Action, Deontology, and Risk: Against the Multiplicative Model*

Sergio Tenenbaum

Deontological theories face difficulties in accounting for situations involving risk; the most natural ways of extending deontological principles to such situations have unpalatable consequences. In extending ethical principles to decision under risk, theorists often assume that the risk must be incorporated into the theory by means of a function from the product of probability assignments to certain values. Deontologists should reject this assumption; essentially different actions are available to the agent when she cannot know that a certain act is in her power, so we cannot simply understand her choice situation as a "risk-weighted" version of choice under certainty.

Given the power and generality of decision theory, it is very tempting to think that it provides a model of how we should understand any kind of decision under risk. We must find the outcome, act, or whatever gets a value in our theory and multiply it by the probability of its obtaining given a certain choice. Moreover, given some basic assumptions about the nature of empirical knowledge, it seems that risk is omnipresent; arguably, whenever we act or choose, there is some uncertainty involved. Thus, it seems that any ethical theory, or at least any ethical theory that can guide an agent or at least evaluate the principles guiding the agent's ethical decisions,[1] must be a theory of decision under risk and thus that any such theory must accept what I call the multiplicative model, roughly, the view

1. Philosophers often distinguish between theories of "objective rightness" and "subjective rightness." In these terms, my main concern will be with "subjective rightness."

that risk must enter the theory by multiplying some value by the relevant probability. Although it is not clear how the multiplicative model applies to deontological theories, it might seem that the model must apply on pain of otherwise making deontological theories either vacuous or nihilist. That is, if deontological prohibitions are not sensitive to degrees of risk, then they permit either everything or nothing at all. If any risk of violating a prohibition makes an action impermissible, then nothing is permissible. If only certainty of violating a prohibition makes an action impermissible, then, everything is permitted.[2] But straightforward ways to incorporate probabilities in deontological theories seem to face serious difficulties.[3] My general view is that the multiplicative model has its home in consequentialist theories, and it cannot be adapted to deontological views, or at least not to important versions of deontology. More specifically, I argue, first, that there is no clear rationale within deontological theories to accept the multiplicative model. I will then argue for a rather straightforward way of dealing with risk that typically will not involve any representation of numerical probabilities. The basic idea is that deontological rules, prohibitions, and permissions apply primarily to intentional acts; risk changes the nature of the act, not the probability that the same act will be performed. I then look at how this basic model applies to important cases of risk.

I. STRUCTURES FOR DEONTOLOGICAL THEORIES

Traditional forms of consequentialism start from a theory that attributes value to certain states of affairs and defines right action in terms of maximizing this value. Some deontological theories accept this starting point and define deontological prohibitions and permissions within this framework. On this approach, deontological rules permit, forbid, or require certain actions that do not promote the greatest good.[4] A more radical departure from the consequentialist model denies that we start from evaluating states of affairs. I'll be concerned here with deontological theories that follow this more radical departure.

2. For versions of this dilemma, see, e.g., Robert Nozick, *Anarchy, State and Utopia* (New York: Basic, 1974); Dennis McKerlie, "Rights and Risk," *Canadian Journal of Philosophy* 16 (1986): 239–51; and Frank Jackson and Michael Smith, "Absolutist Moral Theories and Uncertainty," *Journal of Philosophy* 103 (2006): 267–83.

3. See McKerlie, "Rights and Risk"; and Jackson and Smith, "Absolutist Moral Theories and Uncertainty." See Seth Lazar, "In Dubious Battle: Uncertainty and the Ethics of Killing" (unpublished manuscript, Australian National University, 2016) for an attempt to have deontological restrictions interact with probabilities. I discuss briefly some of the views below, but my focus will be in providing a positive alternative.

4. Samuel Scheffler's discussion of agent-centered restrictions and prerogatives follows roughly this model. See Samuel Scheffler, *The Rejection of Consequentialism* (Cambridge: Cambridge University Press, 1994).

A classical version of such deontology would be a Kantian theory. For Kant, the primary object of evaluation is a *maxim*. For our purposes, we take a maxim to be a general rule that connects circumstances, means, and ends ("Whenever in C, I'll pursue M in order to pursue E"). On this view, deontological rules are second-order rules on the choice of maxims (most notably, for Kant, a requirement to will only universalizable maxims). The focus on maxims is characteristic of a Kantian view, but we can have a more general notion of *act types* and take deontological rules to forbid or require the instantiation of certain act types.[5] For instance, deontological prohibitions might govern acts of killing, so one could have a deontological rule prohibiting the instantiation of acts of killing. The rule would be expressed as "Do not engage in acts of killing" or, more simply put, "Do not kill."

A couple of features of this way of understanding deontology are particularly important. First, this rule should not be understood as a rule forbidding the bringing about of a type of state of affairs that could be just as easily formulated as "Do not bring about a state of affairs that involves your killing another human agent." As it stands, this alternative formulation lacks generality; it is a rule that applies only to the current reader of this article ("*your* killing"). To reintroduce generality, we need to either state it as a generalization of many impersonal rules ("For every agent **a**, **a** should not bring about a state of affairs such that **a** kills another human agent") or make it an agent-relative rule ("**a**: Do not bring about the state of affairs in which **a** kills another human agent"). But we have independent reasons to avoid understanding deontological rules as prohibitions on bringing about states of affairs or as ought statements ranging over propositions. This kind of approach makes deontological rules agent relative (and possibly time relative) and leaves them open to the paradox of deontology: if we are interested in preventing certain states of affairs from obtaining, such as the state of affairs in which the agent is killed, why would the fact that the killing was brought about by me (and brought about by me now, rather than later) be of such momentous importance?[6]

5. And, of course, in some cases deontological rules permit act types. My main focus will be on requirement and prohibitions, and I will discuss permissions only when permissions generate notable complications.

6. Scheffler, *Rejection of Consequentialism*, famously argues for the paradoxical nature of deontology, but I think the expression "paradox of deontology" first appears in Christopher McMahon, "The Paradox of Deontology," *Philosophy and Public Affairs* 20 (1991): 350–77. Mark Schroeder, "Ought, Agents, and Actions," *Philosophical Review* 120 (2011): 1–41, argues on semantic grounds against understanding all "oughts" as impersonal ones, and he points out this implication for deontology. For related semantic arguments, see Hector-Neri Castañeda, "The Paradoxes of Deontic Logic: The Simplest Solution to All of Them in One Fell Swoop," in *New Studies in Deontic Logic*, ed. R. Hilpinen (Amsterdam: Springer, 1981), 37–85. It is worth noting that there is a kind of agent relativity that is "preserved"

At any rate, the form of deontology I want to defend takes prohibitions and requirements to range over acts and not over states of affairs, facts, or sets of possible worlds.[7]

Of course, a deontological theory that does not start from the value of certain states of affairs must have some other explanation of the validity of its rules. Different theories would, of course, provide different explanations. I don't aim to provide an exhaustive list of the possibilities here, but listing some of them will be useful for the next sections:

1. *Foundationalism*: Deontological rules are foundational, known by intuition and not grounded on anything else.[8]
2. *Kantianism*: Moral rules are the constitutive conditions of rational (and thus morally good) acts; an act is rational only if willed according to certain rules.[9]
3. *Second-Personalism*: Moral rules are expressions of duties that obtain because of the special relations that persons bear to each other and in virtue of which they can make demands on each other.[10]
4. *Pluralism*: Different values call for different attitudes. Moral rules specify how we ought to relate to certain forms of values for which the attitude of promoting is not (typically) the appropriate one.[11]

———————

in understanding deontological prohibitions as ranging over act types. For instance, we are (I assume) under a prohibition against neglecting the welfare of *our* children. But these are exactly the cases in which there is nothing puzzling about the fact that I have a special moral relation to my children. I don't mean any of these considerations to be decisive but just to provide a possible rationale for avoiding a certain understanding of deontology.

7. Of course, once we have a deontological theory in play, we might find out that it is extensionally equivalent to a form of agent-relative theory whose rules range over state of affairs and use the latter theory in order to model risk. I am only claiming here that the latter theory is unlikely to provide a satisfactory rationale for deontology. For my reservations about trying to "consequentialize" deontological theories, see Sergio Tenenbaum, "The Perils of Earnest Consequentializing," *Philosophy and Phenomenological Research* 88 (2014): 233–40.

8. See W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930). The references in this and the next three footnotes roughly, but not always precisely, correspond to types of views. I am ignoring some important differences among these views.

9. Christine Korsgaard, "Reasons We Can Share," in her *Creating the Kingdom of Ends* (New York: Cambridge University Press, 1996), 275–310; and Stephen Engstrom, *The Form of Practical Knowledge: A Study of the Categorical Imperative* (Cambridge, MA: Harvard University Press, 2009).

10. Stephen Darwall, *The Second-Person Standpoint* (Cambridge, MA: Harvard University Press, 2006); Ron Aboodi, Adi Borer, and David Enoch, "Deontology, Individualism and Uncertainty," *Journal of Philosophy* 105 (2008): 259–72.

11. Philip Pettit, "Consequentialism," in *A Companion to Ethics*, ed. Peter Singer (Oxford: Wiley-Blackwell, 1993), 211–40 (not endorsing the view). Christine Swanton, *Virtue Ethics: A Pluralistic View* (Oxford: Oxford University Press, 2003), uses it in defense of virtue theory.

