Legal Theory, **17** (2011), 301–317.
© Cambridge University Press 2011 0361-6843/11 $15.00 + 00
doi:10.1017/S1352325211000176

DEONTOLOGY, THRESHOLDS, AND EFFICIENCY

Christopher T. Wonnell*
University of San Diego School of Law

This article explores four topics raised by Eyal Zamir and Barak Medina's treatment of constrained deontology. First, it examines whether mathematical threshold functions are the proper way to think about limits on deontology, given the discontinuities of our moral judgments and the desired phenomenology of rule-following. Second, it asks whether constrained deontology is appropriate for public as well as private decision-making, taking issue with the book's conclusion that deontological options are inapplicable to public decision-making, whereas deontological constraints are applicable. Third, it examines the issue of the relationship between deontology and efficiency, asking whether deontological constraints should yield in situations where everyone would expect to benefit from their suspension, either *ex ante* or *ex post*. Finally, the article concludes that constrained deontology is susceptible to political abuse because of the many degrees of freedom involved in identifying constrained actions and the point at which those constraints yield to consequentialist benefits.

Theoretical scholarship on the normative evaluation of legal institutions tends to resemble warring camps. On the one side are the consequentialists, especially the legal economists, who speak of Pareto and Kaldor-Hicks efficiency, cost-benefit analysis, and social-welfare functions.[1] On the other side are the deontologists, who refer to categorical obligations, agent-centered permissions, retributive and corrective justice.[2] Into this debate enters a recent book by Eyal Zamir and Barak Medina seeking to integrate the two systems of ideas into a single coherent perspective for evaluating legal rules.[3]

*I would like to thank Matt Adler, Larry Alexander, and Emily Sherwin for their helpful comments.

1. *See generally* Anthony Boardman, David Greenberg, Aidan Vining & David Weimer, Cost-Benefit Analysis: Concepts and Practice (1996) (discussing normative economic concepts).

2. *See generally* Shelly Kagan, Normative Ethics 70–78 (1998) (discussing deontological constraints and options).

3. Eyal Zamir & Barak Medina, Law, Economics, and Morality (2010).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

Zamir and Medina are "threshold deontologists" who believe that there are certain actions that are forbidden even if their consequences are on balance desirable, unless those consequences are extremely desirable.[4] In their view, deontologically constrained actions include the deliberate killing or torturing of innocent persons, lying or promise-breaking in commercial relations, discriminating on the basis of race or sex, intervening paternalistically in the lives of others, and censoring speech on the basis of its content.[5] The challenge of the position is to explain why consequentialist advantages in general are not sufficient to justify the constraint violation and then to identify and defend a threshold of consequences beyond which it would indeed be justified to violate the categorical constraint.

The book's principal purpose is not so much to answer the question of where to place the threshold as it is to identify the structure of such an answer. Zamir and Medina defend the idea of "threshold functions" that would add additional arguments to conventional social-welfare functions.[6] Violating a constraint, such as intentionally killing an innocent person, could be justified only if the net consequential benefits exceed a numerical threshold. For example, the threshold might consist of a constant C, so that, say, at least four lives would need to be saved on balance if an action involved intentional killing of innocents, if C were set equal to 3.99. The function might also include another constant, T, that would vary with the number of constraint violations, so that, say, the action of deliberately killing n people would be justified if the net benefits of that action (B) exceeded $C + Tn$. In principle, more complicated, nonlinear threshold functions could also be considered.[7]

This approach raises at least three issues. First, are numerical threshold functions the right way to think about deontological constraints? Second, are governmental actors properly susceptible to analysis in terms of deontological options and constraints? Third, how does deontological constraint interact with specifically economic norms based upon *ex ante* or *ex post* efficiency? This review is divided into a consideration of those three issues and then concludes with a few remarks about the potential for abuse of threshold deontological norms in practice.

4. *See* Larry Alexander, *Deontology at the Threshold*, 37 SAN DIEGO L. REV. 893 (2000) (discussing challenges facing theories of deontological obligations that can be trumped by good enough consequences).

5. *See* ZAMIR & MEDINA, *supra* note 3, at 127–347.

6. *See id.* at 79–104.

