack of information or time or ability to deliberate further. Principles of subjective right are needed precisely to assist agents in deciding what to do in these circumstances."[21]

Indeed, "ought" can be used not only in an *objective* sense but also in a *decisional* or *subjective* sense. The objective "ought" figures in such fundamental moral principles as those establishing the duty not to harm and the duty to aid. This "ought" is used to evaluate behaviors and states of affairs by reference to perfect deontic possible worlds. For instance, a moral proposition that forbids punishing an innocent considers any world in which an innocent is punished as excluded from the set of perfect deontic possible worlds. Therefore, a second-best objective "ought" (e.g., a contrary-to-duty imperative) is meaningless, for it would have to be determined by reference to a nonperfect deontic world. Unlike the abstract, objective "ought," the subjective or concrete "ought" picks out an option as the deontically optimal option within the *feasible set of options* for an agent in a given choice context. Every feasible set of options is constrained by various factors. Among those factors are the physical and social possibilities of various actions and their outcomes, personal and technical capabilities, the strategic interactions with other agents, the kind and degree of available evidence, the extent of the agent's beliefs, and so on. Since the subjective, decisional "ought" is a *constrained ought*, it has different logical relations to other propositions; for instance, unlike the objective "ought," the subjective "ought" entails "can."

I will use the letter "O" to stand for the objective, abstract "ought," and will adjust its font to "○" in order to express the decisional, subjective "ought."

[19] Ibid. Jackson's supposition that a physical theory can leave the "passage to action" to a different discipline is very controversial. Such actions as experimentation and technical application must be determined by physics as well.

[20] Michael Smith, "The Right, the Good and Uncertainty," in Terry Horgan and Mark Timmons, eds., *Metaethics after Moore* (Oxford: Clarendon Press, 2006), 145. Smith disagrees with Jackson. He thinks that "the story from the inside of an agent" is not the business of ethics, but rather "the story of rational decision-making as outlined in the diagram above, the story that is formalized — or anyway partially formalized — in decision theory." This seems like a semantic dispute. Decision theory in ethical contexts is the business of ethics.

[21] Holly M. Smith, "Subjective Rightness," *Social Philosophy and Policy* 27, no. 2 (2010): 74.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

For simplicity's sake, I will treat both "oughts" as propositional operators.[22] I will also assume that any nonconsequentialist theory establishes a set of fundamental principles that determines, for every option ϕ available to an agent, that if ϕ instantiates the kind of actions *K*, the agent has an abstract, pro tanto moral obligation to do ϕ. In other words, I will assume that any fundamental nonconsequentialist principle applies to each available option ϕ according to the following logical schema:

$K\phi \supset O\phi$

This proposition has the form of a narrow-scope deontic conditional, because the deontic operator ranges over the consequent. I think that fundamental moral propositions lead to moral requirements in the following sense: if a nonconsequentialist theory endorses an objective moral proposition of the form $K\phi \supset O\phi$, it also endorses a wide-scope decisional principle that requires, for each option ϕ available to you, that [if you believe $K\phi$, you do ϕ]. The idea is that, if a nonconsequentialist moral theory establishes an objective moral proposition of the form $K\phi \supset O\phi$ about any option ϕ in a menu of mutually exclusive options, it also includes a moral requirement concerning that option that is derivable from the objective proposition according to a normative principle that I call *subjectivist transform* (*ST*). Let "B" represent "you believe" and let "D" represent "you do." According to ST, a fundamental moral proposition of the above form is transformed into a moral requirement in the following way:

1. $K\phi \supset O\phi$

_____ by ST

2. $\circ (BK\phi \supset D\phi)$

Within its restricted scope, ST functions as a normative bridge principle that inputs an objective moral proposition and outputs a moral requirement that imposes on your practical reasoning the following constraint: if you believe about any of your options ϕ (within a set of mutually exclusive options) that ϕ instantiates the class of actions *K*, then you do ϕ. While the premise expresses an objective ought, the conclusion is a *moral requirement*, that is, a wide-scope subjective ought conditional that ties your believing that any of your options ϕ instantiates a fundamental moral kind of actions *K* to your doing ϕ.[23]

[22] In this respect, I follow Williams's approach. See Bernard Williams, "Ought and Moral Obligation," in *Moral Luck* (Cambridge: Cambridge University Press, 1981), 114–23. Although there are more complex ways to analyze the logical form of "ought," I think that this simplification does not imply any significant explanatory cost related to the risk sensitivity of moral obligation.

[23] Prichard argued that the subjective "ought" does not dictate an obligation to *do* something, but "to perform an activity of a totally different kind, that of *setting or exerting ourselves*

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

A moral requirement establishes obligations on sets of beliefs and actions and relative to a feasible set of options. The conclusion obligates moral agents either to abstain from forming or holding a certain belief or to do a certain action.[24] In fact, (2) is equivalent to:

(3) ○ (not-BKϕ or Dϕ)

Whereas objective "ought" principles typically have the form of a narrow-scope deontic conditional, moral requirements are wide-scope "ought" statements, because their function is to guide practical reasoning by providing constraints on combinations of beliefs and actions.[25] Why could conclusion (2) not be expressed in a narrow-scope form? One fundamental difference between a wide-scope "ought" conditional and a narrow-scope conditional is that the latter authorizes factual detachment. And in this context factual detachment would give any agent the power to create and extinguish unconditional obligations.

Consider the following *reductio* to illustrate why moral requirements cannot have a narrow-scope form:

1. If someone goes to the premiere, he ought to purchase a full-price ticket unless he is on the guest list.
2. If someone goes to the premiere and believes that he is not on the guest list, he ought to purchase a full-price ticket.
3. I go to the premiere and I believe that I am not on the guest list

---

I ought to purchase a full-price ticket

Suppose we could obtain (2) from the objective proposition (1) (for example, by strengthening the antecedent). Then the narrow-scope form of (2) would allow factual detachment, that is, we could draw from conditional (2) and premise (3) the conclusion that we have an unconditional obligation to

---

to do something." He claimed that what we directly do is our intending (not intending) to do something, everything else being a more or less likely result of what we intend or do not intend. This is a plausible claim. However, I do not want to consider here the probability that the agent's intention to do *x* will actually bring about his doing *x*. This would demand expository complexities that exceed the scope of this essay. Cf. Prichard, *Moral Obligation*, 34–35; italics added.

[24] I discuss the problem of "epistemic obligations" later on.

[25] Thomas Hill, "The Hypothetical Imperative," *Philosophical Review* 82, no. 4 (1973): 429–50; Greenspan, 1975 Patricia Greenspan, "Conditional Oughts and Hypothetical Imperatives," *The Journal of Philosophy* 72, no. 10 (1975): 259–76; Carlos Alchourron, "Detachment and Defeasibility in Deontic Logic," *Studia Logica* 57 (1996): 5–18; John Broome, "Normative Requirements," *Ratio* 12, no. 4 (1999): 398–419; Jonathan Dancy, *Practical Reality* (Oxford: Oxford University Press, 2000), 49–76; R. Jay Wallace, "Normativity, Commitment, and Instrumental Reason," *Philosophers' Imprint* 1, no. 3 (2001). For an application to judicial reasoning, see: Hugo Zuleta, *Normas y justificacion* (Madrid, Barcelona, Buenos Aires: Marcial Pons, 2008).

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

purchase a full-price ticket, which we do not necessarily have. Suppose that one instant after I draw that conclusion, someone comes along and informs me that the main actress has added my name to the invitee list. It seems I still have the unconditional obligation to purchase the ticket, precisely because it is unconditional. (Am I also under that obligation if I change my mind in the last minute and go somewhere else?)

It is obvious that something has gone wrong. The problem is that we have used a narrow-scope structure in order to formulate the requirement that follows from (1), and moral requirements only say that inasmuch as your epistemic state is such and such, you ought to do so and so. The logical structure of a moral requirement does not allow factual detachment, and, hence, does not permit the conclusion that I have an unconditional obligation to purchase a full-price ticket. Therefore, I think it is clear that moral requirements must have wide-scope conditional form.