Of course, there is not much more that can be said in explaining why foundationalists accept the rules that they do. But we can give very rough explanations of why the other views take acts, rather than states of affairs, as the focus of prohibitions and requirements. The Kantian focus on maxims derives from the fact that maxims express the "principle of volition" or the rule guiding the agent, and only principles that conform to the moral law count as rational principles. On this view, the prohibition on lying does not stem from the special negative value that attaches to states of affairs that involve lying. A maxim of lying has a structure that prevents it from being universalizable and thus prevents it from being rationally willed. For "second-personalists," an act puts the agent in a direct relation to the patient, and such relations must be governed by the legitimate demands that rational beings (or possibly sentient beings) can make on each other. I ought not to kill Jane, because this is a demand that Jane is entitled to make on my action irrespective of an independently determined value of the state of affairs in which I kill Jane. For pluralists, although some values call for promotion, other values, such as the value of persons, do not call for promotion but for some other kind of attitude (e.g., respect). We treat persons with respect by acting toward them in certain ways, not by bringing about certain types of states of affairs that are independently valuable. Again these are rough formulations of positions and rationales,[12] but they're good enough for our purposes.

II. MOTIVATING THE MULTIPLICATIVE MODEL

A. *Raising and Lowering Risks*

Let us consider the following vignette:

> **Lower Risk Farmland**: I need to drive my truck from Farawayland to Nowheresville to sell my wheat. Although the towns are not far from each other, there is no road connecting them. In fact, the road connections are so bad that it'd take me a few days to get from Farawayland to Nowheresville even though they are only a few miles apart. However, I can go from one town to the other very quickly if I just drive through farmland. I consider this option but I realize that the local children often run across the farmland. Given how short they are and how high my truck rides, I would not see the children if they're running around.

Of course, it is not permissible to drive through farmland in this case. Suppose now that a few months later I need to drive my truck back from Farawayland to Nowheresville. The situation is identical in all relevant re-

---

12. Another important rationale is the special value that attaches to persons, according to some deontological views. This rationale is arguably shared by these three positions.

spects except for one. It is July, and the local children are on vacation. It is thus significantly more likely that there'll be children running across the farmland around this time. Let us call the situation in my return trip **Higher Risk Farmland**. Suppose I say to myself: "I was convinced that it was morally impermissible to drive through the farmland because of the risk of injury and death to the children. But the situation is importantly different now. The children are on vacation and the risk is much higher. So, perhaps, driving through the farmland is permissible in this situation."

Obviously, something has gone awry here; the risk of killing children being higher could not make the action permissible. Moreover, suppose that the reason why I need to get the truck across farmland is instead a very morally significant one. My truck carries the food for those affected by a devastating famine in Nowheresville. Suppose there's a significant risk that if I do not cross through the farmland, my delivery will come too late. I assume that even if we add to **Higher Risk** this information about the purpose of my trip, it is still impermissible to drive through farmland. But let us now consider **Lowest Risk Farmland**. It is September, and it is the harvest season; the children are not allowed to play in the farmland during harvest season, and this rule is strictly enforced. I am very confident that there will be no children on the farmland, though, of course, I am not certain that there'll be no children there.[13] It seems plausible to say that in such a situation, I would be required to drive through the farmland. Let us suppose now that I learn that the famine situation has worsened (let us call this new situation **Higher Reward/Lowest Risk Farmland**). Due to the spread of diseases, there is a significantly higher chance that the food supplies will not arrive in time if I take the road. Given this new information, the requirement to drive through the farmland must be no less stringent; it could not be right to conclude on the basis of the new information alone that I was no longer required to drive through the farmland.

It is tempting to conclude that as we change the risks to the children and the comparative chances of saving those starving in Nowheresville, the permissions and requirements will covary with the relevant probabilities. Holding fixed the probability of running over the children (assuming a suitably low probability) as we raise the probability that this is the only way to deliver the food in time, we will hit a point in which we are required (or at least permitted) to drive the truck through the farmland. And if we now hold the probability that we can only deliver the food if we drive through the farmland constant, as we raise the probability that there'll be children in the farmland, we'll eventually hit a point in which

13. This is actually more controversial than the "of course" implies here. I am assuming that philosophers sympathetic to what I call the "multiplicative model" would not want to say that I am certain in such situations, but other philosophers might disagree.

it is no longer again permissible to drive through the farmland. It seems easy to conclude that the permissions, prohibitions, and requirements in such a case are weighed by their risk so as to generate an indifference curve, or some approximation of it. Let me define the multiplicative model more precisely as follows:

> (MULTIPLICATIVE MODEL) The (subjective) rightness of a decision or an act is determined by means of a function from the product of probability assignments to certain possibilities and the assignment of certain values to the same possibilities.

In formulating the multiplicative model, I'm trying to be as neutral as possible regarding the general shape of the theory. "Possibilities" could refer to states of affairs, acts, or any other entity. "Values" here should be read as the values of a function rather than "value" as it appears in an axiological theory.[14] The multiplicative model could take the form of a function from the products of the value of violating a prima facie duty and the probability of violation to the value of the action and a rule for choosing among possible actions. Another possibility would be to have two dimensions representing permissions and requirements and have a calculus in which varying strengths of permissions allow acting in face of obligations of varying strengths.[15] Finally, we could have a view according to which we have various absolute duties such that (*a*) the duties are in force only above a certain threshold of confidence, and (*b*) the threshold varies according to the seriousness of the duty. Such views would also count as an instance of the multiplicative model; they could also be represented as a function that multiplies the probability of a certain act with the seriousness of the wrongness (violation) of the act and establishes a cutoff point such that actions above the cutoff point are permissible and actions below the cutoff point are not permissible.[16]

Let us grant a couple of claims that are extremely plausible in light of the cases above; the examples do seem to show that the principles below are true:

14. I assume that the probability assignments express credences, but this is not essential to the argument.

15. For a version of this model in the context of practical rationality more generally, see Joshua Gert, *Brute Rationality* (New York: Cambridge University Press, 2004). Gert does not generalize the model to cases of decision under risk.

16. Nozick, *Anarchy, State and Utopia*, 72–75, puts forward a view roughly like this in relation to rights but raises some difficulties for the view. See also McKerlie, "Rights and Risk," for an excellent discussion of similar views. Of course, even a view in which there is a single threshold of risk for all duties could be represented in this way, but it would be a quite trivial version of the model. At any rate, I find varying threshold views much more plausible than single threshold views.

WEAK RISK PRINCIPLE *[Permission]*: If it is permissible to ϕ in C because by ϕ-ing one is likely to do A,[17] then it is permissible to ϕ^* in C if the only relevant difference between ϕ^* and ϕ is that one is more likely to do A by ϕ^*-ing than by ϕ-ing in C.

WEAK RISK PRINCIPLE *[Requirement]*: If one is required to ϕ in C because by ϕ-ing one is likely to do A, then one is required to ϕ^* in C if the only relevant difference between ϕ^* and ϕ is that one is more likely to do A by ϕ^*-ing than by ϕ-ing in C.[18]

I also think that the related principle below is correct:

MORAL SIGNIFICANCE OF PROBABILITY:[19] There are cases in which whether one is required (permitted) to ϕ as a means to doing A depends on the likelihood that ϕ-ing will bring about that I do A even when all other relevant factors are held equal.

Now these principles do not imply the multiplicative model; at best, they imply that the correct ethical theory exhibits some kind of monotonicity with respect to risk. I'll argue that arguments one can canvas in favor of the multiplicative model depend on assumptions that the deontologist should reject. Of course, I can't here present an exhaustive list of arguments. But I hope they're representative enough to lead us to suspect that there is no argument for the multiplicative model that should appeal to deontologists, or at least the types of deontology proposed above.[20]

Let me start with a very general consideration. The multiplicative model is quite compelling if moral action aims to bring about a certain value to the highest degree. If a moral agent aims to bring about as much of a certain value as possible, it seems natural to compare options by a sum of the various possible outcomes of each alternative weighted by their probabilities in the same way that you compare the expected return of different investment strategies. Although it is not uncontroversial how to make sense of this idea, we seem to be appealing to a clear notion of expected value in this form of reasoning. But obviously this kind of justification is not open to the deontologist. There is no similar obvious mea-

17. I am using the more convoluted "by ϕ-ing one is likely to do A" instead of just "likely to do A" because, as it will be clear below, in my preferred view basic acts do not admit of probabilities.

18. For simplicity, in what follows I refer only to the "WEAK RISK" and let the context determine to which one I'm referring. There is a version of the principle that would apply to actions being permissible despite risking some undesirable act or outcome, but I'm leaving this aside.

19. In what follows, I refer to this principle as SIGNIFICANCE.

20. With the exception of the foundationalist, whose position is obviously neutral on this issue.

surable value that the deontologist is trying to bring about to the greatest degree. So when a deontologist multiplies a probability by a value, it is far from clear what this product represents. If I am bound by a rule that prohibits lying because lying expresses disrespect, does raising the probability of lying always express disrespect (rather than at most incurring a risk of expressing disrespect)? And if risking disrespect is a form of disrespect, is it disrespectful in proportion to the risk? Nothing in the theory implies that we need to answer these questions in the affirmative. In the next subsections, we look more closely at a couple of possible justifications for accepting the multiplicative model within the framework of deontology.