7. Indeed, it does seem likely that linearity is a severe straightjacket on intuitions. We might feel that we should kill ten innocents only when this would save the lives of ten thousand. But the ratio of 1000:1 might seem quite inappropriate if this required us to refrain from using an atomic bomb that would kill 10 million innocents but that would thereby save the entire human race of 7 billion people. The issues seem morally separate and not susceptible to any linear solution.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

I. ARE THRESHOLD FUNCTIONS THE RIGHT WAY TO CONCEPTUALIZE DEONTOLOGICAL CONSTRAINT?

Zamir and Medina attempt to make deontological constraint familiar to more traditional consequentialist scholars through the device of threshold functions. These begin with more conventional social-welfare functions but then add additional elements based upon the constraint. In the example given above, where lives are being both taken and saved by an action, one might say that the social goal should be to maximize SW = B - C - T*n*, where B is the benefit of net lives saved. C and T would be numbers chosen to match one's considered moral judgment about what it would take to justify the deliberate killing of *n* innocent people. Perhaps it would take five net lives to justify any killings, and then an additional three net lives saved per killing, so that C and T could be set just below 5 and 3.

There is something immediately odd about using functions of this sort to describe deontological constraint. Mathematical signs such as the equals sign and pluses and minuses make sense only if one is dealing with entities of the same kind and common units of measurement. One might think that deontology involves entities of two different kinds, actor-neutral consequences and actor-relative side constraints on action, and that adding the two would be like adding five wars and three aardvarks.[8] Now, the threshold function above does not actually suffer from this problem. C, T, and B are all measured in lives, so they can be coherently subtracted from each other. However, B involves real lives gained and lost; their names can be given if need be. If C and T introduce any additional lives, they are artificial. No one died who is not already named as part of B. Comparability and common units are introduced by fictional device; we are told to act *as if* there were additional people dying as a result of one's action. Then the actor can make an "agent-neutral" decision to maximize "consequences" straightaway, instead of making the agent-relative maximizing decision subject to constraints applicable to her, as one would have thought deontology involves.

This fictional device does some violence to the phenomenology of the deontological approach. In this representation of the problem, the killing and torturing of innocents are always regarded as live options in a maximizing game. They will be chosen less often than otherwise because of the fictional weight of the imaginary additional lives lost from taking those actions. But the agent is urged always to consider those possible means on a par with any other means for maximizing the function. This analytical approach may

8. The concepts of incommensurability and incomparability (distinct concepts but both required for adding) are addressed in RUTH CHANG, INCOMMENSURABILITY, INCOMPARABILITY, AND PRACTICAL REASON (1997).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

produce the right results, but its method of getting to those results seems in tension with the spirit of deontology.[9]

Moreover, it is not clear that the fictional device of imaginary consequences will indeed always replicate the actions one wants. The numbers C and T were of course gerrymandered precisely to replicate our considered moral judgment about the appropriateness of killings that involve net savings of lives. But will those fictions continue to serve if killings produce other consequences? Zamir and Medina themselves state that they believe the killing of innocents can be justified only to prevent other deaths and not other social harms.[10] This is not because Zamir and Medina believe that the number of deaths produced by a policy always has lexical priority *as a consequence* over other consequences such as reducing the suffering of the living. Rather, their moral intuitions tell them that intentional killing of the innocent is the type of wrong that can be justified only by (sufficient) net savings of lives and not by other types of consequences, even consequences that would otherwise be regarded as better than a net savings of lives.

Their moral intuition may be plausible, but it is in tension with their mathematical approach. Let us divide the social-welfare function into two components: the number of lives saved and the quality of life for those who remain. An unconstrained social-welfare function might instruct us to maximize something like $SW = L + Q$.[11] Now Zamir and Medina ask us to introduce fictional additional consequences and maximize a new "social-welfare" function, $SW2 = L + Q - C - Tn$. Unfortunately there is no value one can assign to C and T that would replicate Zamir and Medina's moral intuitions. Whatever the value of C and T, the function is going to tell us to be more tolerant of intentional killings when Q is higher, and if Q is higher enough, even if L itself is lower.

This problem can be fixed by introducing a new fictional consequence. Imagine another variable, Z, that is always exactly equal in absolute value to Q but opposite in sign. The variable Z comes into play only when a deontological constraint is violated, which we can mathematicize by introducing an indicator function, D, that is equal to 1 when the constraint is violated and 0 when it is not. Now one can produce a social-welfare function $SW3 = L + Q + DZ - C - Tn$. This magically has the effect of precisely counterbalancing the gain in quality of life with an equivalent utility loss only in cases where the constraint is violated. These are fictional antiutils; no one actually experiences them. And it is quite an extraordinary fiction.

9. *See* CHARLES FRIED, RIGHT AND WRONG 10 (1978) (arguing that wrongs such as promise violations might be justifiable by avoiding catastrophic consequences but that these should be seen as extreme situations in which normal categories of right and wrong do not apply).

10. *See* ZAMIR & MEDINA, *supra* note 3, at 86 (lexical priority of lives over pecuniary values).