Once we have a normative principle that allows an agent to draw a moral requirement from a fundamental moral proposition, we need to explain how a moral requirement can prescribe a certain option (as something different from a combination of a belief and an action). At this point, the advantage of the wide-scope conditional form confronts us with a problem. The problem is precisely that, unlike narrow-scope conditionals, wide-scope deontic conditionals do not allow detachment of an unconditional deontic statement by modus ponens. Fortunately, this does not mean that no detachment is possible, and one form of detachment must be possible for the moral requirement eventually to prescribe a certain course of action. Among the various detachment rules proposed by deontic logicians, there is one that can be tailored to our purposes. That is *deontic detachment*, which has the following form:

$$\begin{array}{l} \mathrm{O}\,(p \supset q) \\ \mathrm{O}p \\ \hline \mathrm{O}q \end{array}$$

Deontic detachment licenses inference of an unconditional obligation from a wide-scope deontic conditional provided that the statement is coupled with another proposition that obligates you to see to it that the antecedent of the conditional is true.[26] Deontic detachment is preferable over factual detachment (i.e., modus ponens) as a model of moral inference. For example, deontic detachment is capable of avoiding Chisholm's paradox because contrary-to-duty imperatives cannot be inferred from wide-scope deontic statements if the antecedents are "not obligatory"

[26] Patricia S. Greenspan, "Conditional Oughts and Hypothetical Imperatives," *The Journal of Philosophy* 72, no. 10 (1975): 265; Mark I. Vorobej, "Conditional Obligation and Detachment," *Canadian Journal of Philosophy* 16, no. 1 (1986): 16.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

(e.g., if they are "forbidden").[27] By applying a variation on deontic detachment, I want to consider a different model of moral inference that I call *epistemic deontic detachment* (EDD). It contains two premises: first, a wide-scope ought statement, and, second, an epistemic ought statement. With EDD we can infer a *moral prescription* from a wide-scope moral requirement as long as we assent to an epistemic norm that requires upholding the antecedent belief:

1. $\bigcirc (BK\phi \supset D\phi)$
2. $\bigcirc (BK\phi)$

_____by EDD

$\bigcirc D\phi$

There are two potential objections to my proposal for connecting objective moral propositions and subjective moral propositions. First, it might be argued that deontic constructive dilemma and deontic modus ponens still produce paradoxical results when combined with ST and EDD in *Three Gates*. For instance, it might be suggested that a consequence of fundamental moral principles is that if an action ϕ (for example, closing gate 1) instantiates the relevant kind of action (for example, saving all the trapped miners), it is obligatory to do ϕ, and that the same applies to action ψ (for example, closing gate 2). Then, by deontic constructive dilemma, a deontic dilemmatic conclusion could be drawn:

1. $K\phi \supset O\phi$
2. $K\psi \supset O\psi$

_____ by CD

$K\phi \vee K\psi \supset O\phi \vee O\psi$

The objection could go on by suggesting that, if we applied ST to this dilemmatic conclusion, we could obtain the following puzzling moral requirement:

$\bigcirc \{B (K\phi \vee K\psi) \supset D\phi \vee D\psi\}$

This supposed moral requirement would still be problematic because it would permit drawing a wrong moral prescription ($\bigcirc D\phi \vee D\psi$), given that we ought to believe, as a result of the assumptions in *Three Gates*, that either closing gate 1 instantiates the kind of action of saving all the trapped miners, or closing gate 2 instantiates that kind of action. So where does the error lie?

[27] Roderick Chisholm, "Contrary to Duty Imperatives and Deontic Logic," *Analysis* 24 (1963): 33–36.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

Actually there are two errors. The first error is that this is an illegitimate application of ST. In fact, ST does not allow us to convert a disjunction of two deontic objective operators into a single subjective practical operator. In the consequent of the dilemmatic conclusion there is a disjunction of obligations, and ST cannot be applied to this disjunction. A disjunction of two objective obligations cannot be turned into a unique subjective obligation with a disjunctive content, and ST and EDD certainly do not permit this kind of conversion.

The second error lies in thinking that the purported moral requirement presumably obtained by ST is a moral requirement. In fact, it is *not* a moral requirement. By definition, moral requirements are obtained by ST, and ST only applies to fundamental objective propositions of the form: $K\phi \supset O\phi$. The material implication is surely an inadequate formal resource to express those fundamental connections, but it can be used for expository purposes as long as we exercise caution. Thus, we must bear in mind that the dilemmatic conclusion obtained by deontic constructive dilemma is not a fundamental objective proposition that connects one option instantiating a morally relevant kind of action with one pro tanto moral obligation to perform *that* option. Moral requirements can be obtained by ST when you assert that a particular option of an agent instantiates a fundamentally relevant kind of action. ST cannot be applied to draw moral requirements when you assert propositions that refer to more than one option, such as the proposition that either an option instantiates a relevant kind of action, or *another* option instantiates that kind of action. In nonconsequentialist moral theories, objective propositions of the form "$K\phi \supset O\phi$" establish an injective or one-to-one function between kinds of action and moral duties. The dilemmatic conclusion obtained by deontic constructive dilemma does not preserve the needed distinctness.

The second potential objection emphasizes that the key to detach unconditional subjective obligations from moral requirements is to assert that we have "epistemic duties," that is, subjective obligations to form or uphold certain beliefs. But in talking about "epistemic duties" it seems that I am committed to a controversial claim, because "ought" is generally thought to entail "can" and the beliefs we form, hold, and relinquish are often not within our *direct* voluntary control. So if we suppose that an agent ought to form or relinquish a certain belief, but he cannot do so, it seems that our supposition is false and he does not have that duty. Though this is not a work in epistemic deontology, I want to make three comments that render epistemic oughts plausible. First, if we say that someone ought to believe p, we just mean that she believes p in the best world out of a set of possible epistemic worlds relative to certain constraints. We do not imply that the passage to that epistemic world is under her direct control. Second, epistemic duties require both *evidentiary* and *inferential* actions. When we say that someone ought to believe p, we mean that he ought to gather and examine a relevant body of evidence (from which p follows) and that he ought to infer p as

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

a conclusion from that available evidence. Both evidentiary and inferential actions seem to be under the agent's control. Third, the evidentiary and inferential actions required by epistemic duties need not be identical to the formation of beliefs as states of mind. In a practically consequential sense, a belief is an input in a decision-making process. When the Holy Office ordered Galileo to relinquish his belief that the sun is the center of the world and that the earth moves, I take it that the inquisitors were not assuming that Galileo had direct control over his beliefs as states of mind. The famous injunction should rather be understood as ordering Galileo to reexamine a body of theological-cum-physical evidence and to conclude from that evidence that the geocentric dogma is true. What the church was interested in was not Galileo's states of mind, but rather his use of Copernican heliocentrism as an auxiliary reason for taking heretic decisions such as holding, defending and teaching that the earth moves. In an operational sense, the injunction ordered Galileo to act in all practical ways as if he subjectively believed in the geocentric dogma (which presupposes that the church was not seriously worried about Galileo's potential insincerity).

For a simple world in which each available option open to an agent in a given context of choice instantiates one single moral obligation, and the agent has access to perfect factual information about whether each of his options instantiates a moral obligation, the upshot of my discussion about the decisional "ought" can be summarized as follows. *Objective moral propositions related to each available option (within a menu of options) lead to unconditional moral prescriptions via moral requirements when the agent ought to believe that the factual condition for the obligatoriness (or permissibility) of an option obtains.* More often than not, however, moral agents are confronted with options that instantiate various competing moral obligations and/or with options about which they can only form or uphold beliefs with a level of confidence that is lower than certainty. How can a moral requirement and, ultimately, a moral prescription be obtained in order to provide guidance for moral decision making in a real world of deontic conflicts and risky choices? It is obvious that ST and EDD do not work in the real world. However, it is possible to revise those principles in order to infer decisional oughts from competing abstract oughts and probabilistic beliefs. Before embarking on this task, we must delve into the ways in which nonconsequentialism can address risk-imposing choices.

III. Decisional Nonconsequentialism

Consequentialism has an obvious solution to the problem of epistemic constraints in moral deliberation. In Frank Jackson's words, consequentialists can just "take a leaf out of a decision theory's book."[28] Instead of

[28] Jackson, "Decision–Theoretic Consequentialism and the Nearest and Dearest Objection," 463.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

overall actual utility, the consequentialist turns to *expected utility*, which is defined as the result of multiplying the utility of each possible outcome by the agent's subjective probability of that outcome. Then decisional consequentialism requires the agent to choose and perform the action that maximizes the overall expected utility. On this view, risk attitudes are exogenous to utility. Jackson and other ethical consequentialists say that their notion of utility corresponds to modern decision theory's concept of utility.[29] But they get it wrong. Indeed, the notion of utility used by consequentialism is different from the concept of utility defined in modern expected utility theory. The latter notion was introduced by Daniel Bernoulli and given a modern formulation by Von Neumann and Morgenstern.[30] Von Neumann and Morgenstern define the utility function in terms of preferences over lotteries (i.e., choices between ordered pairs of monetary outcomes and probability distributions). The modern concept of utility encodes a preference ranking among risky prospects and, therefore, risk attitudes are built into the utility function. Consequentialists still use a premodern notion of utility according to which moral choices should be grounded on utility, rather than a notion according to which utility is just a representational device of moral choices.