B. Extension of the Weak Risk Principle?

Let us start with a very direct attempt to arrive at the multiplicative model from the WEAK RISK PRINCIPLE. Let us assume that doing A and doing B are things that I should do, all other things being equal. Let us suppose that in cases in which I am certain that I can either do A or do B but not both, I must always choose A. Let us now assume that in a situation in which the likelihood that I do A by ϕ-ing is n and that I am certain that I do B if I ψ, I must ϕ. Now, holding constant my certainty that I can do B, by WEAK RISK, for any circumstance C^* in which the likelihood that I will do A by ϕ-ing is greater than n I must ϕ. But what about circumstances $C^\#$ such that the likelihood I will do A by ϕ-ing is lower than n in C? Now as we consider lower and lower likelihoods at some point it must no longer be required that I ϕ. This would establish our indifference point at a certain probability p so that we could say that ϕ-ing in order to do A should have the same value as the product of p and the value of ψ-ing in order to do B. This would be basically an adaptation of the Ramsey procedure for establishing some kind of continuity in our deontic evaluations in decisions under risk. I think there are many problems with using this procedure, but I want just to point out one that will be particularly important for our purposes. Let us look at a possible principle that a deontologist might find appealing:

> (PERMISSION) If all the relevant actions that the agent might perform ($\phi_1, \ldots, \phi_n$) are such that they carry similar chances that the agent will do A, then if the agent is not permitted to perform at least one of ($\phi_1, \ldots, \phi_n$), then the risk of doing A cannot on its own make any of ($\phi_1, \ldots, \phi_n$) permissible.[21]

---

21. A more precise version of PERMISSION must exclude cases in which an act is impermissible exactly because there is an alternative action in which I incur a lesser risk of causing the same harm. So, e.g., if life-saving pill A causes certain bad side effects in 1 out of 1 million possible scenarios while life-saving pill B causes bad side effects in 2 out of the same 1 million possible scenarios, then it might be impermissible to give pill B to a patient ex-

PERMISSION says basically that very small differences in risk cannot make an impermissible action permissible.[22] Now PERMISSION is compatible with WEAK RISK but not compatible with this kind of continuity. In order to see this, let us assume that the following moral rules are correct:

> (R₁) It is not permissible to seriously endanger a bystander in order to defend oneself against an attacker.

> (R₂) It is permissible to impose trivial risks on bystanders in order to defend against an attacker.

Rules (R₁) and (R₂) are compatible with PERMISSION, at least as long as there are no precise cutoff points for 'minimal risk' or 'seriously endanger'. Suppose that I'm facing a situation in which I could shoot at my attacker, but there is a bystander running from the west toward the trajectory of my bullet; it is very likely that the bullet will hit her if I shoot at my attacker, and since this counts as seriously endangering a bystander, (R₁) deems this action impermissible. But suppose I realize that if I aim a millimeter further to the east, I'll still hit my attacker's heart and very

---

actly because we have available the less risky alternative of giving pill A. I discuss the significance of such cases below, but I am using the simplified version of PERMISSION (for the purpose of the article, all I need is that there is a principle like PERMISSION that is reasonably plausible).

22. It is important to note that PERMISSION is put forward just to show that the weak risk principles do not imply the multiplicative model. Although I take PERMISSION to be an attractive principle for deontologists, it is notoriously difficult to understand in a plausible way the progression from permissibility to impermissibility through options that are pairwise very similar. In particular, we need to block the sorites-like iterated inferences that would lead to the conclusion that if taking some risk is permissible, taking any risk is permissible. There are a number of ways of blocking these inferences that are compatible with PERMISSION. We could say that the truth-values of claims about permissibility are indeterminate within a penumbral area. Or we could have the truth-value of subjective permissibility claims in the penumbral area partly depend on the context of the choice (and, more specifically, on the options available) and perhaps also on the agents' judgments. In other words, the truth of these claims in penumbral areas would be relative to (i) the agent's judgments about the permissibility of the other options available within the range of small variations in risk and (ii) the options available at that point. On this proposal, the permissibility of an option in a context of choice in which the risk ranges between n and n + 2δ would not imply the permissibility of the same option in a context of choice in which the risk ranges between n + δ and n + 3δ—for a related proposal for understanding tolerance principles with respect to vague predicates, see Diana Raffman, *Unruly Words* (New York: Oxford University Press, 2014), 174–75. Of course any such approach comes with costs and difficulties, but that is true of any attempt to understand permissibility in the context of vague thresholds (see Miriam Schoenfield, "Moral Vagueness Is Ontic Vagueness," *Ethics* 126 [2016]: 257–82, for related difficulties). At any rate, my aim here is not to defend a particular version of PERMISSION but just to argue that the weak risk principles do not imply the multiplicative model.

slightly reduce the probability that I'll shoot the innocent runner. According to PERMISSION, it must be the case that shooting a millimeter to the east is also impermissible; this seems quite plausible. But continuity, and more specifically the axiom of continuity (in Von-Neumann and Morgenstern's [VNM] decision theory),[23] implies that this cannot be true in every such case; in some cases, a very slight difference would move me to the indifference point and thus would change the deontic status of my action. On the other hand, this feature of PERMISSION does not make it incompatible with WEAK RISK. Nothing about PERMISSION requires that in any particular circumstance I am permitted to act simply by raising the risk of an option that made my act impermissible.

Of course, I'm not denying that there might be ways of representing a decision procedure that endorses PERMISSION compatible with the multiplicative model; the point is that accepting WEAK RISK on its own does not commit us to the multiplicative model. However, (R₁) and (R₂) also illustrate something important about how a deontologist might understand the relevance of risk. There are acts of killing, acts in which I might possibly kill someone, acts in which I seriously endanger someone, and acts in which I expect no one will get killed but I cannot fully guarantee their safety. The differences among these acts are arguably differences in kind; they are different types of act that I might perform. It seems perfectly coherent to say that it is disrespectful to seriously endanger someone's life, but that we do not owe it to him or her to guarantee his or her safety. The fact that we can place these two acts in a continuum of risk so that we can move smoothly and progressively from one end to the other just shows that, like many other distinctions in kind, this is one whose vague boundaries are crossed by (at least in part) lowering and raising the degree of something; it does not show that fundamentally the difference in degree is the one that matters. It is also worth noting that saying that there is a difference in kind among these acts does not imply that imposing a risk is a wholly sui generis evil. The wrong of imposing a risk is obviously connected to the wrongness or badness of the act in the case in which the risk eventuates.

One might protest that we are having troubles moving from WEAK RISK to a full-blown multiplicative model because we are assuming that the multiplicative model is committed to a level of precision that is un-

---

23. The axiom of continuity says that given three outcomes A, B, and C with values such that A > B > C, there exists a probability p such that the value of B is equal to the value of (p(A), 1 − p(C)). (Of course, the axioms are usually stated in terms of preference; I'm changing to "value" so as to conform to the rest of the article.) So, independently of the values we give to "defending oneself against an attacker," "not defending oneself against an attacker," and "harming a bystander," there must be such an indifference point and thus, given (R₁) and (R₂), PERMISSION could not be right.

realistic. Once we allow that our judgments (or ethical values them-
selves) are imprecise and that our measures are only rough measures,
we can see that the multiplicative model is a natural extension of WEAK
RISK. A more realistic version of the multiplicative model would, for in-
stance, disregard small differences or treat nearby values as (roughly)
equal. Of course, the plausibility of this move will depend on the details
of the proposal, but I want to mention briefly some reasons this might
not fit certain deontological conceptions of the nature of moral obliga-
tion.[24] Let us look at the following case:

> (DROWNING) While riding my boat, I see three people drowning in
> the left bay waving at me asking for rescue and three people waving
> at me in the right bay in exactly the same situation. I can reach at
> most one of the bays in time. I'm very likely to succeed whichever
> way I go, but there is a small chance that I won't make it in time.
> Because the people on the left bay are a couple meters farther away,
> the chances are slightly higher that I will arrive at the right bay in
> time (if I reflected on it, I would estimate the probability as .98 as
> opposed to .97 of arriving in time in the left bay).

Now, I am certainly permitted to go right in DROWNING. But would it be
permissible to reason as follows: given that the differences in probability
are so small, I'll flip a coin and give each group a fair chance of being res-
cued? I find a positive answer intuitive, but, more important, given a typ-
ical deontologist understanding of the value of persons, a positive answer
should be a live possibility. But it is hard to account for this permission by
appealing to the imprecision of the relevant value. Let us say that going
left and going right have roughly the same value.[25] But we also need a
function that will allow us to say that you ought to save three people over
one in a choice between saving three and saving one, and save five in a
choice between five and three, and so on.[26] So we need the relevant value
to be higher as the number of people being saved goes up. And yet, the
permissibility seems to be invariant with the number of people; there
were three people (rather than one) in our example, but there could have

24. Another, more important reason to be skeptical will become evident later; impre-
cise values would not help to account for the "bundling" phenomena I discuss in the final
section. It is also worth noting that on a standard supervaluational account of imprecise
values, PERMISSION would come out as definitely false, given that it would be false under
any precisification of the values. I think this is the wrong result, but one might insist that
PERMISSION is a rough guide in real life situations; a more rigorous account can deny PER-
MISSION without wreaking havoc on our ethical outlook.

25. This already seems strange. The values should be exactly the same.

26. Someone like John Taurek would deny this. See "Should the Numbers Count?"
Philosophy and Public Affairs 6 (1977): 293–316. But Taurek's views would not be any friend-
lier to the multiplicative model.

been four, five, or twenty. But how can the model imply that as the numbers go up, the relevant values are still roughly equal, while at the same time ascribing higher values to saving n people as n goes up?[27]

My qualms here assume that multiplicative models will employ a relatively simple function that gives a value to an act and multiplies it by a probability. One could use much more complex functions, appeal to nonstandard analysis, and employ all sorts of mathematical fireworks. Given enough ingenuity, for almost any verdict the deontologist wants, one could find some mathematical function from probabilities and acts that would match all these verdicts. But the more unwieldy the function, the less likely that it is tracking the right explanation of why certain acts involving risk are wrong; such mathematical functions would be at best tracking judgments that would be justified in a different manner. The idea that what explains the wrongness of certain acts is the risk-weighted (dis)value of certain outcomes seems plausible enough. The more complex the mathematical function employed by a theory of rightness in risk situations, the less likely it is that the theory captures a similarly plausible explanatory relation.