11. Of course, in some way this function would itself face an issue of comparability. Perhaps the function would involve maximizing total as opposed to average utility, since the total would be affected both by the number of lives and by the average happiness of whatever lives there were. Total utility is in some ways not very attractive, however, and there might be better ways of converting these two consequential values into a common denominator.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

One is supposed to imagine that one's constrained action always happens to produce just enough disutility to offset whatever utility it provided to the living, so that the better its actual consequences for the living, the worse its imaginary consequences in precise proportion. But it produces the right result, at least until a different type of choice needs to be confronted.

In fact, Zamir and Medina propose very complicated rules for the consequences that can outweigh constraint violations. They argue that low probability consequences cannot justify constraint violations and that losses prevented can justify more constraint violations than can prospects for gain.[12] They also contend that small consequences to individuals should not justify constraint violations, no matter how many individuals they affect.[13] Again, Zamir and Medina do not argue that foregone benefits are always less important *as consequences* than costs inflicted nor that small harms to many people are always less significant as consequences than large harms to few people. Rather, the moral significance of these different types of consequences varies depending upon whether one is contemplating a deontologically constrained or an unconstrained action.

At some point, the world of unconstrained consequentialism and the world of constrained consequentialism just do not seem to resemble each other. Of course, even this can be fixed. We can take the ordinary unconstrained social-welfare function and introduce a new fictional consequence that completely removes it and starts over. If SW is the ordinary social-welfare function, we just need a variable $Q = X - SW$, where X is some completely different social-welfare function, and then we can tell the actor to maximize SW + QD. In effect, we would be telling the actor that when she is contemplating a constrained action ($D = 1$), she should throw the old welfare function, SW, out of the window and substitute the new one, X. I suppose this is "integrating" social welfare and deontology in form, but it remains a pretty thorough apartheid between the two worlds in substance.

There may indeed be a more fundamental problem with Zamir and Medina's mathematical approach. It may be impossible to represent the types of normative judgments that Zamir and Medina want to have in mathematical form. Economists have experience with these issues from consumer theory. People have all kinds of preferences; they may prefer apples over oranges, going to a church over going to a synagogue, or working on a law review article rather than preparing for class. Consumer theory says that on certain assumptions about our preferences, it is possible to reproduce the consumer's choices with a model in which the consumer is trying always to get the highest mathematical score on a single "utility" function.[14] The claim is not really that there is one unitary thing called "utility" to which all

12. ZAMIR & MEDINA, *supra* note 3, at 87–91.

13. *Id.* at 87. The issue is discussed in Samantha Brennan, *Thresholds for Rights*, 33 S. J. PHIL. 143, 160 (1995).

14. *See* ANDREU MAS-COLELL, MICHAEL D. WHINSTON & JERRY R. GREEN, MICROECONOMIC THEORY 46–47 (1995).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

preferences are ultimately reducible.[15] There may be many types of human objectives, but on certain assumptions about preference structures, a set of util numbers for every possible object of choice can be posited such that with those numbers, people act *as if* they were trying to maximize their mathematical total of those utils.

However, consumer theory says that this mathematical representation can work only if preferences obey certain rules. One of those rules is continuity. Suppose that a consumer is choosing among bundles of three commodities—apples, blueberries, and cherries. Suppose that she strictly prefers a bundle that has two apples and one blueberry over another bundle that has two blueberries and one apple, whether her bundle also includes a million cherries a year or fifty or one cherry every millionth year. However, if she has zero cherries, she reverses her preference and strictly prefers the bundle of two blueberries and one apple over the bundle of two apples and one blueberry, never having experienced a point of indifference between the two bundles as the number of cherries declines. These discontinuous preferences cannot be represented mathematically by a utility function assigning a fixed number of utils to each commodity bundle.[16]

To be sure, Zamir and Medina do not purport to be describing preferences. Nevertheless, they are trying to produce a kind of ordering that a morally sensitive person ought to affirm for any given bundle of social-welfare consequences and constraint violations by the agent.[17] The agent might be said to be trying to maximize a mathematical function called "moral value," and there would not necessarily have to be any entity in the world corresponding to that term for the mathematics to work. However, the mathematics will not work if the moral tradeoffs between social-welfare elements and deontological constraints do not exhibit continuity.

As it happens, Zamir and Medina's concrete moral judgments are full of discontinuities and lexical orderings. Suppose that as a matter of pure consequentialism we would find a policy that generated one extra death but X extra utils per living individual to be moral progress. However, if constraints are violated, Zamir and Medina would reverse the judgment. From all they say, that would remain true if the agent has to violate the constraint a million times a day or once every million days in order to produce these consequences. But if the agent does not violate the constraint, the consequentialist ordering is reinstated. Moreover, these are not tiny discontinuities. A world of fewer deaths but sufficiently many additional utils per living person could, under some circumstances, be far superior on

15. *See* JOHN BROOME, WEIGHING GOODS 66 (1991) ("By 'the utility of A' I mean the number assigned to A by a function that represents an ordering, and nothing else.").