Jackson argues that a theoretic-decision version of consequentialism must require the maximization of *expected* utility, rather than the maximization of *actual* utility. He gives the following example, which is similar (but not identical) to *Three Gates*:

> *Three Drugs*
> Jill can choose among three different drugs, *a*, *b*, and *c*, to treat her patient, John, who has a minor skin condition. Careful consideration of the literature indicates that she can obtain John's partial cure with drug *a*. One of drugs *b* and *c* will provide a complete cure of John; however, the other will kill him. There is no evidence Jill can rely on to tell which of the two is the complete cure and which the lethal drug. (John's informed consent cannot be obtained.)[31]

Jill knows that she can maximize utility by giving John drug *b* or by giving him drug *c*. Although Jill knows for sure that *b* or *c* will yield the greatest

[29] Jackson, "Consequentialism and the Nearest and Dearest Objection," *Ethics* 101, no. 3 (1991); Michael J. Zimmerman, *Living under Uncertainty. The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008); Ingmar Persson, *From Morality to the End of Reason* (Oxford: Oxford University Press, 2013), chap. 10.

[30] John Von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1953), chap. I; Kenneth J. Arrow, "Exposition of the Theory of Choice Under Uncertainty," in *Essays in the Theory of Risk-Bearing* (Amsterdam: North-Holland, 1971).

[31] Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 462–63.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

health benefit for John, we all would agree that she is bound to give him drug *a*, which has no chance to be the objectively right action. A consequentialist following expected utility would reach the correct conclusion in *Three Drugs*. If he assigns a value of 1,000 to John's total cure, and, prescribing drugs *b* or *c* has a 0.5 chance of causing this benefit, the expected utility is 500. Given that this option also has half a chance of causing John's death, the expected utility of the option must be discounted by the negative expected utility of John's death. Suppose we attach a disvalue of -100,000 to this outcome. Then the option of prescribing *b* or *c* has an expected utility of -50,000, which means that the net expected utility of the option is -49,500. Of course, the consequentialist can conclude that this (negative) utility must be lesser than any (positive) utility to be derived from the sure but partial cure caused by prescribing drug *a*.

Although Jackson's approach is obviously tailored to fit consequentialism, Michael Zimmerman argues that all substantive theories of moral obligation, including nonconsequentialism, should maintain what he calls the *Prospective View*, which is based on the notion of *expected value*.[32] Zimmerman follows Jackson in (wrongly) suggesting use of the modern economic notion of expected utility: "One obvious proposal (which for the time being I will assume to be correct) is to construe it in terms of the maximization of expected value."[33] According to the Prospective View, "an agent ought to perform an act if and only if it is the prospectively best option that he has."[34] Zimmerman says that an objectivist approach to moral obligation gives a wrong answer in *Three Drugs* because it requires Jill to give John either drug *b* or *c*. On this view, Jill's giving John drug *a* is not a way of minimizing the risk of doing wrong, but a sure way of doing wrong.[35] For the Prospective View, in contrast, giving John drug *a* is prospectively best, "in that it provides Jill with a better prospect of achieving what is of value in the situation (namely, the restoration of John's health)" than giving him either of the other two drugs.[36]

Zimmerman argues that the Prospective View can be applied to a theory of moral rights.[37] However, the extension to moral rights results in a convoluted position, in which "maximal moral suitability to some aspect of the agent's situation" substitutes "best." He takes the concept of "moral suitability" from David Ross: "To say that an act suits its agent's situation is to say that it serves some actual value that is at stake in that situation."[38] However, Zimmerman does not say much by way of a procedure to calculate "prospective maximal suitability." His Risk Thesis (in its definitive

[32] Michael J. Zimmmerman, *Living under Uncertainty*, 17–42.
[33] Ibid., 19.
[34] Ibid., 19.
[35] Ibid., 20.
[36] Ibid., 18.
[37] Ibid., 33–34.
[38] Ibid., 74.

formulation) makes rights dependent on the information of the agent at the moment of action:

> We have moral rights against others that they not impose on us any risk by virtue of their failing to do what is prospectively most suitable to some aspect of their situation, when such failure constitutes the failure to satisfy an associative obligation that they owe to us.[39]

For Zimmerman, all moral rights are correlative with oughts relative to available evidence. Therefore, whether an agent violates a right of another depends on the expected value of the action, which, in turn, depends on the probability that the right will actually be violated. If an agent's behavior has a high prospect to respect another's right, then he has not violated the right. On the contrary, if the behavior has a high prospect to violate another's right, he has violated the right, even though the act did not actually constitute a violation of the right. On Zimmerman's view, *expected* rights violations are the only legitimate rights-based grounds for complaint. Therefore, a guilty person has a ground to complain about the expected rights violation committed by a court that applied an unreliable procedure for ascertaining his guilt, a procedure that would too often result in the conviction of innocent defendants.

The difficulty with this account of rights is that it yields consequences that sometimes are inconsistent with legal and moral rights. For instance, Zimmerman says that if I promise to return a book to you by Friday, "my obligation is to do that which will most probably result in your receiving the book by Friday, whether or not you in fact receive it."[40] This example is counterintuitive, since in contract law both the obligation to return the book and its correlative right are extinguished only if the book is actually returned. Zimmerman must introduce the mysterious idea of your right being "rekindled" when I discover that, despite my good efforts, you have not received the book.[41] His other example is more palatable. Zimmerman says that, if Brenda shoots Alf in self-defense, her justification is that "Alf infringed a right of Brenda's having to do with her being (at risk of being) harmed."[42] Criminal law can accommodate more easily the concept of *subjective rights*.[43] It sounds plausible to argue that Alf infringed a right of Brenda's whether or not he actually caused the harm if, in light of the evidence available to Brenda, he imposed a grave risk on her, and that precisely this is the reason why her defensive shooting was justified.

[39] Ibid., 87.
[40] Ibid., 88.
[41] Ibid., 94. Other rights in contract law might be interpreted in subjectivist fashion (e.g., contractors' right to a good faith performance of the contract by the other party.)
[42] Ibid., 99.
[43] This notion has nothing to do with the Continental expression "subjective rights"; in English all rights are subjective in the latter sense.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

Nozick's dualist approach seems an attractive alternative to the objective and prospective accounts of rights. He initially appeals to natural procedural rights against the use of unreliable and unfair procedures of guilt ascertainment in order to justify the prohibition of independents' self-help procedures of enforcement of their rights. However, he acknowledges that the "the notions of procedural rights, public demonstration of guilt, and the like, have a very unclear status within state-of-nature theory."[44] The problem is that natural procedural rights are not among classic natural rights, and, hence, procedural rights must be separate sources of moral obligation. For this reason he also tries to ground his case on individuals' "jurisdictional" right against risk-imposition.[45] This right entails that punishing a guilty person on the basis of unreliable or unfair evidence is morally wrong. But what right could ground his possible complaint about the punishment? Since risk is a subjective concept (under the Bayesian approach), Nozick must discuss how to recover a right against risk-imposition from objective moral rights.

Nozick then proposes an "epistemic principle of border-crossing": "If doing act *A* would violate *Q*'s rights unless condition *C* obtained, then someone who does not know that *C* obtains may not do *A*." A consequence of this principle is that no enforcer or self-enforcer of natural rights has a power to punish *Q* unless he knows that *Q* is guilty. In order to ban the use of unreliable enforcement procedures, Nozick presents two different versions of this idea:

(1) Enforcers do not have a right to impose punishment on an agent unless they *know* that the agent is guilty.
(2) Enforcers have a right to impose punishment on any offender but their exercising this right is wrong (impermissible) unless they *know* that the agent is guilty.[46]

According to principle (1), we violate an individual's right (whether he is guilty or innocent) if we punish him without knowing that he is guilty. According to principle (2), if we punish an innocent, we do not violate his rights, but what we do is wrong.[47] Nozick challenges the idea of subjective rights through puzzling questions:

> Does a person who did violate another's right himself have a right that this fact be determined by a fair and reliable procedure? It is true that an unreliable procedure will too often find an innocent person guilty. But does applying such an unreliable procedure to a guilty

[44] Ibid., 96.
[45] Nozick, *Anarchy, State, and Utopia*, 105.
[46] Ibid., 106.
[47] Nozick offers different interpretations of the epistemic condition an agent must have in order to be permitted to cross the jurisdictional borders of another individual (ibid., 106–7).

> person violate any right of his? May he, in self-defense, resist the imposition of such a procedure upon himself? But what would he be defending himself against? Too high a probability of a punishment he deserves?[48]

Since Nozick does not like the idea of subjective rights, he picks out principle (2) and concludes that an unreliable punisher violates no right of the guilty person, but he does something wrong if he punishes him.[49] He claims, accordingly, that "*many procedural rights stem not from rights of the person acted upon, but rather from moral considerations about the person or persons doing the acting*."[50] This approach is inconsistent with the conceptual correlativity between rights and duties, but Nozick shows no concern for this well-known thesis.