One could object that the multiplicative model might have different ambitions: it might not be put forward as explaining the correctness of our moral judgments (or determining what makes them true) but simply as a decision procedure. But decision theory does not seem to be particularly plausible as a decision procedure.[28] Orthodox decision theory requires a level of precision and mathematical calculation that could not be demanded from a moral agent under ordinary circumstances, let alone in the kind of conditions under which one cannot remedy one's uncertainty; such cases often involve severe time constraints that would further limit the possibility of using decision theory as a decision procedure. And,

27. Again, I do not mean this to be a conclusive reason to reject the multiplicative model. One obvious way to accommodate the permission is to make "roughly equal" comparative; the value of two acts is roughly equal if one is not proportionally much higher than the other. But this suggestion would also face problems. First, a small technical problem: such a suggestion would require a model to use a ratio scale rather than an interval scale (given that proportionality is not invariant through all linear transformations). More important, we would be committed to accepting that the same holds even when there is no risk involved, so that the choice (under certainty) between saving 10,100 or 10,200 would also possibly be one in which both actions are permissible. At the very least one should be able to accept the permissibility of saving the right bay without having to accept that it must be permissible to save many fewer lives if the numbers are large enough.

28. As decision theorists themselves often point out. See, e.g., James Joyce, *The Foundations of Causal Decision Theory* (Cambridge: Cambridge University Press, 1999), 80: "No sensible person should ever propose expected utility maximization as a decision procedure. . . . The expected utility hypothesis is a theory of 'right-making characteristics' rather than a guide to rational deliberation."

of course, if we replace standard decision theory with theories that employ more complex mathematical functions, such theories will be even less likely candidates to serve as a useful decision procedure. At any rate, my argument is against the explanatory adequacy of the multiplicative model; I certainly do not wish to argue that there are no possible mathematical functions, no matter how complex, that would not be extensionally equivalent to the best deontological theory.

C. *Representation Theorems*

Standard utility theory (as well as some variations and extensions) is often justified by means of representation theorems. Representation theorems show that agents satisfy certain plausible constraints if and only if they maximize expected utility (or that the satisfaction of certain constraints correspond to some other structure). Of course, I cannot show here that no such representation theorem is possible. However, if my arguments above are correct, no such representation theorem for deontology is going to be a simple variation of the expected utility theorem.[29] Moreover, there is no clear way of representing the structure of permissions and obligations that most deontologists favor as a partial ordering as needed for standard utility theory. If we represent a situation in which I am permitted to ϕ or ψ as indifference, we'll get problematic results.[30] For instance, for many deontologists, I am permitted to refrain from going for dinner in order to save one person and I am permitted to refrain from saving two persons in order to go for dinner, but I am not permitted to refrain from saving two persons in order to save one person.[31] We could doubtless try to present separate representation theorems for separate orderings of obligations and permissions, and then find some way of combining these

---

29. Graham Oddie and Peter Milne, "Act and Value: Expectation and Representability of Moral Theories," *Theoria* 57 (1991): 42–76, prove a representation theorem that seems to support a version of expected utility for choice of acts that is supposed to apply to deontological theories. But Oddie and Milne take for granted that we can assign real numbers to the values of the acts in a way that is incompatible with many deontological views. For a detailed discussion of the various ways in which this assumption is incompatible with classic deontological views, see Martin Peterson, *The Dimensions of Consequentialism: Ethics, Equality and Risk* (Cambridge: Cambridge University Press, 2013), 174–80.

30. This does not necessarily speak against the project of showing that we can "consequentialize" deontological theories. We can relativize the measure to choice situations to get the right ordering, but this would trivialize the model in such a way that we cannot systematically multiply risks by acts (as acts necessarily vary from one choice situation to another). The "consequentializer" might not care about such trivialization, but as far as defending the multiplicative model goes, such relativization would make the model pointless.

31. Of course, not all deontologists agree with this last judgment (and some might disagree with the other two). But this is a substantive dispute among deontologists.

orderings, but it is not clear how the representation theorems would jus-tify the resulting theory.[32]

D. Pre-Keynesian Rationale and Pragmatic Arguments

The pre-Keynesian rationale says that the multiplicative model has the best results "in the long run." In other words, people faced with risky choices repeatedly through an indefinitely long time would do better if they always followed the multiplicative model. As it stands, the pre-Keynesian rationale is not very plausible, and not only because "in the long run we are all dead." It is unclear how the rationale can be applied to any particular individual case (after all, exceptions here and there do not significantly alter the long-run prospects), let alone cases that are not repeatable. I'll just assume that there is a version of the pre-Keynesian rationale that does not fall prey to the problem.

As noted above, this rationale is most at home in a consequentialist framework. It explains at best the adoption of a multiplicative model in terms of the best expected compliance with a rule or principle or the ex-pected reduction in terms of the number of violations. But since the de-ontologist does not think that maximizing compliance or violation reduc-tion is the aim of following the moral principles she advocates, this kind of rationale would not vindicate the multiplicative model by her lights. Moreover, the problem of the pre-Keynesian rationale extends to any at-tempt of providing a pragmatic vindication of the multiplicative model. Pragmatic defenses of practical principles will show that following certain principles leaves you better off with respect to a certain value, even if they do not directly promote the value in question. So a pragmatic argument in favor of, say, having stable intentions would show that an agent is better off in achieving her goals if she has stable intentions, even if she forgoes some chances of advancing her goals in particular instances in which she does not revise her intentions. But since the forms of deontology we're considering deny that there is any value that is being promoted to a larger degree by complying to the right moral rules, this kind of pragmatic ar-gument cannot vindicate the multiplicative model.

A full rejection of the multiplicative model would require explain-ing better the rationale of various choices of the theory. I also did not con-sider an obvious rationale: the intuitive results that some form of the mul-

---

32. However, there are well-known counterexamples to transitivity for permissions, such as Kamm's case of a surgeon who might face pair-wise choices among playing golf, keeping a promise to her child to spend time with him, and saving a patient (Frances Kamm, "Superer-ogation and Obligation," *Journal of Philosophy* [1985]: 118–38). If one accepts the kind of argu-ments Larry Temkin raises against the transitivity of 'better than' (see, e.g., Larry Temkin, "A Continuum Argument for Intransitivity," *Philosophy and Public Affairs* 25 [1996]: 175–210), one could also generate cases of intransitivity involving obligations.

tiplicative model delivers. Here I rely on various attempts to show that such views deliver a number of counterintuitive verdicts.[33] My aim here is to show that the multiplicative model is not a natural model for important forms of deontology. Next, I want to look at a more natural way to incorporate risk into deontological theories and see whether it gives rise to any difficulties.

III. THE NORMAL CASE AND ITS EXTENSION

In classic formulations of expected utility theory, the correctness of one's acts must depend on one's epistemic states: whether an act maximizes utility ultimately depends on the agent's subjective probabilities.[34] Since we are not certain about any empirical matters, and our actions are part of the empirical world, it seems that there is really no such thing as choice under certainty. So in determining the right act, one must take into account not only the value of each outcome but also the probability that the agent assigns to each possible outcome given each possible act.

Deontological theories, on the other hand, propose rules prohibiting, permitting, and requiring certain acts. At least in some cases, let us call them "the normal cases," the epistemic states of the agents are irrelevant in determining what they ought (are permitted) to do.[35] Let us start from a simple requirement that you cook a vegetarian meal if your guest is vegetarian. It would be strange, to say the least, to try to reformulate the rules in terms of a rule that examines what you know. It would certainly be absurd to think that what you need to do is to cook a vegetarian meal if you know that your guest is vegetarian. One can imagine a meat lover deciding not to call her guests in fear that she'll learn that they're vegetarians. A rule that enjoins the agent to cook a vegetarian meal if she does not know that her guest is not vegetarian would not do much better. The requirement is really to cook a vegetarian meal for your guest if she is vegetarian. And if you don't know whether she is vegetarian, but

---

33. See, e.g., Michael Huemer, "Lexical Priority and the Problem of Risk," *Pacific Philosophical Quarterly* 91 (2010): 332–51; Jackson and Smith, "Absolutist Moral Theories and Uncertainty"; and McKerlie, "Rights and Risk." Lazar, "In Dubious Battle," tries to provide a model without counterintuitive implications, but it provides at most sufficient conditions for permissibility.

34. Assuming that in the relevant version of decision theory the probabilities in question would be subjective probabilities.

35. Yoaav Isaacs, "Duty and Knowledge," *Philosophical Perspectives* 28 (2014): 95–110, proposes that the deontologist's subjective ought should be paired with a knowledge-first epistemology. I agree with Isaacs that a knowledge-first epistemology fits deontology better than traditional decision theory's framework (although I think that the less contentious way of putting the point is that the only epistemic states that matter for a deontologist are knowledge states). Isaacs does formulate the deontological norms as having knowledge states as part of the content of the norms; I find this problematic for the reasons outlined below.

you suspect she might be, you must settle the issue before you can cook dinner;[36] the relevant moral rule tells us that we need to ensure vegetarian meals for our vegetarian guests, and we can comply with such a rule only if we gather the relevant information. Similar things can be said about a rule that prohibits killing innocent bystanders; if you are not sure whether the person next to the terrorist is just passing by or is her bodyguard, you need to check and ensure that it is the bodyguard before shooting. Let us say that a situation in which all the relevant information is accessible to the agent and the agent is in a position to know or at least ought to be in a position to know that this is the case is "the normal case." At least for many versions of deontology, most of us rarely face anything but normal cases. I am almost always in a position to know which promises I made, I can check with my guests whether they have dietary restrictions, and for the most part I'm in a position to know that I am not killing anyone with my actions.