16. A proof of this proposition appears in MAS-COLELL, WHINSTON & GREEN, *supra* note 14, at 46.

17. The idea that normative orderings of goods can be treated on an analogy to descriptive orderings of preferences is stated in BROOME, *supra* note 15, at 165 (1991) (principle of personal good instead of Pareto-based theory of preferences).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

a consequentialist calculation, so that the discontinuity of judgment as the constraint violations go from a small positive number to zero could be quite serious.

I am not necessarily suggesting that Zamir and Medina's enterprise of giving mathematical form to deontological constraint is hopeless. If the discontinuities are removed, utility theory gives reason to believe that some kind of mathematical representation should be possible, though it may be frightfully complex. On the other hand, if moral judgments in the deontological area ought in fact to exhibit discontinuities and lexical priorities, then presumably that is a reason to abandon the mathematical approach of constraint functions. The whole enterprise of extra deaths and extra antiutils and moral value points is fictional to begin with and makes sense only as a way of representing through an orderly set of numbers the concrete normative judgments we want to affirm. If the mathematics would force us to alter our moral judgments just to preserve the fictional notion that we are single-mindedly maximizing one goal, so much the worse for the mathematics.

II. IS DEONTOLOGY SUITED TO PUBLIC AS OPPOSED TO PRIVATE MORALITY?

In an important chapter of the book, Zamir and Medina address the question of whether deontology might be an ethic of private morality only and therefore unsuited to governmental affairs.[18] In particular, the authors take up a challenge laid down by Cass Sunstein and Adrian Vermeule, who take the position that deontology is indeed inappropriate for governmental actors and use that rationale as a way of defending capital punishment against certain deontological objections.[19]

Deontology has two facets: agent-relative permissions not to maximize the good, and agent-relative prohibitions on certain types of actions that would maximize the good. As an example of the former, an agent may have the option to pursue her own projects rather than devote her time to alleviating the distress of another, even if the agent is not at liberty affirmatively to cause equivalent distress in another. As an example of the latter, an individual might lack the right to kill someone in retaliation for murder committed against one's loved ones, even if such retaliation would be good for the community as a whole through improved safety.

Sunstein and Vermeule take the position that these normative intuitions of morally optional conduct and moral constraints on intentional harms with utilitarian benefits are much more plausible for private than for public conduct. An individual is not responsible for alleviating the distress of

18. ZAMIR & MEDINA, *supra* note 3, at 57–78.

19. Cass R. Sunstein & Adrian Vermeule, *Is Capital Punishment Morally Required? Acts, Omissions, and Life-Life Tradeoffs*, 58 STAN. L. REV. 703 (2005).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

random strangers, but the state does take on such a responsibility for the citizenry at large. A private individual may not have the right or the duty to engage in vigilante justice, but the state does have the right and the duty to control crime by taking retributive actions on people's behalf. Sunstein and Vermeule argue that if executions actually do reduce the murder rate appreciably, as some studies suggest, then the government bears full responsibility not only for the deaths it brings about by executing but also for the deaths it fails to prevent by not executing.

Zamir and Medina take the position that deontological options do not apply to the government but that deontological constraints do apply.[20] Deontological options are grounded in liberty values that do not apply to the state when it fails to be of service to people in need. But it is important for the state not to engage in wrongs, such as torture, that would also be wrong for private individuals. Indeed, the systematic and organized nature of such violence might make it more important for the government than for private citizens to be kept within deontological constraints on particular types of egregious actions.[21]

I want to argue that deontological considerations, properly understood, apply equally to the state and to private behavior. Indeed, I do not see how it could be otherwise, since the state is not an otherworldly being but merely a particular instrumentality through which some individuals are acting on others. However, the two contexts result in different concretizations of the same abstract deontological principles for private and public behavior, and it is important not to confuse the concrete rights and duties of private citizens with the abstract principles on which they are based.

Let us look again at individuals who exercise their alleged moral option not to be of affirmative assistance to their fellow citizens in need. How are we to think about this case? One possibility is that people have a genuine right not to be affirmatively useful to others in certain circumstances because their liberty is at stake. Another is that people have an abstract obligation to do their part toward the needy but that this obligation cannot be effectively concretized without a coordinated scheme enforcing each individual's separate obligation. In other words, a norm that looks like a deontological permission not to be charitable is really a difficult-to-enforce duty to participate equitably in a scheme of poor relief that the individual is not in a position to create on her own. These two conceptions of deontological options are not mutually exclusive. It is possible, indeed it seems probable, that we have duties to be of some assistance to the needy but have a genuine moral option not to be of the maximum possible assistance.