My own account is different from Zimmerman's and Nozick's. I claim that we can both (1) retain the correlativity thesis, and (2) hold that objective rights do not vary with subjective probabilities or epistemic considerations. The trick is to argue that rights are correlative with the objective ought and independent of epistemic or deliberative concerns. However, moral principles that establish rights must be subjectively transformed into moral requirements by ST, and then into moral prescriptions by EDD. Moral prescriptions that are ultimately inferred from moral requirements and epistemic norms are, of course, sensitive to epistemic considerations. Moral prescriptions can establish both *subjective obligations* and *subjective rights* and, therefore, there is a sense in which a guilty suspect can complain about a court having breached his subjective right to have his responsibility determined through a reliable procedure, even though his objective rights have not been violated at all.

Objective rights are correlative with the objective ought and, therefore, their content does not vary with the subjective probabilities of the relevant agents' beliefs. In this sense, a guilty defendant does not have a right not to be punished, and, consequently, if a court punishes him through an unreliable fact-finding procedure, or just by a utilitarian calculus, he has no moral grounds to complain. But this does not preclude the possibility, at the level of moral prescriptions, of conceiving of a right that is correlative with the subjective ought. Nozick's dismissal of subjective rights is correct only in the sense that objective rights are much more common in moral and legal language than are subjective rights. Because of their historic evolution in legal and political parlance, rights are naturally framed in an objective sense. However, in procedural law some rights can be derived as the correlatives of subjective obligations established in moral prescriptions. Nozick says that "many procedural rights" stem from moral

[48] Ibid., 103.
[49] Ibid., 107.
[50] Ibid., 107; italics added.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

considerations about the agent (that is, obligations), but it could equally be argued that they stem from underlying subjective rights. In criminal law a good deal of responsibility (for example, attempts) is attributed to violations of subjective obligations, but it could equally be claimed that responsibility is justified by the subjective rights of (potential) victims. On the other hand, Zimmerman's suggestion to relativize all moral rights to the evidence of the agent leads to a complete distortion of the concept of moral and legal rights. Although subjective rights are not unintelligible, they are often a bit of an oddity in many areas of law (for instance, contract law). The truth of the matter is, I think, that normative language has enough flexibility to accommodate revisionary uses. For example, it is intelligible to say that citizens have a general subjective right against public officials that such officials seriously consider citizens' objective constitutional rights when making decisions that may affect such rights.

While "ought" can have an objective or subjective meaning, this does not entail that an agent must consider both oughts when making a decision. Although the objective ought is important in *grounding* a prescriptive, subjective ought, the concept of ought that must provide direct guidance is the subjective one. The objective ought is practically inert insofar as it does not ground a decisional ought. This means that the recognition that "ought" is ambiguous between the two distinguished senses does not raise the problem Zimmerman adduces against deontic objectivism. He says that an agent must have an *unequivocal* answer to the question of what ought must guide her choice as a conscientious person.[51] There is no gainsaying that one single ought must eventually settle the practical question. The subjectivist ought is the last normative element dictating a conscientious agent's decision. Any apparent conflict between both senses of ought must be dissolved, as far as practical guidance is concerned, in favor of the ought dictated by the moral prescription.

It is not easy to visualize how subjective transform (ST) and epistemic deontic detachment (EDD) can possibly operate with probabilistic beliefs and/or conflicting moral obligations. I will also follow Ross, but not his *Foundations of Ethics*. Rather, I will draw on *The Right and The Good*, where Ross, without mentioning Cicero, develops the Ciceronian view that conflicts of duties must be resolved by "calculating" their respective weights.[52] Indeed, Ross is clearly concerned with the decisional domain of ethics in introducing his conception of moral deontologism as a theory characterized by an irreducible plurality of pro tanto obligations of different kinds:

> When I am in a situation, as perhaps I always am, in which more than one of these *prima facie* duties is incumbent on me, what I have to do is to study the situation as fully as I can until I form the considered

[51] Ibid., 7.
[52] Cicero, *De Officiis*, Bk I, paragraph 18.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

> opinion (it is never more) that in the circumstances one of them is more incumbent than any other; then I am bound to think that to do this *prima facie* duty is my duty *sans phrase* in the situation.[53]

However, in this passage Ross refers to clashes of duties, not to the problem of probabilistic factual propositions relevant to moral choice. (He addresses this problem explicitly in the *Foundations of Ethics*, where he embraces subjectivism under the influence of Prichard.) To discuss how probability distributions can affect moral values in nonconsequentialism, it is necessary to develop Ross's view that the "duty proper" or "actual duty" is "the duty which is in the circumstances more of a duty."[54] Specifically, we must reconceptualize "more of a duty" so that it contemplates both the degree of moral importance *and the degree of probability*. We must, so to speak, complete the lacking half of Ross's cursory but insightful account. Naturally, constructing a decisional theory requires the introduction of a new a set of logical, notational and conceptual tools.

To start with, we need to map moral values (or disvalues) onto the set of real numbers in order to obtain numerical measures for pro tanto obligations. Therefore, instead of the qualitative conception of value underlying deontic logics, we must interpret pro tanto obligations in quantitative terms. How can quantification be integrated with the conceptual structure of pro tanto obligations? My proposal is to say that each pro tanto moral obligation is associated with a function that relates a relevant quantifiable feature (for example, external utility) to moral value. Since pro tanto obligations can be either positive or negative (that is, prohibitions), every obligation *applicable* in a given choice context is indexed to a real number ($\pm v$) that indicates moral value or disvalue. Obligations fall on doing ϕ ($D\phi$) or not doing ϕ (not-$D\phi$). Conventionally, an obligation has positive moral value if and only if it is an obligation with positive content ($D\phi$), and negative value (disvalue) if and only if it is an obligation with negative content (not-$D\phi$). In addition, for simplicity's sake I will assume that a liberty or permission (to do ϕ and/or not to do ϕ) has moral value 0. (In many cases, a permission curtails the scope of application of an obligation or changes the moral value of an option, but I will ignore these complexities here.) One difficulty is that the expression "applicable obligation" allows two *opposite* options to have the same numerical value, thus leaving indeterminate which option the agent ought to choose. The reason is that "applicable obligation" does not disambiguate between satisficing and breaching an obligation (negative or positive). For example, if we have two options, to aid someone and not to aid her, it seems that both options get the same positive number because the positive obligation to aid is applicable to both (one option satisfies the obligation, the other breaches it).

[53] W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), 19.
[54] Ibid., 18.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

By the same token, the option not to kill and the option to kill can get the same negative number because the same negative obligation applies to both. To avoid this kind of indeterminacy in the decision-making procedure, I will say that an option *switches on* a positive obligation (or permission) when it has a non-zero probability to satisfy that positive obligation (or permission), and that an option *switches on* a negative obligation when it has a non-zero probability to breach that negative obligation. The phrase "to switch on an obligation" has the goal of avoiding an ambiguity in ordinary language that stands in the way of assigning consistent numerical values to opposite options. It is a simple, indirect way of stipulating that the concept of "satisfaction" is limited to positive obligations (or permissions), and that the concept of "breach" is circumscribed to negative obligations. For example, in *Three Gates* the options of closing gate 1 and of closing gate 2 switch on the positive obligation to save one hundred lives. In *Three Drugs* the options of giving John drug *b* and giving him drug *c* switch on the negative obligation not to harm (that is, not to kill).

Consider any option ϕ within a set of mutually exclusive alternatives open to an agent. With respect to each option ϕ so conceived, I can define the key concept of *expectational obligation* (EO).[55] The expectational obligation relative to ϕ results from the overall moral value of the obligations switched on by ϕ (according to the moral value function of each obligation) weighted by the probability that ϕ will actually satisfy or breach each relevant kind of obligation. So one can obtain an expectational obligation for a given option by adding the moral value of every obligation or right (or permission) switched on by that option weighted by the probability of that option actually fulfilling or violating each switched-on obligation.

In order to design a detachment rule for expectational obligations, we must use two different subscripts: the subscript in operator o is a variable for the moral value (positive number) or disvalue (negative number) of the corresponding obligation O, and the subscript in the operator *B* indicates the credence or level of confidence attached to the belief (a probability number). (When the credence is 1, the latter subscript is left out.) Because credences are real numbers within the interval [0,1], we can revise EDD in such a way that it permits drawing an unconditional expectational obligation whose value is calculated on the basis of the (positive or negative) value v of every (positive or negative) obligation O i that is switched on by a given option and the credence c relevant to that O i that is obligated by appropriate epistemic norms. Accordingly, every expectational obligation must be indexed to an *expectational deontic value*, $\pm d$. The expectational deontic value of each expectational obligation, $\pm d$, is obtained by adding

[55] Although his enterprise is different from mine, it is relevant to point out that Ralph Wedgwood proposes an expectational semantic analysis for deontic propositions. See Ralph Wedgwood, "Objective and Subjective 'Ought'," in Nate Charlow and Matthew Chrisman eds., *Deontic Modality* (Oxford: Oxford University Press, 2015).

up the products of the values $\pm v$ and epistemically obligated credences c corresponding to every duty O_i that is switched on by each option.