So far as the normal cases are concerned, deontological views will tell us to ϕ or refrain from ϕ-ing. Of course, even in the normal cases, states of knowledge play an indirect role in determining which moral rules apply to us. If deontological rules are rules concerning intentional acts or maxims, in complying with a rule or choosing a maxim, an agent is doing something intentionally. But doing something intentionally, or even setting yourself to do something intentionally, is typically possible only if you know that you can perform the act intentionally. In the typical case, if I do not know that I am ϕ-ing, I cannot be ϕ-ing intentionally. A rule that tells me to ϕ is not a rule I can comply with intentionally if I cannot know whether I will ϕ if I try to ϕ (of course the rule could still tell me to try to ϕ, but this would be a different rule). I could do related things: I could try to ϕ, and I could do various things in the hope that I ϕ (even though winning the lottery is not something I can do intentionally, I can do many things in the hope of winning the lottery).[37] This is true even for simple commands. Suppose my officer tells me "Run at a six minute mile pace." It seems appropriate for me to respond that I don't know whether I can do this. A reasonable officer would revise the command to try my best.[38] Here is one way

36. I am assuming here that, for some reason, you have good reason not to cook a vegetarian meal if your guest is not vegetarian.

37. I am bracketing a number of difficult questions on the nature of intentional action, including, of course, the extent of the knowledge requirement on intentional actions, i.e., whether A-ing intentionally requires that I know that I am A-ing (I discuss some cases in which the knowledge requirement might not hold below). Although many philosophers doubt the knowledge requirement, there is a general agreement that in a wide range of cases your state of knowledge makes a difference to which intentional actions you can perform. This weaker claim is all I need for my argument.

38. Of course, the officer could answer by saying: "I don't care; just do it." But then she would be telling me that she considers me strictly liable to my success; she would not be imposing on me a proper duty to run at this pace.

to put this point (but other deontologists might prefer to formulate it in other ways): moral rules apply primarily to what we can will directly (rather than by means of willing something else), but we can will directly only what we know that we can do by (in) willing it.[39] The requirement to keep one's promises applies to an agent only insofar as she knows that by doing certain things, she'll be keeping her promise.

It is worth noting that the relevance of the agent's knowledge to her deontic situation is a direct consequence of the nature of the intentional acts that are the objects of the relevant requirements and permissions. Focusing on knowledge rather than credence is not some ad hoc move to avoid issues raised by the alternatives; rather, what the agent knows partly determines which acts are available to the agent.

We now have a more direct argument that the multiplicative model cannot be of help to the deontologist. If what is evaluated is the choice of the maxim or the act itself, there is strictly speaking nothing like acting under risk. The rules apply directly to items in the theory that are similar to acts in VNM: there are no probability assignments to acts in classical decision theory;[40] choosing an act is what we do knowingly. In classical decision theory, acts are not evaluated directly but in terms of their possible outcomes; the values of these outcomes thus can be weighted by their probabilities. But since deontology's rules apply directly to the acts, there is no similar space for weighing risk. Let us take, for instance, the case in which I remember that I promised to do something for my wife, but I could not remember what it was. It might seem that in this case I must choose by evaluating the various options in terms of the probability that various acts will bring about that I keep my promise to my wife. But if the primary locus of evaluation is the (intentional) act itself, we must understand this case as involving a change in which acts are available to me, rather than a change in the probabilities of various possible outcomes. Given that I do not know what I promised my wife, there is no act of the form "I am keeping my promise to ϕ" available to me. Of course, I can still, say, buy yogurt in order to avoid the risk that I might not do what I promised my wife, but this is a different act; changing our knowledge states in this view does not change the situation by forcing us to assign different probabilities to the same actions

39. It's worth noting that I'm not claiming that only the intended (as opposing to foreseen) aspects of my action matter to moral evaluation. A rule not to kill can be understood as a command to refrain from killing. I fail in obeying this command if I knowingly kill even if I merely foresee the victim's death.

40. See Brian Weatherson, "Knowledge, Bets and Interests," in *Knowledge Ascriptions*, ed. Jessica Brown and Mikkel Gerken (New York: Oxford University Press, 2012), 75–101, on classical decision theory assumptions regarding what the agent knows.

but changes our situation by changing the nature of the acts available to us.[41]

Now this seems so far quite intuitive; suppose it is somewhat more likely that I promised to buy yogurt than that I promised anything else but that it not very likely that this is what I promised to my wife. It does not seem that I ought to buy yogurt simply because it is the highest value for "expected promise keeping."[42] The rule against breaking promises commands me to find out what I promised my wife, but if this is not possible, then, the rule "keep your promise" no longer applies to me in this instance as it is no longer possible for me to be intentionally keeping my promise to do what I promised her. In cases in which, because of my negligence, I forgot about a promise I made, I have arguably already violated the duty to keep the promise, and certain compensatory obligations apply to me now. If it is not my fault, my situation would be like any other in which I cannot keep my promise because of no fault of my own. If I promise to drive my friend to the airport but my car unexpectedly breaks down, I am no longer under this obligation (or I am at least excused from meeting it), but there might be "remainders" to this duty that impose other obligations on me (a duty to apologize, perhaps a duty to meet my friends' needs in other ways, etc.). Similarly, if it is not my fault that I don't remember what I promised, I might have a number of "remainder duties," but I am no longer under any proper duty to keep my promise.

This also provides an immediate answer to a question pointedly raised by Holly Smith, namely, how deontological theories can justify an appropriate duty to gather information before acting.[43] In fact, she makes a compelling case that a deontological theory that tries to derive duties to collect information by assigning deontic value to the fulfillment and the violation of various duties would deliver very counterintuitive verdicts with respect to our duties to gather information. Smith asks what kind of derivative duty we could have to acquire information (rather than spend our time doing other things and possibly fulfilling other duties). The duty would

41. One could complain here that even on this picture, we'll have to reintroduce risk as part of the content of the new act ("I *might* not do what I promised my wife") so that the resulting view is just a notational variant of the multiplicative model. I address this worry below.

42. Of course, a multiplicative model could avoid this consequence. So, e.g., Lazar, "In Dubious Battle," argues that the expected utility model provides only sufficient conditions for permissibility. There is no doubt that I am permitted to buy the yogurt, but this is a permission that I had anyway; it is far from clear that this expectation of promise keeping could ground a permission on its own. A threshold proposal could also avoid this consequence, but it would have other problems (see below).

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

be derivative since it would be derived from the more fundamental duties whose fulfillment depend on having the information that I now have a duty to procure. She suggests that the deontologist needs a principle along the following lines:

> **SD**. An agent has a subjective derivative prima facie duty to do what he believes is acquiring information if and only if he believes that doing what he believes is gathering information would lead him subsequently to produce the maximum amount of deontic value.[44]

Smith then argues then that under any plausible assignment of deontic value, we get very counterintuitive results.[45] But if what I said above is correct, there is a much more straightforward answer to this question. If I have a duty to refrain from serving meat to vegetarian guests, I must collect information about my guests' dietary restrictions until *I know* whether they are vegetarians. Stopping short of settling this question (if, e.g., I stop gathering information as soon as I learn that my guest lived all her life in Texas, which makes it likely that she is not vegetarian but does not settle the issue) puts me in a situation in which my own negligence prevents me from complying with my duty; if I continue gathering information after I settled the issue (if I ask her friends about her dietary preference after she already told me that she's vegetarian),[46] I'm just wasting time.[47] Of course, one could always gather more evidence in favor of a proposition, but unless we embrace some kind of radical skepticism, in most typical situations, the agent will have at some point acquired the requisite knowledge.[48] But, of course, not all situations are typical, and we need to examine at least two cases in which they are not typical. First, there are cases in which acquiring the requisite information would be impossible or too costly. Second, there are cases in which the agent, due to no fault of her own, might have failed to have the requisite knowledge without her realizing that this was the case.

---

44. Ibid., 24.

45. Smith also argues that a general nonderivative duty of gathering information will fail to deliver the right results. If Smith is right that understanding the duty to gather information in terms of deontic value will not provide us with the correct duties, this would be a serious problem for the multiplicative model, as it would be hard to see how to extend the view to cases of uncertainty without relying on such a deontic value.

46. Assuming that this is a normal case in which her telling me would suffice for me to gain knowledge that she's vegetarian.

47. But doesn't the amount of "research" I need to depend on the seriousness of the obligation? I discuss this objection in Sec. IV below.

48. Total skepticism would cause problems for the multiplicative model as well, as even an expected utility calculus might depend on the agent having knowledge at least of the relevant decision matrix. On this issue, see Weatherson, "Knowledge, Bets and Interests."

Let us say that my guest is completely incommunicado, and she'll remain so until she appears at my doorstep for dinner. Suppose that either (*a*) I have no way of finding out whether she's vegetarian or (*b*) the only way I can tell whether she's vegetarian would involve torturing her evil twin (who knows that my guest is vegetarian but who would refuse to cooperate with me). These would be again cases in which (*a*) the duty does not apply to me because it is not in my power to act in the required way, or (*b*) the duty is defeated by another duty (the duty to refrain from torture). To go back to our promise analogy, these would be cases like the cases in which (*a*) I can only keep my promise if I let someone die on the road, or (*b*) due to no fault of my own, it is no longer in my power to keep my promise (like our example of my car breaking down above). And, like in the case in which I forgot what I promised my wife discussed above, I would not have the same duties, even though I might have similar duties and "remainder" obligations.

What about the case in which I innocently, but wrongly, think that I know my guest's dietary restrictions? Suppose I call my guest, but I talk instead with her evil twin (I do not know she has a twin, and I cannot tell their voices apart) who pretends to be my guest and tells me (falsely) that my guest has no dietary restrictions. Have I violated my duties to my guest? I think there are two possible ways to go, depending on how one understands the modality in question. On the one hand, one could say that it was possible for me to refrain (intentionally) from serving a nonvegetarian meal to the vegetarian guest, as I could have acquired the information that would allow me to perform the relevant act (I could have called again later or investigated further until it would become clear to me that the person on the phone was not my guest). But since my ignorance was blameless, my failure to perform my duty would be in such a case excused. On the other hand, one could say that given my predicament, it was not possible for me to know that my guest was vegetarian and thus that the duty did not apply in such a case as it was not in my power to perform the relevant act. Since for the purposes of the article it does not matter which option we take, I'll not try to settle this issue.