20. *See* ZAMIR & MEDINA, *supra* note 3, at 60–61 ("It is *logically* possible and *normatively* plausible that governments are morally required to actively promote the well-being of all of their people and at the same time are subject to the moral duty to refrain from actively harming them.") (emphasis in original).

21. *See* Christopher Kutz, *Torture, Necessity, and Existential Politics*, 95 CAL. L. REV. 235, 256 (2007) (moderate deontology at governmental level creates risk of overuse of torture).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

The government can supply a coordinated scheme of poor relief and enforce more or less fair contributions from each person. Nevertheless, the introduction of the government does not change the fundamental morality of the situation. If individuals have a genuine moral option whether to be of assistance, then they have a genuine moral option whether to tax themselves to provide the coordinating scheme they might prefer, and the government cannot justly cater to the consequential claims of beneficiaries by forcing taxpayers into a system they do not willingly embrace.[22] By contrast, if individuals have an abstract obligation to be of assistance that becomes feasible only when concrete governmental coordination and enforcement devices are available, then they have an obligation to set up and support a well-designed redistributive institution and can be appropriately made to participate in such a system by a government responsive to the just demands of the beneficiaries.

The issue is not government's options or the personal options of government officials but citizens' moral options to provide or not provide certain types of poor relief through the government, along with their lack of moral options to resist justly other types of government poor relief. As an ideal, government is a faithful agent of citizens who are acting wisely and justly in the claims they make upon the government. There are difficult second-best questions about what government officials should do when their constituents make demands upon them that the officials believe are unjust or are based upon flawed empirical premises. The aspiration, however, is to have a government willing and pressured only to execute the optional wishes of people in cases where they have those options and the demands of people upon others where justice imposes affirmative obligations.[23]

A test of moral intuitions about deontological options is how one would handle cases where exigent circumstances mean that the governmental safety-net apparatus is unavailable—perhaps a drowning swimmer in a remote location. The question then is whether individuals who happen to be in a position to help have a personal duty to rescue. It would be consistent with the approach taken to poor relief in general to imagine that individuals have a moral obligation in such cases to make an easy rescue but have a genuine option not to incur large costs to their own projects and lives in carrying out such a rescue. This of course would leave room for the supererogatory rescue that would go beyond the call of duty and actually maximize good consequences impersonally considered.

22. *See* MILTON FRIEDMAN, CAPITALISM AND FREEDOM 192–194 (1962) (defending the concept of the top 90% imposing a tax on themselves for the benefit of the bottom 10%). This seems to imply that the rich are doing something that no one else has the right to make them do, which may be true for some categories of relief and not for others.

23. Poor relief that is optional is somewhat analogous to contributions for public goods. A person may have a moral obligation to contribute his fair share toward a public good such as national defense, which is hard to enforce against individuals in the absence of a coordinating state apparatus.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

With respect to deontological constraints, I would say once again that government constraints are fundamentally derived from moral constraints attached to human action generally, though the same abstract principles will be concretized differently to apply to facts applicable to government. In particular, legal systems are often quite heavily tilted against the purposeful use of private violence, even in circumstances of retaliation for wrongs where violence might significantly lower the crime rate.[24] Is this constraint maintained because human beings have no moral right to use violence in retaliation for wrongs committed against themselves, their loved ones, or strangers? Or is it rather that citizens have a right (and perhaps in some cases even a duty) of retaliation but that a coordinating mechanism such as the state provides in most cases an orderly and controlled way to carry out that violence?

There is no doubt that private retaliatory violence undertaken upon unproven suspicion of previous wrongs can be a recipe for unhealthy cycles of vigilantism and revenge. Still, if at some fundamental level the state operates under the same constraints as private citizens, there must be circumstances where individuals cannot realistically count on the state and therefore have rights to engage in suitably proportioned force in response to previous wrongs. Parental rights of discipline may be an example of legally tolerated private retributive violence. And almost surely the law in practice is likely to be more tolerant of such behavior than the law on the books. Athletes are often affirmatively expected to "protect" their teammates by retaliating on their behalf, whether through high sticks in hockey or inside pitches in baseball. It is hard to believe that a schoolchild who is the victim of endless bullying and finally lashes out against his tormentors in a proportionate way would receive no de facto mitigation of his own act. Lawyers may overestimate the practicality of judicial remedies in many settings, and it is likely that norms of personal conduct and property rights are often enforced in fact by individuals who raise their fists with the tolerance of the community.