The normative principle that we can then propose to substitute for EDD, which I will call *expectational epistemic deontic detachment* (EEDD), allows detachment of expectational obligations to do-ϕ (not to do-ϕ) for each mutually exclusive option open to an agent in a given choice context. The EEDD rule can be formulated as follows:

$$EO_{\pm d}\ D\phi/\text{not}\ D\phi,\ \text{where}\ \pm d = \sum_{i=1}^{n} \pm v(\circ_i) \times c(\circ_i)$$

When applied to an option that switches on two duties, a positive duty (Om) and a (conflicting) negative duty (O_{-k}), EEDD operates in this way:

1. $\circ_m (B_h\ p \supset D\phi)\ \&\ \circ (B_h\ p)$
2. $\circ_{-k} (B_j\ q \supset \text{not-}D\phi)\ \&\ \circ (B_j\ q)$

_____ by EEDD

$EO_{\pm d}\ D\phi/\text{not-}D\phi$, where $d = (m \times h - k \times j)$

Armed with those resources, I can finally enunciate the revised core thesis of decisional nonconsequentialism: *every agent ought to choose and perform the available option (within his range of mutually exclusive options) that switches on the highest expectational obligation (or permission)*. Put another way, she ought to pick out the option that maximizes expectational deontic value.

Let's see how decisional nonconsequentialism works in *Three Drugs*. Jill should conclude that the highest expectational obligation switched on by any of her options is to give John drug *a*, rather than drug *b* or *c*. This conclusion follows if Jill's killing John has a moral disvalue that is equal to or greater than the moral value of Jill's discharging her duty to give John a complete cure for his condition. This assumption seems altogether reasonable because the harm of death is much greater than the benefit of curing a minor skin condition. In *intrapersonal* cases like *Three Drugs*, one need not be morally concerned with the fact that John is confronted with a duty to aid and a duty not to harm. In fact, Jill is a physician and the negative or positive character of her medical duties toward John is *morally* inconsequential because her professional responsibility can be characterized as one of optimizing John's health. In this particular example (as much as in intrapersonal cases with competing positive duties, such as *Three Gates*) nonconsequentialism normatively overlaps with consequentialism because the stringency of the relevant moral duties directly depends on the extent of the harms and benefits at stake. This is shown by the fact that Jill's highest expectational obligation would be to save John from dying in a hypothetical in which she could not do so unless she gave John a drug that would cause him to develop a mild skin disorder. Of course,

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

this does not preclude assigning negative and positive deontic values to the switched-on duties by considering their negative or positive character. The latter matter is merely notational.

Remember that we have three options: giving John drug *b* (*x*); giving John drug *c* (*y*), and giving him drug *a* (*z*). The relevant factual propositions and corresponding credence levels are: "Drug *b* will cure John completely" (*Cb*), 50 percent; "Drug *b* will kill John" (*Kb*), 50 percent; "Drug *c* will cure John completely" (*Cc*), 50 percent; "Drug *c* will kill John" (*Kc*), 50 percent, and "Drug *a* will provide a partial cure for John" (*PCa*), 100 percent. Thus, the obligations that each option switches on are:

(1) $\circ_m$ ($B_{0.5}$ *Cb* ⊃ D*x*) & $\circ_{-k}$ ($B_{0.5}$ *Kb* ⊃ not-D*x*)
(2) $\circ_m$ ($B_{0.5}$ *Cc* ⊃ D*y*) & $\circ_{-k}$ ($B_{0.5}$ *Kc* ⊃ not-D*y*)
(3) $\circ_l$ (B*PCa* ⊃ D*z*)

In options *x* and *y*, assuming proper norms obligating the epistemic antecedents, we obtain by EEDD expectational obligations *not* to choose *x* and *not* to choose *y* that have a negative value: EO_{-d} not-D*x* and EO_{-d} not-D*y*, where $-d = (m \times 0.5 - k \times 0.5)$. Since the credences supporting the conflicting obligations cancel each other out, we obtain a negative value $(m - k)$ for such negative expectational obligations if the difference between the value of curing John completely (*m*) and the disvalue of killing him (*-k*) is lower than 0. As I said, this is a reasonable assumption. In option *z*, the switched-on moral requirement, that is, "O_l (B*PCa* ⊃ D*z*)," has a positive value *l* (John's partial cure has a value *l*) and the expectational obligation "EO_d *z*" that can be inferred by EEDD from that moral requirement and O (B*PCa*) has a positive value *d*, where $d = l \times 1$.) Therefore, Jill ought to do *z* because this option switches on the highest expectational obligation incumbent on her. (By similar reasoning, it can be concluded that we ought to close gate 3 in *Three Gates*.)

IV. Two Modes of Speech and the Probabilitarian View

Nonconsequentialism can address the problem of decisions under risk on the basis of four different approaches. In this section, I will address three of those approaches, *Zero risk*, *Low risk*, and *Probabilitarianism*, with a focus on negative obligations (Fig. 1). (I will discuss *Epiphenomenalist functionalism* in the next section.) All the approaches can be stated in terms of two modes of speech: a *probability-focused mode* and a *value-focused mode*.

The Zero risk approach requires total certainty that the agent will not violate someone else's right.[56] This implies that the correlative moral

[56] Frank Jackson and Michael Smith, "Absolutist Moral Theories and Uncertainty," *The Journal of Philosophy* 53 (2006): 271; Michael Huemer, "Lexical Priority and the Problem of Risk," *Pacific Philosophical Quarterly* 91 (2010): 333; 337.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

Approaches to risk (moral prohibitions)	Probability-focused mode of speech	Value-focused mode of speech
1. Zero risk	p = 0	Infinite disvalue
2. Low risk	p = changing probability	Large disvalue with no lower and upper bounds
3. Probabilitarianism	p = t (threshold)	Lower and upper bounds on expected value
4. Epiphenomenalist functionalism	p = changing probability	Marginalist value functions

Figure 1: Approaches to risk sensitivity of (negative) moral obligation

prohibition is maximally demanding. For instance, in punishing a person an agent must ascertain that the probability of his being innocent is equal to 0 (Fig. 1). Therefore, in the probability-based mode of speech, any probability level of innocence that is higher than 0 is impermissible as far as the moral prohibition against punishing an innocent is concerned. This approach can also be formulated in the value-based mode of speech by saying that the violation of the prohibition has infinite moral disvalue.[57] In fact, the expected disvalue of any violation whose probability is greater than 0 will then be minus infinity, because any real number times infinity is infinity. As is obvious, this position is impracticable because one can at most only very rarely be certain about any factual condition in making a moral decision.

The Low risk approach is more clearly stated in the value-based mode of speech (Fig. 1). The moral disvalue of violating the prohibition is very large, which means in practice that the permissible risk, though variable, tends to be low. Consider the following example proposed by Jackson and Smith:

> *One Skier*
> A skier is heading in a direction that, you know for sure, will trigger an avalanche that will kill ten people. You know that the only way to save the ten people is for you to shoot the skier. You don't know, however, whether the skier is guilty (that is, whether he intends to kill the ten people.)

Unlike *Three Gates* and *Three Drugs*, Jackson and Smith's hypothetical confronts us with an *interpersonal* moral situation, that is, one in which expected harms and benefits accrue to different persons. In these cases,

[57] Jackson and Smith, "Absolutist Moral Theories and Uncertainty," 272.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

the nonconsequentialist contrast between negative and positive duties does matter. Suppose we assign a large negative value to killing an innocent, and that this negative value more than doubles the negative value of failing to save the ten (i.e., it exceeds by more than 50 per cent the positive value of saving the ten). Though the level of permissible risk is variable here, we can establish that it will be less than 0.5. This equation seems reasonable in the context of the example. The problem for nonconsequentialism in general is that, if one of the variables inversely related to the moral disvalue of breaching the deontological constraint is sufficiently increased (for example, the number of people to be saved), the prohibition gets exhausted. In fact, Jackson and Smith argue that, if the negative value of letting people die is augmented with the number of lives lost, it is easy to construct an example in which the negative value of failing to save people exceeds the negative value of killing an innocent to an extent that is sufficient for making the expected disvalue of failing to save higher than the expected disvalue of killing an innocent.[58] (This is the problem of "loss-inflation" that I mentioned at the beginning.) In this way, the Low risk approach does not capture the core of nonconsequentialism.