Finally, what if we have something more like a Gettier case? Suppose my guest's twin wrongly believes that my guest eats meat but wants to prevent her from eating meat and tells me that my guest is vegetarian (and in fact she is vegetarian). Given my (innocent) lack of knowledge, did the duty to cook a vegetarian meal for a vegetarian guest apply to me in this occasion? And if it did, did I comply with it? Here again there are a number of ways one could go. As I said above, philosophers mostly agree that, at least typically, one can only do something intentionally if one knows that one can do it. But one might accept that in such cases it is enough to have a true belief in order to be able to intentionally cook an appropri-

ate meal for a vegetarian guest.[49] On this view, this case would just be like a normal case; even if the background condition of knowledge is not present, having the appropriate belief here suffices for my being able to comply with the duty.[50] However, if the correct view of intentional action rules out the possibility that I intentionally perform the relevant action, then I cannot be complying with this particular duty. But, there'll still be closely connected duties that I do comply with, most obviously the duty to try to serve a vegetarian meal to a vegetarian guest.

It is important to note, first, that the different views are not, at least in the first instance, different views about which epistemic states are relevant in performing independently described intentional acts. They are different views about what acts are available to the agent. Given that the deontological theories we are discussing take deontic rules to have act types as their content, different conceptions of action will lead to different (though often related) views about the duties involved in a particular situation. But in the end this amounts to no more than different forms of bookkeeping, as all the plausible views can accommodate in some way the main judgments about blameworthiness, compliance, and permissibility that we want to make in these situations.

One might object that the view I am proposing depends on implausible commitments about the transparency of certain mental states. So, for instance, Jackson and Smith discussing a similar proposal say: "Although it has been suggested to us in discussion that absolutisms of the kind in question avoid the troubles we have been raising on the ground that subjects can be certain of what they intend, this is a mistake. We do not have incorrigible knowledge of our mental states and the point is especially obvious and widely agreed in the case of intentions."[51] But this would misunderstand the nature of the proposal. Jackson and Smith are assuming here that the content of the choice includes "I will intend to ϕ" as one of its possible outcomes, as if "I will intend to ϕ" were one

---

49. Anscombe herself, who famously defended that in acting intentionally one knows what one is doing (without observation), clearly allows for this possibility: "if one does certain things, say ABC, then it is possible to have the intention of doing Z in doing ABC; and if the case is one of knowledge *or if the opinion is correct*, then doing or causing Z is an intentional action" (G. E. M. Anscombe, *Intention* [Cambridge, MA: Harvard University Press, 2000], 50, emphasis mine).

50. On this view, of course, in cases in which the agent's epistemic standing is compromised (such as in Gettier cases), (all-out) belief states might be relevant at least in evaluating whether the agent complied with a duty. For simplicity sake, I will assume the view described in the following sentences in the remainder of the article, but nothing in the overall argument depends on this assumption.

51. Jackson and Smith, "Absolutist Moral Theories and Uncertainty," 279.

of the outcomes in a VNM decision matrix.[52] But the claim is that we choose by acting with an intention (or by choosing a maxim), and this act (or choice of maxim) is directly assessed. Of course, in reflecting on my choice I might make all kinds of mistakes. I might think that I am trying to run at a seven-minute-mile pace, when in fact I am barely putting an effort to move my legs. I can even be a guilt-ridden version of Davidson's climber who thinks that he is letting go of his colleague intentionally when in fact my grip on the rope was relaxed by the spasms that resulted from my awareness of certain beliefs and desires. But all this is irrelevant; according to the deontologist, whether I complied with my duty is determined by what I did (or failed to do) intentionally, not by what I believe I did intentionally.

Of course rejecting the idea that an agent apply deontic rules by calculating risks of violation is a purely negative move; we still need to explain how a deontological theory can make room for cases in which differences in risk clearly make a difference to our obligations. In fact, so far it is not clear how deontology can accommodate SIGNIFICANCE or WEAK RISK. Worse, we might seem to be back to square one. Unless all relevant differences in risk affect the background knowledge required to perform an intentional action, it might seem that the only way we can accommodate SIGNIFICANCE is by restoring the multiplicative model by specifying in some way the content of the relevant acts in terms of probabilities.

If we're right about the advantages for the deontologist to have act types as the primary object of evaluation, the easiest way to preserve this advantage in extending the theory to cases of risk is to preserve this structure. That is, we need to take rules governing "risky actions" as governing act types different from, even if related to, "nonrisky" act types. Of course, in describing the relevant act types, the types will very likely be essentially connected to risk. But there is no reason to think that we'd have to go back to formulating action in terms of bringing about states of affairs multiplied by a certain probability. We'll need to describe acts as "endangering," "imposing risks," and so on. But we need not assume that such language is just a convenient way of describing precise, or even imprecise, probabilities.

Let us take, for instance, a deontological theory that accepts the following rules:

52. Mark Colyvan, D. Cox, and K. Steele, "Modelling the Moral Dimension of Decisions," *Nous* 44 (2010): 503–29, make a similar assumption (in fairness, they're trying to find an expected utility model that fits deontology, not criticizing deontologists for assuming transparency).

1. One ought not kill (One ought to refrain from killing).[53]
2. One may not endanger someone's life except in order to secure a very significant good.
3. One may impose trivial risks to others in order to secure small goods.

Rule 2 might allow me to engage in a high-speed car chase to stop a murderer but not to ensure that another driver will apologize for cutting me off. Rule 3 might allow me to take my child to a public place despite her severe nut allergies or to let her come back on her own from school or even to take her rock climbing or white-water rafting for our vacation. But probably neither 2 nor 3 allows me to hold my baby outside the window of the top floor of a high-rise so that she can get a bit of fresh air.

One might suspect that the difference between 1 and 3 is a difference in the probabilities of the death of the possible victims of one's action (and of course a difference in the amount of good consequences brought about by my actions). In 1, we are talking about cases in which my action is certain, or nearly certain, to result in the death. Rule 2 covers cases in which the probability is not as high but still high enough, and 3 covers cases in which the probability is very low. But once we think about matters this way, it seems that this way of talking is just an approximation. A more precise version of these rules would assign numerical probabilities to the risks and more precise values to the acts; we would be essentially endorsing the multiplicative model.

However, this identification of these difference act types as probability intervals of possible outcomes brought about by my actions is too quick. First, there is the issue that, as we said above, whether a prohibition to kill applies to me depends on what I know. But it is at least dubious that the difference between knowing and not knowing can be captured in terms of differences in credences. Second, it is not clear that the differences in probabilities or credences are the only relevant ones in determining whether an act falls under the category of "endangering." Given my driving and shooting skills, it might be more likely that someone will die while I am driving him somewhere than if I fire a few shots at him from ten meters away.[54] Still, my giving him a ride home does not count as endangering him, but my shooting at him does count as a form of endangering. And it does count just as much as endangering even if my ultimate aim in both cases is the same (I would do either to entertain my baby, who laughs just as much when I'm firing at someone as when she sees some-

53. An absolute prohibition on killing of any kind is rather extreme and implausible even among absolutist views. But the extreme version makes the presentation much simpler.

54. Just to make it clear: my driving skills are below average but not so bad as not to allow me to drive; on the other hand, my shooting skills are truly phenomenally bad.

one nervously riding in my car). And even in our examples above, I have no idea what the differences in risk are between holding a baby out of the window and white-water rafting, and, similarly, I don't have a clear sense of how the benefits of fresh air for a baby compare with the benefits of rafting for a young child. But I am reasonably confident that the former but not the latter count as endangering a minor in a morally objectionable way. Of course, it is no miracle that we see a difference between these cases; I would hazard a guess that the enjoyment of white-water rafting is inseparable from the risks involved in the activity, and that driving is an essential means to many activities, makes a difference to what counts as endangering in the relevant sense. The question is whether "endangering" reduces to the combination of a probability and an independent morally relevant factor.[55]

Finally, and most important, even if such a reduction were possible, we have no reason, independent of prior allegiance to the multiplicative model, to deny that the important differences are genuine qualitative differences in the nature of my actions. That is, even if we think that the quantitative differences in risk determine the classification of the different actions in one of these categories, the quantitative differences might matter only because they correspond to these qualitative differences and not the other way around. If this is the case, the deontologist is still on safe ground when she says that the rules apply primarily to act types. Moreover, this order of explanation is independently plausible; in fact, it explains the plausibility of PERMISSION. If the qualitative categories are essentially vague,[56] then there is not going to be a precise cutoff point of risk for an action to qualify or fail to qualify as case of endangering. In such cases, it is possible that a small quantitative difference between the probability of two options cannot justify assigning them to different deontic categories.

55. For some interesting parallels here, see Timothy Williamson's discussion of a similarity-based, as opposed to a probability-based, conception of safety and danger in "Probability and Danger," *Amherst Lecture in Philosophy* 4 (2009): 1–35. In fact, I suspect that whether these acts count as endangering depends on how they "bundle" with other acts. That is, if we in general pay no attention to small risks such as holding a baby out of the window, the cumulative risk to the baby would be very high (while activities such as white-water rafting don't "bundle" in the same way as the risk is internal to the benefit). However, as I discuss later, the relevance of such bundles is itself a problem for the multiplicative model, as the probabilities of harming my baby in the case I am holding him out of the window is independent from the probabilities of bad outcomes of other risky actions. So it is unclear how such "bundling" could justify distinguishing this action from the white-water rafting case. This parallels a difficulty for how expected utility theory treats repeated gambles in nonmoral cases. On this issue, see Sergio Tenenbaum and Diana Raffman, "Vague Projects and the Puzzle of the Self-Torturer," *Ethics* 123 (2012): 86–112. See also Sec. V.

56. Epistemicism will make us reformulate the claim but will not change anything fundamental.

Whether this basic framework is plausible depends on how well it can answer a number of objections against attempts to account for risk in deontological situations. In the next sections, we look at how the proposal fares in these challenging cases.