Zamir and Medina are especially concerned to avoid any conclusion that government officials are somehow immune from constraints against the worst types of behavior, such as torturing innocent relatives of a suspected terrorist. In this case there appears to be a genuine moral constraint against a particular type of human action and not merely a constraint on private action that facilitates orderly and controlled use of force by individuals acting as state agents. If there is a genuine moral constraint against some type of action, Zamir and Medina are surely correct that the state has no immunity from that moral constraint.

24. The distinction between retaliation and preemptive self-defense is stressed in Larry Alexander, *A Unified Excuse of Preemptive Self-Protection*, 74 NOTRE DAME L. REV. 1475, 1476 (1999).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

Overall, there is a question of whether the act/omission distinction, which seems so fundamental to deontology, is applicable to public conduct.[25] The state can act unjustly, but is it not just as culpable for failing to act as it is by tolerating violence against some of its citizens when it could punish sufficiently to deter that violence? Again, however, the dichotomy between private citizens and the state is overdrawn. Private citizens surely have an obligation to contribute affirmatively their fair share toward a law-enforcement apparatus that extends its protection to all citizens. And when private citizens cannot act affirmatively, neither can the state.[26] Neither private citizens nor state officials are authorized to deter crime by punishing people whom they know to be innocent, even if more innocent lives are thereby spared than taken.

To summarize, Zamir and Medina are correct to say that genuine moral constraints that apply to private citizens apply equally to state officials. They are incorrect in claiming that state officials have a duty to act when the underlying citizenry has a genuine moral option not to act and has indicated its desire to stand on its option. The appearance that the state is morally different from private actors stems from the fact that many of the practical options and constraints of private citizens are not genuine moral options and constraints but are instead duties to act through coordinated and orderly processes.

III. WHAT IS THE RELATIONSHIP BETWEEN ECONOMIC EFFICIENCY AND DEONTOLOGY?

An important aspect of the relationship between welfare economics and deontology is the challenge posed by Louis Kaplow and Steven Shavell.[27] They argue that deontological obligations (including such nonconsequentialist norms as corrective and retributive justice) will in some cases require people to act in a way that makes everyone worse off and thus to violate the Pareto principle. It is worth quoting the response of Zamir and Medina to this argument at some length:

> There are several persuasive responses to this claim. To begin with, deontological theories do not *necessarily* conflict with the Pareto principle. A deontological theory may plausibly qualify any constraint or option such that it would not apply whenever everybody would be better off without this constraint or option. More fundamentally, the conflict between deontology and the Pareto

25. *See* Sunstein and Vermeule, *supra* note 19, at 721 ("When Homeowner B sues Factory A over air pollution, a decision not to rule for Homeowner B is not a form of inaction: it is the allocation to Factory A of a property right to pollute. In such cases, an apparent government omission is an action simply because it is an allocation of legal rights.").

26. *See* ROBERT NOZICK, ANARCHY, STATE, AND UTOPIA 6 (1974) ("What persons may and may not do to one another limits what they may do through the apparatus of a state, or do to establish such an apparatus.").

27. LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE 52–58 (2002).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

> principle is tautological: it merely restates the conflict between deontology and consequentialism. It may therefore be read as establishing the weakness of the Pareto principle. Finally, in the context of deontological constraints, Kaplow and Shavell's argument holds only behind an imaginary contractarian veil of ignorance, for in real life, it is the *removal* of a constraint against harming people that necessarily violates the Pareto principle.[28]

Let us start with the last point: that there is no possible real-world conflict between the Pareto principle and deontological constraints because those constraints are protections against being harmed, which necessarily means that the Pareto principle has already been violated. This view appears problematic on a number of levels. The first point is that iterated repetitions of constraint violations can benefit everyone even if each constraint violation, considered separately, imposes harm. For example, Zamir and Medina argue that there are deontological obligations not to deceive transacting partners or not to breach promises.[29] If the efficiency arguments are to be believed, however, it is sometimes at least Kaldor-Hicks efficient to allow nondisclosure of deliberately acquired information[30] and breaches of contract when unforeseen material changes in circumstances have occurred.[31]

That does not necessarily mean that there are no losers from the transactions involving constraint violations considered individually. But we are participants in large numbers of transactions over our lifetimes, especially considering our indirect stakes as consumers, employees, or investors in the institutions involved in litigating these cases. If we find ourselves on different sides of the issue in a more or less random manner,[32] and the gains of the winners outweigh the losses of the losers, then the rule can be Pareto efficient despite the harm from each constraint violation.[33] Even if the consequences are not completely random, so that there are still net losers, the incidence of costs can be altered by redistributive measures to make the overall policy Pareto efficient or at least not harmful on balance to the people whose deontological constraint was violated.