Probabilitarianism, the third approach, establishes a probability threshold t that works as follows: when the risk of killing an innocent reaches t, the expected negative value of the violation is at the bottom, that is, the maximal possible disvalue (Fig. 1).[59] Therefore, when the prohibition is activated by t, the expected disvalue of allowing people to die cannot possibly exceed the expected disvalue of killing the innocent (that is, the expected value of saving people from dying cannot offset the expected disvalue of killing the innocent when t is reached). In the example, we are prohibited from shooting the skier if the risk of his being innocent exceeds a certain probability threshold, in which case our shooting is classified within the forbidden class of killing an innocent. In this absolutist version, the fact that our action *is* the killing an innocent rules out trade-offs with positive duties.

Jackson and Smith argue that the probabilitarian view gives rise to a moral gap when we have two symmetrical skiers, X and Y, each posing a risk of avalanche to two groups of ten people. The example is a variant on *One Skier*:

> *Two Skiers*
> Now two skiers are heading in the avalanche-direction and threatening two separate groups of ten people. Again, you know that the only way to save the two groups is for you to shoot the skiers, but you don't know whether each skier intends to kill the respective group. The probability of each skier being innocent is below t, say

[58] Ibid., 273.
[59] Ibid., 275.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

> 0.95 (that is, the probability of both skiers being guilty is higher than 1-*t*). Therefore, it is obligatory or merely permissible to shoot. But the probability of one or another being innocent exceeds *t* (that is, the probability that both skiers are guilty is lower than 1-*t*.)

If we have the kind of relaxed consequentialist attitude toward sacrificing innocents that is entailed by Jackson and Smith's exemplified threshold, it seems that we are licensed to shoot each skier separately, but, since we are not allowed to shoot the two skiers jointly, we are not allowed to shoot them separately either. Jackson and Smith draw on a well-known rule of probability. It is possible to calculate the chances of two or more independent events by multiplying the chances of each event. For example, if each of two independent events has a probability 0.5, the probability of both events is 0.25. Jackson and Smith claim that all the courses of action open to us in *Two Skiers* are impermissible, because we may not shoot the two skiers, we ought to shoot each skier, but it is impossible to shoot each skier without shooting both. Therefore, the argument goes, the probabilitarian construal of deontologism cannot single one option out as morally right. This implication seems intolerable, for the function of decisional nonconsequentialism is precisely to guide action.

One possible answer to the objection based on *Two Skiers* comes from criminal procedural law. Consider an analogous example:

> *Two Defendants*
> We are permitted to convict a single defendant if the probability of his being innocent is below *t*, say 0.1 (or any number you might deem fair). Suppose now we have two defendants who have acted independently. The probability of each defendant being innocent is below *t*, but the probability of one or another being innocent exceeds *t* (that is, the probability that both defendants are guilty is lower than 1-*t*).

Criminal law procedures protect the rights of each defendant *individually*, which means that we are allowed to condemn both defendants. As long as the innocence threshold of each defendant is respected, no right is violated and the legal solution is well determined. Along this line, Aboodi, Borer, and Enoch say that individualistic patient-based deontological theories can rebut Jackson and Smith's paradox by arguing that both skiers (and both defendants) cannot avail themselves of a defense that is unavailable to each of them on his own. The joint normative status of both skiers (and both defendants), they argue, is simply derivative from the status of each of them, since deontological constraints shield individuals, not sums of individuals.[60]

[60] Ron Aboodi, Adi Borer, and David Enoch, "Deontology, Individualism, and Uncertainty: A Reply to Jackson and Smith," *The Journal of Philosophy* 105, no. 5 (2008). They acknowledge that other deontological theories might be affected by the paradox.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

I think that Aboodi, Borer, and Enoch misdiagnose the problem. To reject the paradox they need to defend the following thesis: "The normative status of the complex action shooting-X-and-Y is completely determined by the normative status of the more basic actions shooting-X and shooting-Y."[61] Jackson and Smith dismiss the possible reply that deontological obligations apply to each separate action, not to conjunctions. They argue that in that case we would be allowed to shoot *X* and *Y* with one single bullet, instead of shooting twice, which is clearly absurd.[62] Aboodi, Borer, and Enoch claim that it is the moral status of each person that justifies treating the deontic status of the complex action as derivative from the status of each action, and they accordingly conclude that "individualistic patient-based deontological theories are off the hook."[63] (In this respect, Aboodi, Borer and Enoch reverse Nozick's suggestion that the moral status of risk-imposing actions is based on considerations relative to the agent.) Though the comparison with criminal convictions is attractive, I think that this paradox has no relation to the patient-based or agent-based character of nonconsequentialist duties. In either case, what matters for the moral justification of a decision is the *probability of the minimally sufficient set of factual conditions of the deontic status of the respective option*. Whether the deontic status of the option is grounded on the status of patients or agents does not matter at all. In *Two Skiers*, it must be assumed that the potential actions targeted on *X* and *Y* are *independent* events, otherwise the probability calculus would not yield the paradox by multiplication of the probabilities of each event. Because they are independent events, the decisions we should make in response to each independent threat are consequentially *independent decisions*. Whether we execute our two independent decisions by killing two birds with one stone is a purely technical issue. The minimally sufficient set of factual conditions for determining the permissibility or obligatoriness of the *X*-shooting-decision does not include factual conditions related to the *Y*-shooting-decision, and vice versa. Therefore, our obligation or permission to shoot *X* and shoot *Y*, or shoot *X* and *Y*, is derivative from the moral prescription that applies to each decision.

Suppose now that we have evidence that *X* and *Y* act in concert. Since the two threats are *not* independent, the probability assessment of the complex threat could not be made by multiplying the probability of each one. What we should assess in this case is the probability of the minimally sufficient set of conditions for determining the permissibility or obligatoriness of our response to the joint threat. Our procedure would then be the converse. The status of each basic action would derive from the moral status of the complex action, regardless of whether we ground the latter status on considerations related to *X* and *Y* as patients, or to ourselves as agents.

[61] Ibid., 264.
[62] Jackson and Smith, "Absolutist Moral Theories and Uncertainty," 278.
[63] Ibid., 270.

V. The Grounding of Deontic Risk Sensitivity in Marginal Moral Value

My objection to probabilitarianism is not related to Jackson and Smith's apparent paradox. Probabilitarianism takes the probability-focused mode of speech as basic and derives permissible and impermissible levels of risk from degrees of probability. Although the probability-focused and the value-focused modes of speech are intertranslatable, the latter is normatively fundamental. Probability-focused approaches are merely derivative from the value-focused approaches. The reason is that risk attitudes are reducible to a marginal functional relation (that is, the properties of the moral value function). We should distinguish two kinds of epiphenomenalism coupled with a functionalist approach to the problem of risk: *descriptive epiphenomenalism* and *normative epiphenomenalism*. As is well known, risk aversion and risk preference express themselves through the shape of the utility function (concavity or convexity). Risk aversion can be interpreted as an epiphenomenal aspect of the diminishing marginal utility of some good (graphically, the concavity of the curve), and risk seeking can be expressed by increasing marginal utility (graphically, the convexity of the curve). Descriptive epiphenomenalism holds that risk attitudes express themselves by the curvature of the relevant utility function, which also sums up other epiphenomena (for example, diminishing marginal utility "explains" consumers' preference for diversity). In any event utility relations are logical relations in the context of decision theory.

Normative epiphenomenalism rejects the probabilitarian assumption that it is possible to ground the permissibility or impermissibility of risk-imposing actions on mere probabilities. It is arguable that probabilitarians violate so-called Hume's Law, because no normative conclusion can follow from a purely nonmoral premise (such as a probabilistic proposition).[64] Nozick says in this connection: "It is difficult to imagine a principled way in which the natural-rights tradition can draw the line to fix which probabilities impose unacceptably great risks upon others."[65] My explanation for this difficulty is that premising a normative argument on a certain level of risk puts the cart before the horse. Levels of risk are just credences, that is, mathematical representations of our levels of ignorance. Such epistemic levels cannot justify a principled difference between moral rightness and wrongness according to nonconsequentialism. This is a natural corollary of Ross's characterization of deontologism. Ross says that "[w]hether an act is a duty proper or actual duty depends on all the morally significant kinds it is an instance of."[66] Now epistemic levels are not intrinsic features

[64] Hume did not subscribe to this law. For Hume's actual view, see my article "Hume's Theory of Justice," *Rationality, Markets and Morals* 5 (2014): 47–63 (freely available on the web).

[65] Nozick, *Anarchy, State, and Utopia*, 75.

[66] W. D. Ross, *The Right and the Good*, 19–20.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

of any morally significant kind of act. The reason is that kinds of actions are defined by their objective features, and epistemic levels are not among those features. *Epiphenomenalist functionalism*, the fourth approach (Fig. 1), tries to solve this problem by grounding the relative moral values of expectational obligations in the nonlinear value functions associated with each morally significant kind of action. On this view, the risk sensitivity of a kind of moral obligation is defined as the rate of change of moral value of that obligation with respect to one or more quantifiable attributes of that kind of moral obligation. Obligations have many quantifiable attributes whose marginal valuations vary from one kind of obligation to another. So, for instance, one key attribute of the obligation to aid is the external utility caused by the aid, and one key attribute of the obligation not to kill is the death toll taken by the killing.