IV. RAISING THE STAKES

Let us say I own a llama, Larry, and I charge children $20 for petting it. Little Laura is up next, and Laura's parents explain to me that she hates it when llamas spit on her. That happened to her last summer, and she was really upset; in fact, it completely spoiled their vacation as she kept crying for days. I explain to them there's nothing to worry about; Larry does not spit. They tell me that it would be wrong of me to earn my livelihood this way if this involves a risk to Little Laura. I (honestly) agree. But I tell them that there is no risk; Larry is not like the nasty llama Laura tried to pet last summer. Larry does not spit. Let us assume that I know that Larry will not spit. So my preferred deontologist would explain the rightness of my actions by noting that (I know that) Larry will not harm Laura, so I am not infringing on a prohibition to harm children for profit.

However, the advocate of the multiplicative model can now raise an important objection. She can say that it must be my *credence* that Larry will not spit that is playing a crucial role here. For suppose a friend tells me: "you know, her parents refuse to believe it, but Laura is actually severely allergic to llama saliva. If a llama spits on her again, she'll be in agonizing pain for the rest of her life." On learning this, I might now have an obligation to refrain from letting Laura pet the llama. The obvious explanation for the changing obligations seems to be that the relevant epistemic state in evaluating my obligations was my high credence in, rather than my knowledge of, the proposition that Larry would not spit. Once I found out that the potential harm was much higher than I thought at first, I realized that the prospect of a small benefit for myself could not override or cancel the expected violation of a much more serious duty. It seems that the position I defend cannot explain the situation. After all, my justification for my belief did not change; I still have the same knowledge that I had before.

However, it is far from clear that I still have the "same knowledge." This is a case in which rising stakes might have undermined my knowledge that Larry would not spit. Of course, I cannot settle here why, and even whether, knowledge is hostage to such issues,[57] so I will rest con-

---

57. It is worth noting that one can accept that rising stakes will undermine knowledge in such situations even while denying that knowledge is sensitive to pragmatic considerations. See Jennifer Nagel, "Epistemic Anxiety and Adaptive Invariantism," *Philosophical Perspectives* 24 (2010): 407–35.

tent in pointing out that our judgment that I should not let Laura pet Larry seems to go hand in hand with the judgment that I no longer know that Larry will not spit. Thus, in such a case, I would be on my account in violation of a prohibition to refrain from endangering children for profit. The important point is that the deontologist need not concede here that there is a continuous increase in "degree of wrongness" or on the value of a function that determines whether the action is wrong; insofar as we are no longer in a position to say that I know that the llama will not spit, the nature of the acts available to me has changed.

V. BUNDLING AND CLOSURE

Some of the most serious obstacles for accommodating risk in a deontological theory are cases of repeated risk, more particularly, cases in which one or many agents impose relatively small risks repeatedly on the same or various agents. Suppose I conclude that I am not permitted to shoot Mary even if I take two bullets out of my six-bullet chamber in my gun. Disappointed with this result, I ask myself whether it would be permissible to shoot six separate guns with one bullet in each. The answer is obviously "no." Given that the risk is (approximately) the same, it seems to make no difference whether I do it in one step or carry out my plan performing multiple actions. And, of course, the problem is not that we did not spread the risk through a significant enough number of actions. It would make no difference whether I did it by shooting two-hundred-and twenty-two-hundred-bullet guns with one randomly placed bullet in each of them. Reduction of compound lotteries follows from standard expected utility theory's axioms, but the kind of deontologist we described seems committed to the view that it must make a difference whether a risk of killing is spread through multiple actions or borne all at once.

The deontologist might argue that this is an advantage of the view. A system of criminal punishment will over time almost certainly condemn some innocent people. Still, it seems plausible to say that despite the serious harm the justice system will cause to certain people, there is no general deontological prohibition against a system of criminal punishment.[58] However, things would be different if the risk were concentrated in one person and in one action. We would not replace our system with one such that we would punish Jane who is (almost certainly) innocent but that would somehow allow us to reduce the number of people who are wrongly condemned overall.

---

58. See Aboodi, Borer, and Enoch, "Deontology, Individualism and Uncertainty," for this example and a similar point.

I think this is essentially correct, but there is a serious difficulty here.[59] If there is a threshold probability for permissibility of, say, killing an innocent person in self-defense, as we look at multiple actions near the threshold we must find cases in which we are permitted to perform each of a certain series of actions, but we are not permitted to perform the conjunction of all these actions. Now this problem might seem not to apply directly to a view for which only knowledge states are relevant. Let us assume that there is indeed an absolute prohibition to kill innocent threats (but not culpable aggressors). If I hold a threshold view, I am committed to the view that prohibitions against killing innocent threats hold when the probability that a potential victim is an innocent threat is higher than a certain threshold probability p. But if this is the case, there must be possible scenarios in which I am facing two threats and such that the likelihood in each case that the attacker is an innocent threat is p − ε; in such cases, given that the probability that at least one of the attackers is an innocent threat is higher than p, I'll be in a position in which I am able to kill the first attacker and, after killing the first, I am allowed to kill the second attacker, but at the same time, I am not allowed to kill both attackers. But if we rely on a knowledge norm, this problem does not arise. Let us formulate the principle as follows:

> (INNOCENT) It is permissible to kill a culpable aggressor in self-defense, and otherwise killing is absolutely forbidden.

The permission in INNOCENT will apply to me as long as I know that the threat I face (and the victim of my defensive action) is a culpable aggressor, otherwise the permission does not apply. In the case above, if I know in each case that I am shooting at a culpable aggressor then I am permitted to shoot each of them. Otherwise, the absolute prohibition on killing applies to my case. Since the numerical probabilities play no role in determining the permissibility, the possibility that they will add up in the problematic way described above does not come up here.

However, there is a similar, if less extensive, problem generated by difficulties with epistemological closure. Suppose I go to battle committed not to violate INNOCENT. Let us assume that at each time, early on in the war, I know that my victim is a culpable aggressor. But I also know that the longer the war goes on, the more likely it is that I'll make a mistake and use lethal force against an innocent civilian. If the war is long enough, at some point I won't know that I haven't killed any innocent civilians. There are difficult questions in epistemology about how, and

---

59. See Jackson and Smith, "Absolutist Moral Theories and Uncertainty"; and McKerlie, "Rights and Risk," on this issue.

whether, these various epistemic states are compatible. But however one approaches these issues, these kinds of scenarios pose a problem for how our deontologist can formulate the rules that apply to me in such cases. Before I go to war, if I take the long view of things, I would possibly conclude that it is impermissible to go to war. After all, in going to war I am engaging in a course of action in which I'll be killing people, and since I don't know that I'll kill only culpable aggressors, the absolute prohibition applies to me. But if I don't reflect on it, go to the war, and shoot only at the people that, at the time of the action, I know to be culpable aggressors, I do nothing wrong. It seems that whether it is permissible for me to go to war depends on whether I reflect on the combined effects of my actions or restrict myself to thinking about them one at a time. But this is absurd.

Aboodi et al. try a solution to this problem for forms of deontology that focus on the claims that particular persons have on my actions.[60] So I ought not to break my promises to Jane because Jane has a particular claim on, or right against, me. If I make a promise to Barry, then Barry also has a claim on my actions. But there is no claim that the "combined agent" Jane-and-Barry has on me. Aboodi et al. propose a view in which the threshold for a permission is "moral certainty";[61] in our case, I am permitted to kill each person as long as I am "morally certain" that she is an enemy combatant. In reflecting on the series of actions ahead of me, I might come to the conclusion that it is very likely that an innocent person will die. The probability that A_1 or A_2 or A_3 or . . . is an innocent threat is very high. But since claims are made individually, there is no claim that the disjunctive agent A_1-or-A_2-or-A_3 has against me.

Now there are a number of problems with this proposal. For instance, it is not clear that it solves the problem in its full generality, as it does not seem to extend to cases in which I will repeatedly face the possibility of imposing a very minimal risk on a particular person whom I know to be innocent.[62] But I want to focus on a different and, I think, more telling issue. When a certain form of deontology defends nonaggregative claims, it makes sense to appeal to the idea that claims, rights, and so on, need to pertain to particular individuals. So I can violate Mary's rights, or I can violate Kerry's rights, but I cannot violate the rights of a conjunc-

---

60. Aboodi, Borer, and Enoch, "Deontology, Individualism and Uncertainty."

61. I think there is a way of understanding the idea of "moral certainty" that would not make this a threshold probability view. "Moral certainty" might be a view about how much evidence is needed to settle on the question whether **p** or form the belief that **p**. Since Aboodi et al. use the notion as a probability threshold, I am ignoring this possibility.

62. See Huemer, "Lexical Priority and the Problem of Risk," for this criticism. See Lazar, "In Dubious Battle"; and McKerlie, "Rights and Risk," for a number of other difficulties with this approach.

tive or disjunctive agent Mary-Kerry. But, as we said above, imposing a risk is not a sui generis form of wrong; ultimately the wrongness of imposing a risk must be explained in terms of the wrongness or badness of the act in the case that the risk eventuates. If the risk eventuates, there will be someone in particular who is the patient of the act. So we cannot just say that there is no particular person that we are endangering when we distribute the risk and thus no one in particular that we're wronging. At least in the case that the risk eventuates, there is someone in particular that has a claim on being wronged in a way that you must have anticipated.