A related point is that deontological constraints can harm their putative beneficiaries by virtue of long-term effects. Zamir and Medina argue that there is a deontological constraint against discrimination on the basis of

28. ZAMIR & MEDINA, *supra* note 3, at 51 (emphasis in original).
29. *Id.* at 257–311.
30. *See* Anthony T. Kronman, *Mistake, Disclosure, Information, and the Law of Contracts*, 7 J. LEGAL STUD. 1, 9–18 (1978) (important to encourage deliberate searches for productive information).
31. *See* RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW 119 (7th ed. 2007) ("In many cases it is uneconomical to induce completion of performance of a contract after it has been broken.").
32. *See* A. Mitchell Polinsky, *Probabilistic Compensation Criteria*, 86 Q.J. ECON. 408, 409 (1972) (bundles of changes with randomness of distribution).
33. *See* JOHN D. GRAHAM, SAVING LIVES THROUGH ADMINISTRATIVE LAW AND ECONOMICS 395, 414 (2008) (single vs. repeated applications of Kaldor-Hicks).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

race, sex, and perhaps other characteristics such as personal handicaps.[34] These constraints are said to be based upon respect for the autonomy or choice set of the victims and are said to include not just intentional discrimination but also actions taken for otherwise unobjectionable reasons that have disparate impact.[35] Now, a serious school of thought suggests that highly aggressive antidiscrimination norms of this sort can impose long-term harms on stated beneficiaries. Firms are aware that they are more likely to be sued for their firing than for their hiring decisions, so the safe way for the firms to minimize antidiscrimination suits is not to hire the people who might bring suits.[36] Moreover, the inefficiencies from forcing firms to ignore otherwise useful practices with disparate impact mean that companies that are required to incur these costs will suffer competitively against firms located in nations or regions of the country with fewer minorities.[37] It is therefore important for Zamir and Medina to address the question of whether their deontological constraint against discrimination should hold up if the expected long-term consequences of the policy turn out to be negative for minorities.

The conflict between Pareto principles and deontology is still more likely if one embraces an *ex ante* approach to the Pareto norm.[38] *Ex ante* efficiency would say that a change is desirable if everyone would rationally expect to benefit from it (or at least that no one would rationally expect to lose and some would rationally expect to gain), even if everyone knew that some people would wind up losing *ex post*.[39] For example, allowing private insurance might be defended as promoting mutually beneficial transactions *ex ante* for a risk-averse insured and a risk-neutral insurer, even though both parties understand that the feared loss will either occur or not occur, with

34. Zamir & Medina, *supra* note 3, at 225–256.

35. Basing deontological constraint on autonomy or choice-set values is problematic in general. The right to choose one's contractual associates might be regarded as itself a species of autonomy. This applies as well to other examples from the book. Zamir and Medina defend free speech, even when it causes consequential harm, because of the autonomy or choice set of the speakers, but the harm itself, such as reputational harm from defamation, contracts the choice set of the victims. Perhaps one could say that this simply acknowledges that deontological constraint can apply even if violating the constraint would reduce the number of constraint violations, but it still begs the question of which is the constrained act.

36. *See* Posner, *supra* note 31, at 357.

37. This is especially likely regarding handicap discrimination, where firms are required to expend real resources in make reasonable accommodations. *See, e.g.*, David C. Stapleton & Richard V. Burkhauser, The Decline in Employment of People with Disabilities: A Policy Puzzle (2003).

38. *See* Christopher T. Wonnell, *Efficiency and Conservatism*, 80 Neb. L. Rev. 644, 666–682 (2001) (defending *ex ante* efficiency ideal against objections).

39. The concept is discussed and criticized in Matthew D. Adler, *The Puzzle of "Ex Ante Efficiency": Does Rational Approvability Have Moral Weight?*, 151 U. Pa. L. Rev. 1255 (2003). Broadly speaking, Adler criticizes *ex ante* efficiency because it is not an actual choice that embodies consent-based values nor a decision that would actually be in the person's best interest. The intuition of the *ex ante* approach, however, is that it makes sense to apply moral criteria to decisions that make the best available use of imperfect information about probabilities.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

the result that one party or the other will wind up having lost from the transaction as a whole.