Monist ethical consequentialists attach moral value to one single value, namely general utility or social welfare. This means that moral value can change only as a result of increases or decreases of that value. If ethical consequentialists were to contemplate risk attitudes in their moral value function, risk aversion could be expressed by the concavity of the function of moral value and decreasing marginal value, and risk preference would depend on increasing marginal moral value and express itself in the convexity of the utility function. However, I think that, as a matter of fact, consequentialists normatively endorse risk-neutrality with respect to moral value (that is, they reject both risk aversion and risk affinity). Accordingly, their moral value function could be interpreted as a linear function and be graphically represented by a straight line.[67] Some consequentialists might reject this interpretation, but I think it is fair to say that consequentialists such as Derek Parfit, Frank Jackson, and Michael Zimmerman assume that *any additional unit of utility is equaled by an additional unit of moral value*.[68] In other words, consequentialists claim that any increase in utility yields an increase in moral value in the same proportion, and, conversely, that any decrease in utility yields a proportionally equivalent decrease in moral value. Even if that interpretation were inaccurate with respect to some consequentialists, in monist consequentialism there seems to be no good argument for denying that the marginal increase of moral value is constant. (I think that pluralist consequentialism might try to offer

[67] Consequentialists might reject risk neutrality, but I doubt that they could do so without sliding into some form of nonconsequentialism. Rivera Lopez acknowledges: "I do not know any (consequentialist) argument for abandoning moral risk neutrality and endorsing a degree of moral risk aversion," but he nonetheless thinks "that some degree of risk aversion is defensible on consequentialist grounds." Eduardo Rivera López, "The Moral Murderer. A (More) Effective Counterexample to Consequentialism," *Ratio* 25 (2012): 313–14.

[68] Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), 25; Frank Jackson, "Decision–Theoretic Consequentialism and the Nearest and Dearest Objection," 468; John Broome, *Weighting Goods*, 124–26; Michael J. Zimmerman, *Living under Uncertainty*, 19; Ingmar Persson, *From Morality to the End of Reason*, chap. 10; Eduardo Rivera López, "Probabilities in Tragic Choices," *Utilitas* 20, no. 3 (2008): 326.

this sort of argument.) However, that does not mean that consequentialists defend risk neutrality with respect to any utility function. For instance, one well-known utilitarian case for an equal distribution of income is premised on the declining marginal utility of income.[69] But this is different from holding that utility has diminishing moral value

I think that decisional nonconsequentialism should not follow in the footsteps of consequentialism. On the one hand, it should say that there are various kinds of actions that can generate moral value or disvalue in various ways. In effect, nonconsequentialists maintain pluralism and subscribe to various moral duties (or rights) that are paired with different moral value functions.

On the other hand, it is arbitrary to assume that all value functions (corresponding to various moral reasons) are linear, or put another way, that morality is risk neutral. All risk attitudes, including risk neutrality, are *grounded* in marginal valuations of different sorts — personal, moral, and collective. Conversely, all marginal valuations offer the possibility of risk neutrality, risk aversion, and risk affinity. Which risk attitude is adopted depends on the structure of the relevant marginal valuations. For example, if constitutional drafters were risk averse with respect to liberty, they would attach diminishing social value to each additional unit of liberty. Therefore, if the constitutional drafters were to consider a choice between a sure option of a particular quantity of liberty and an option that offers a greater chance of liberty but also a risk of tyranny, they would rationally choose the sure option (as classical liberal drafters actually did).

According to decisional nonconsequentialism, the risk sensitivity of moral duties is grounded in the moral value function associated with the various kinds of morally significant action that the options available to an agent switch on. The moral value functions associated with different kinds of action may be linear or nonlinear. In the latter case, the shapes of the corresponding curves express nonconsequentialist priorities among different kinds of duties.

How is the shape of nonconsequentialist value functions to be delineated? All depends on the moral duties at stake. In such cases as *Three Gates* and *Three Drugs*, the value functions appear to be linear, because the moral value of the relevant duties only depends on the losses and gains that the relevant duties are addressed to avoid or bring about. However, in other cases the relevant moral duties offer nonconsequentialist contrasts that can only be captured by means of nonlinear functions. Consider the following example:

> *Two patients*
> Jill suddenly takes charge of an emergency hospital in a war scenario. She finds two patients, *a* and *b*, who are unconscious. The medical

[69] Richard B. Brandt, *Ethical Theory. The Problems of Normative and Critical Ethics* (Englewood Cliffs, NJ: Prentice-Hall, 1959), 415–19.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

> records have been partially destroyed and all the information that can be gathered is as follows. One of the patients has a minor condition. The other patient has a severe condition and his survival depends on being injected now with a drug that may become lethal if applied to a wrong patient. The available evidence does not indicate whether *a* or *b* is the patient with the severe condition. Yet Jill knows that if she injects the wrong patient, he will certainly die (the drug is lethal to the patient with the minor condition and life saving to the patient with the severe condition).[70]

Since Jill's medical responsibilities do not cover the *interpersonal* distribution of health and life, the negative or positive character of the switched-on duties does matter now. Jill has three options. First, if she does nothing, she certainly lets one patient die, but she refrains from imposing a new 0.5 risk of death on the other. Second, if Jill injects both patients, she certainly saves one patient, but also certainly kills the other. Third, if Jill injects one patient picked out at random, she imposes a 0.5 risk of death on one patient and reduces the risk of death of the other by half. It seems clear that decisional nonconsequentialism must recommend the first option (namely, Jill doing nothing). For this conclusion to follow, it must endorse some sort of moral asymmetry between Jill's active intervention and Jill's abstention (for instance, the asymmetry between killing and letting die). Jill ought not to impose a sure or high risk of death on one patient just in order to reduce the risk already threatening the other patient. The first option is what Jill ought to do in a moral decisional sense. In other words, Jill's highest expectational obligation is to abstain from (actively) imposing a 0.5 risk of death on the patient with the mild condition. How can we explain this solution in terms of the moral value function corresponding to harming and failing to aid in interpersonal cases where special moral duties do not apply?

Decisional nonconsequentialism has an elegant solution for the problem of moral choice under risk. Value functions associated with different kinds of obligation can supply the numerical values appended to obligations that are switched on by various options in decision-making contexts. Expectational obligations can then be obtained from these values and the corresponding credences in accordance with EEDD.

In decisional nonconsequentalism the level of morally permissible risk varies in accordance with the moral value function of each moral obligation, and, therefore, various levels of permissible risk correspond to moral obligations of different kinds or categories. Our intuitive response in *Two Patients* is at odds with a linear moral value function that applies to both positive and negative moral duties. There must be a breakpoint at zero value of (imposed) external utility for the function to represent the intuitive lack of equivalence between Jill's first and third options. In order to express

[70] I am indebted to Marcelo Sancinetti for this example.

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

the asymmetry that justifies the intuitive solution in *Two Patients*, the moral value function corresponding to Jill changing the patients' health or life by her acts or failures to act must be taken to be different from the standard single straight line depicted by consequentialists. By supposing that the priority of negative duties over positive duties is represented by nonlinear moral value functions, we obtain a recognizable form of nonconsequentialism that can demarcate permissible and impermissible levels of risk.

Let's differentiate negative duties and positive duties in the following way: negative duties are duties not to cause an external negative utility, whereas positive duties are duties to cause a positive external utility. (An external utility is a harm or benefit of whatever kind or gravity that is imposed on another.) Only a nonlinear moral value function with a breakpoint that represents the normative asymmetry between causing negative external utility and causing positive external utility can accommodate the fact that, in interpersonal cases such as *Two Patients*, our response to the breach of negative duties is stronger than our response to compliance with positive duties. How can we explain that moral contrast? We can set aside Jill's second option because it involves certainty and, therefore, a standard nonconsequentialist solution applies. That is, Jill may not kill one patient in order to save the other. When we consider Jill's third option, the following seems true. Jill ought not to inject one patient picked out at random, even though the probability of saving one patient or killing the other is just the same (0.5). A nonlinear moral value function can explain the fact that, with respect to a wide range of probability distributions, risking causing a death generates a moral disvalue that is higher than the moral value generated by taking a chance to save a life. It is the first option (that is, injecting neither patient) that Jill ought to choose from a nonconsequentialist standpoint. This option gives rise to the highest expectational obligation, since the switched-on negative duty not to kill has a disvalue that cannot be made up for by the positive value of the duty to save a patient. The greater nonconsequentialist disvalue attached to harmings as compared with the relatively weaker value attached to aidings is a typical feature of nonconsequentialism. In contrast, the linear moral value function generally suggested by consequentialists cannot explain the moral contrast.