In fact, we can see that this aspect of the view leads to very counter-intuitive results. A 1 in 7 billion chance that I might kill my personal trainer driving him to the athletic store is, I assume, above the "moral certainty" threshold on the prohibition on killing a person in order to get new running shoes. But suppose I am told by city officials that there is a dangerous virus in my car that will be released to the atmosphere for an hour if I drive now. Everyone is immune to the virus most of the time, but at any time there is exactly one person in the world who is not immune to it. Unfortunately, the virus will move from host to host very quickly throughout the entire surface of the earth until it lands on a person who is not immune at the time and then will kill this person (and it will die with her). Since the virus moves incredibly fast, the probability that someone will be killed within an hour is very close to 1. Thus, if I drive to the athletic store, I am virtually certain to kill someone within an hour, but for each particular person the chance that I will kill her is roughly 1 in 7 billion. But surely I am not permitted to drive to the store in this situation.[63] Aboodi et al. are correct that in some cases the accumulation of risk does not seem to make each action impermissible. But they move to the conclusion that the fact that the risk is distributed across potential patients makes the series of actions permissible. However, the example above suggests that if there's a morally significant difference between such cases and cases of a single action with the same total risk, the difference stems from the fact that there are multiple acts that result in a "bundled act" that is essentially different from the original act. In the virus case or the multiple guns cases, the bundled acts are not essentially different acts from the single action in which I endanger someone to the same degree. On the other hand, declaring war and instituting a criminal justice system are acts that are essentially different from engaging in particular acts of war or conducting a trial. But accepting that, even in some cases, we must treat the bundled acts the same way as single acts carrying the

---

63. In fact, I think that you are not allowed to divert a trolley that is about to kill one person, if by doing so you'll release this virus, but I don't need anything nearly as strong to make my point.

same risks seems to lead us back to our original problem. The deontologist now seems to be committed to the view that we could have a series of permissible actions such that each act in the series is permissible but the whole series is impermissible. In the rest of this article, I will try to show that this is an advantage of, rather than a problem for, the view.

Let us distinguish two related types of cases:

a) APPARENT CLOSURE FAILURES: Cases in which I know that **p**, I know that **q**, I know that **r**, . . . , but I seem not to be in a position to know that **p** & **q** & **r** & . . . .

b) MOVING UP AND DOWN QUALITATIVE RISK CATEGORIES: Cases in which, for instance, ϕ-ing poses a minimal risk, ϕ_1-ing poses a minimal risk of doing A, ϕ_2-ing poses a minimal risk of doing A, . . . , but ϕ-ing and ϕ_1-ing and ϕ_2-ing and . . . makes it likely that I (we) will A.

In both types of cases, we have a "bundling problem." Our long war case was an instance of (*a*); in that case, it seems that I reach two different conclusions when I take the long view and when I deliberate act by act. The following case would fall under (*b*): I have a dog that I take to work and I know that one of my coworkers in this large corporation is allergic to dogs, but I don't know who he or she is. However, the chances that anything bad will happen are very small. Her allergic reaction would be triggered only if she and my dog were at a small area for a long period of time, and the only way this could happen in my office building would be if the two of us were stuck in a malfunctioning elevator. But the chances of elevators malfunctioning are very small and the chances that I'll happen to be with this coworker in the same elevator when this happens are even smaller. Let us assume that it is permissible to impose this kind of minimal risk on my coworker. But now I get promoted and my new job is to be the liftman. In my new job, I spend all my working hours going up and down in the elevator. Each time I go up or down the elevator, the risk that the relevant coworker is riding and the elevator gets stuck is, of course, just as small as when I used the elevator to get to my desk. But the cumulative risk of my always having my dog with me is now significant. Again, it seems here that I might form a different conclusion for my actions if I look at them as a bundle or if I simply approach them one by one. It turns out that bundling acts together in this manner can be morally significant; committing, or refraining to commit, to a series of actions, or even having or refraining from having a general policy, is itself a different act; thus, the possibility of bundling might bring a further dimension to the permissibility of actions. In other words, it might be impermissible to perform all the actions in a series, even if each action in the series would be permissibly performed if all these actions were performed; the fact that the moral significance of bundled acts cannot be easily captured by the

multiplicative model (exactly because we are bundling causally independent acts) is actually a problem for the multiplicative model. Of course, once I know about this possibility, I cannot decide to ignore that I'll be performing the bundled act. I can't insist on taking the "short view of things" any more than I can decide to ignore my knowledge that by driving in Farmland I will be endangering the children; the bundled acts will be among the things I'll be doing.

A similar phenomenon is present even in prudential choice. I can't fully argue for this claim here,[64] but I hope to show at least that this view is much more plausible than it might seem at first. Suppose I am planning to go to a concert, and it would be very convenient to bike to the venue. I realize that cars in my city are not very aware of the bikes around them. However, I am a very careful biker; I conclude thus that the risk is small enough that it makes sense to bike. Suppose a week later I get a new job, and my new job happens to be at the same concert venue. Of course, biking each time is just as convenient and risky as on the night of the concert, but I realize that if I bike every day in the next ten years to and from work, there is a significant chance that I'll be involved in a major accident. I conclude that, given these risks, I should adopt a policy not to bike to work. Assuming that the utility of biking and of suffering an accident does not change, and that the probabilities of being involved in an accident at each time are independent, my preferences are not compatible with expected utility theory. However, there is nothing clearly irrational about these preferences; in fact, there is nothing clearly irrational about having a policy of not biking due to the combined risk of biking every day, and making exceptions to the policy here and there, even if the utilities of dying and biking to school in different days are independent of each other (as they plausibly are). These preferences indeed come out as irrational in standard utility theory;[65] the same values for the acts or outcomes multiplied by the same probabilities should result in the same utility.[66] But this is a consequence of classical decision theory idealizing away vague and indeterminate ends, not of the intrinsic irrationality of this combination of attitudes. But if we reject this idealization, it seems possible to have long-term ends not to expose oneself to significant risks, while having short-term preferences that allow you to take minimal risks.[67] In these cases, the rationality of a cer-

---

64. See Tenenbaum and Raffman, "Vague Projects."

65. But not in all forms of utility theory. See, e.g., Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2004).

66. For a classic, formal proof of a more general theorem, see Paul Samuelson, "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia* 57 (1963): 50–56.

67. Or so we argue in Tenenbaum and Raffman, "Vague Projects."

tain pattern of activity does not supervene on the rationality of the constituent act.[68]

If we agree that, in the prudential case, sometimes the rationality of a certain pattern of activity does not supervene on the constituent acts, then we should, analogously, allow that the deontic status of a pattern of activity might not supervene on its constituent acts. Indeed, a structure very similar to these prudential cases of nonsupervenience appears in both our (*a*) and (*b*) cases. Let us start with our elevator example. As I take only the occasional trip up and down the elevator, I do not endanger my coworker in a significant manner. Each trip is morally acceptable, and if the risk eventuates, I would certainly owe an apology or compensation to my coworker, but my action is excused. It would be right for me to answer the coworker's complaint by saying: "I am so sorry for what happened. But that was a tragic coincidence."[69] Once I accept the job, I now know that repeatedly coming with my dog will impose a significant risk that my coworker will suffer a serious allergic reaction. Were the risk to eventuate, the above excuse would be unavailable. My coworker could rightly say that given that I was coming in with my dog every day, this was likely to happen sooner or later. Note that, here too, adopting the policy of not bringing my dog to work is compatible with making a few exceptions. If I very rarely bring my dog, it would still be legitimate to use the above excuse if the risk were to eventuate. If we look at my biking example, given my ends, arguably, the rationality of my pattern of activity does not supervene on the rationality of each momentary choice. Every single time I bike in, I am instrumentally rational. But if I bike in to work every day I am acting irrationally.[70] Similarly in our elevator case, it is permissible for me to bring my dog in at each particular instance, but the series of actions such that I bring my dog in every day is not permissible.

Of course, some times there is no point in engaging in a particular action if we do not intend to repeat it, so the impermissibility of the series of actions will "trickle down" to the impermissibility of each particular action. My war case might have this structure. In general, if we think

68. Of course, you could try to add some value to the pattern of activity itself. But this would be implausible and would not solve the problem as there might be no more determinate way of characterizing the pattern other than by "enough instances of this act" or "not too many instances of this act." For more on this issue, see Sergio Tenenbaum, "Reconsidering Intentions," *Noûs* (forthcoming); and Tenenbaum and Raffman, "Vague Projects."

69. I am representing moral rules as demands that agents make on patients for illustrative purposes, but, of course, nothing hangs on this form of presentation.

70. I argue for the claim that the rationality of a pattern of activity might not supervene on the rationality of each momentary action that constitutes this activity in Sergio Tenenbaum, "Reconsidering Intentions," and "The Vice of Procrastination," in *The Thief of Time: Philosophical Essays on Procrastination*, ed. Chrisoula Andreou and Mark White (New York: Oxford University Press, 2010), 130–50; and Tenenbaum and Raffman, "Vague Projects."

that it is not acceptable in wartime to significantly endanger civilians, we'll have to side with some form of pacifism; after all, there is not much point in going to war with the intention of shooting at the enemy sporadically. On the other hand, the ends of declaring war, or even of volunteering or complying with a conscription order, are not the same as the ends of each particular battle I am engaged in; we are engaged in rather different actions when we do one or the other. It is thus fully compatible with the forms of deontology I am proposing that the prohibitions that apply to declaring and joining a war are significantly different from the prohibitions that apply to one's conduct in war.

VI. CONCLUSION

This is, of course, not a full defense of the plausibility of deontology in dealing with situations involving risk; in particular, those who find deontology implausible in nonrisk contexts are unlikely to embrace it in actions and decisions under risk. But I hope to have shown that, at least with regard to some forms of deontology, those who favor such approaches should not be moved by the lure of the multiplicative model. The natural way to extend deontology to risky situations is to understand risk as a morally relevant aspect of certain situations, rather than as a factor that is present in all situations in different degrees. Moreover, although differences in risk are morally relevant, there is no reason to think that the best way to understand this relevance is invariably through numerical probabilities. Risk is often morally relevant because it changes what it is my power to do, not because it puts different weights on the same things I would otherwise do.