Besides expressing moral asymmetries, nonconsequentialist value functions might also represent an aversion to or a preference for risk taking in moral choices. This depends on the particular shape of the relevant nonlinear moral value functions, that is, on the rate of change of moral value related to harmings and non-aidings. In contrast, consequentialists generally assume risk neutrality with respect to moral value. Nonlinear functions can have various shapes, but no philosophical method can capture the particular shape of the moral value function corresponding to any moral obligation. For instance, we do not know whether the moral value function corresponding to duties to aid encodes a constant or a decreasing rate of marginal moral value for any additional unit of external utility for

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

people aided. We also do not know whether the nonconsequentialist value function corresponding to the duty not to harm encodes a constant or a decreasing rate of marginal moral value. However, our intuitions suggest that nonconsequentialism does not embrace risk neutrality, as consequentialism generally does.

If the nonlinear function applicable to positive and negative duties had a curved shape representing a decreasing rate of moral value corresponding to aidings, and an increasing rate of moral disvalue corresponding to harmings, it would also establish the rejection of risk neutrality with respect to moral obligation. In that case, the nonconsequentialist priority of negative duties over positive duties would also involve risk aversion with respect to the fulfillment of positive duties and risk preference concerning the breach of negative duties. There is provisional evidence that our intuitions go in that direction. Consider the following case:

> *Enemy Game*
> You were captured by the enemy and must choose within the following prearranged pairs of options. In problems *A* and *B*, which would you choose?
> A. 100% chance to save one person from dying, or 8.4% chance to save 12 people from dying
> B. 100% chance to kill one person, or 8.4% chance to kill 12 people

I would prefer the first option in problem *A* and the second option in problem *B*. If you shared my moral intuition, your choices would be inconsistent with the standard decision-theoretic model, which would predict and dictate the opposite choice in each problem. In fact, whatever external positive utility we assign to saving from dying, and whatever external negative utility we assign to killing, in problem *A*, the second option gives greater expected moral value in accordance with expected utility theory, and in problem *B*, the first option yields less expected disvalue, also according to the standard model. Since the expected value difference between the two options in both problems is marginal (8.4 times 12 is equal to 100.8), the standard model could also claim to predict indifference reactions, but it could claim to predict neither a preference for a supposedly lesser expected value (problem *A*), nor a preference for a supposedly greater expected disvalue (problem *B*). Thus, intuitive responses seem sometimes to reverse the preference order that expected utility consequentialism would predict and require. If we could confirm such responses in well-designed experimental research, we would be able to conclude that some moral value functions underlying nonconsequentialist intuitions do not express risk neutrality but rather risk aversion or risk preference. This finding would also show that decision-theoretic consequentialism, which assumes risk neutrality, has counterintuitive results, not only in cases of sure utility maximizing violations of moral side constraints, but also in

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

cases of expected utility maximizing violations of risk-imposing violations of side constraints. Risk-sensitive moral asymmetries have commanded little attention so far. (Whether a nonstandard form of consequentialism could reject risk neutrality is unclear, but this is not the place to discuss this issue.)

A marginalist approach to moral obligation might also avoid the problems of "loss-inflation" and "probability deflation." On the one hand, positive duties might have a diminishing rate of increase that approached a limit and, consequently, it would not be possible to augment indefinitely their moral value so as to justify infringing negative duties. Up to a certain point of negative external utility, negative duties might allow trade-offs with positive duties. For instance, mild harmings could be justified as a way of fulfilling stringent positive duties. But beyond that point, negative duties would have a disvalue that could not be compensated by the value of positive duties, since positive duties could not exceed a certain level of moral value when they entered into the asymptotic plateau. On the other hand, the stronger marginal disvalue attached to violating negative duties would make it difficult that potential conflicts between unlikely killings and less unlikely aidings be resolved in favor of taking odds rather than taking no action. (All these appraisals have a ceteris paribus character, of course.)

The approach casts new light on the examples of defendants in criminal courts. The priority of an innocent's right not be punished over the obligation (or permission) to punish offenders, which is a mark of nonconsequentialism, can be represented through a nonlinear moral value function (regardless of the curvature it might have). The same applies to the "rights" of suspect innocent skiers in Jackson and Smith's examples, though these "rights" are so weak (their existence requires probability thresholds that are so high) that they can hardly count as side constraints, and, therefore, their moral disvalue function does not represent a *typical* deontological aversion to shooting innocent people. Instead of an arbitrary demarcation between levels of probability of innocence (or guilt), my approach to risk in moral choice derives normative conclusions from the various nonconsequentialist moral functions that correspond to different obligations. As I said, the possibility that some moral value functions represent increasing or decreasing rates of marginal moral value, and, accordingly, risk aversion or risk preference with respect to different moral duties cannot be ruled out.

Considering the obligations and rights that are relevant to examples of potentially innocent defendants, it seems that the moral aversion to punishing an innocent yields a relatively high moral disvalue as compared with the potential value arising from punishing an offender. Therefore, the expectational obligation not to punish will predictably offset the expectational obligation (or permission) to punish, except when the probability of the defendant's innocence is very low. This does not mean, of course,

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

that probabilistic assessments are dispensable in criminal law. In terms of my account of expectational epistemic deontic detachment, probabilities are necessary in order to calculate the deontic content of expectational obligations. Yet, which specific degree of probability tips the balance in favor of the judge's or jury's expectational obligation to acquit (that is, not to punish) depends on the relative marginal value of the obligation to punish and on the relative marginal value of an innocent's right to be acquitted. The adequate way to treat the choice to condemn or acquit a defendant is no different from the way in which we must treat risky decisions in every context of moral and legal choice.

VI. Conclusion

Nonconsequentialists seem to be trapped in a dilemma. On the one hand, moral philosophers seem to be concerned with the question "What ought I to do?" On the other hand, in nonconsequentialism, objective moral propositions seem incapable of providing answers to that question in concrete ethical situations, that is, in situations of partial ignorance. So what should we do? Should we give up the practical function of morality (and law), or, rather, dismiss objective moral propositions and shift to subjective prescriptions. The former horn of the dilemma is unacceptable. If we abandoned the practical function of morality, morality would become something unrecognizable. The latter horn is no less palatable. If we replaced objective propositions with subjective propositions, how could we possibly justify the latter?

I have explored an alternative route that maintains both the practical function of morality and the objectivity of moral propositions. The solution is to develop theoretical resources that fill in the gap existing between the objective propositions and moral prescriptions. To accomplish this task we need rules of inference that make it possible to withdraw moral requirements and moral prescriptions from objective moral propositions. Those rules of inference are: subjectivist transform, epistemic deontic detachment, and expectational epistemic deontic detachment. The detachment of moral prescriptions further requires epistemic norms.

The central notion of decisional nonconsequentialism is that of expectational obligation. Expectational obligations are logically drawn from moral requirements and epistemic norms provided that we have, in respect of every obligation switched on by an option open to the agent, (1) moral values supplied by the value function associated with the relevant kind of obligation, and (2) credences attached to every proposition relative to the set of minimally sufficient conditions of applicability of that obligation. The main thesis of decisional nonconsequentialism then is that, among the alternatives open to him, an agent ought to choose and bring about the option that switches on the highest expectational obligation (or permission).

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press

Objective moral propositions are indispensable because they justify the moral value that must be given to every switched-on obligation for an expectational obligation to be detachable. Since I have explained away the risk sensitivity of moral obligation as an aspect of the moral value functions of various kinds of obligations (or rights), in the absence of objective propositions that supply those functions, it would be impossible to determine the proper nonconsequentialist solution to moral problems under conditions of risk. Views that try to reconstruct nonconsequentialism in purely subjectivist terms are, therefore, misguided. Not even risk neutrality is morally neutral, for it is just an aspect of a a moral value function.

It is true that traditional moral principles expressed in binary deontic terms do not have the degree of determinacy that is necessary to sustain value functions that specify the rate of change of moral value of each obligation in response to quantitative variations in the morally significant attributes of the relevant kind of obligation. Further research is no doubt necessary. We must investigate the shapes of various moral value functions by testing people's intuitions on "moral lotteries." Nonconsequentialists defend special priority doctrines that indicate the general form of those value functions. But all these doctrines are only sketchy and fragmentary. We still lack a theory about decisions in moral contexts. It is precisely the rough formulation of objective moral principles that explains the role that intuition still plays in moral decisions. Intuition might be the only resource available to establish the different shapes of moral value functions. But experimental work is needed to establish the real patterns of intuition in hard moral cases.

Law, University of San Diego School of Law and Universidad Torcuato di Tella

https://doi.org/10.1017/S0265052516000121 Published online by Cambridge University Press