There are serious possibilities for a conflict between deontological obligations and *ex ante* efficiency. Consider a well-known variation on the example of the trolley car careening out of control that is about to kill three pedestrians, P1, P2, and P3.[40] An agent A can save the three pedestrians only by pushing an innocent bystander B into the way of the trolley, which will kill the bystander. Deontology says that A is not required—and indeed is not permitted—to take this action, despite the net savings of two lives that would result; B is an autonomous person, entitled to be respected as an end in himself and not to be conscripted as a mere means or resource for the benefit of others.[41]

From an *ex post* perspective, where we know who is playing each of the identified roles, it does appear that B has been taken and used to benefit others. But suppose that the question had been raised earlier, in legislative debates about types of behavior that should be deemed acceptable in rescue situations. Suppose further that the chance that one would wind up in the position of A, B, P1, P2, or P3 was roughly equal, and that the deliberating citizens made a rational, self-interested calculation that they would be willing to forego their deontological right if they ended up in B's position in order to receive protection if they end up in P1, P2 or P3's position.[42] In short, the rule permitting the constraint-violating push maximized the expected utility of each party at the time the legislation was passed.

Perhaps hypotheticals of this type motivate Zamir and Medina's phrase about the "imaginary contractarian veil of ignorance."[43] But there is nothing particularly imaginary about the idea that parties might want to forego their rights against being used as a means in certain circumstances.[44] After all, deontological constraint is important only in cases where the action in question produces more good consequences than bad ones. If those consequences are unpredictable in incidence *ex ante*, it is quite likely that each person would expect to benefit from removing that constraint. Again, it becomes important to know how a system of social evaluation should handle such a conflict.

40. The original trolley hypothetical involved switching tracks so as to kill a smaller number of people, a hypothetical that tends to elicit different attitudes from deliberately tossing an innocent person in the way of the trolley. *See* Judith Jarvis Thomson, *The Trolley Problem*, 94 YALE L.J. 1395 (1985) (presenting the trolley hypothetical to illustrate deontology).

41. The original Kantian sentiment advises: "Act in such a way that you treat humanity, whether in your own person or in the person of any other, never merely as a means to an end, but always at the same time as an end." IMMANUEL KANT, GROUNDWORK OF THE METAPHYSIC OF MORALS 30 (James W. Ellington trans., 3d ed. 1993) (1785).

42. If one thinks of deontology in terms of the agent, A, rather than the victim, B, then the question is whether A should have the power to waive his own deontological objection to participating in the pushing scheme by advanced consent.

43. *See* note 28 *supra*, and accompanying text.

44. *See* JAMES M. BUCHANAN & GORDON TULLOCK, THE CALCULUS OF CONSENT 63–84 (1962) (describing a constitutional stage where uncertainty is sufficiently great that individuals can rationally support public-regarding policies).

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

The remainder of the quotation cited above suggests that Zamir and Medina are willing to imagine either that the deontological constraint is lifted when it conflicts with Pareto or that it continues and trumps over Pareto. This agnosticism is not very satisfying in a book purporting to tell us how to integrate deontology and welfare economics into a unitary system of social evaluation. So let us think through the different possibilities.

One idea would be to say that deontology is a trump card over Pareto considerations. It is wrong to deceive, even if certain categories of deception generate net benefits for everyone when applied over long periods of time. But why? Deception ordinarily violates Kantian norms against the use of another autonomous person as a mere means. But is that true for a category of deception that one is considering embracing only because it is better for everyone over time? It is not obvious to me that some particular person is still being used as a means *only* when the category of behavior is approved on the express condition that *each person considered separately* derives net benefits from allowing that behavior.

Perhaps one could say that only actual consent, and not mere benefit, would suffice for waiving deontological duties. There is a great deal to be said for securing actual consent when it is possible. It provides epistemic confidence in one's judgment that the particular type of deception or other constraint violation really is in each person's interest.[45] It also represents an act of the person's own autonomous will specifically inviting a particular conduct that might otherwise be thought to be a violation of personal space. But what if transaction costs make the expression of actual consent impossible? Would we want a rule that surgeons cannot operate on a person who has fallen unconscious because they cannot secure his actual consent and tampering with a person's organs without consent (and probably earning a fee in the process) is a using and a constraint violation?

Another option would be to say that *ex post* Pareto efficiency has the power to trump deontological constraint but that *ex ante* efficiency does not. The emergency medical care hypothetical could be changed to accommodate this concern. Imagine that the person is unconscious and that the medical community would say that a particular surgical treatment is indicated because of its large potential benefits but that the treatment carries small risks of making the patient even worse off. As luck would have it, the patient was made worse by the treatment, though not because of physician negligence. I suppose that one could say that the doctor has now committed a wrong because she violated a constraint and did not in fact leave the patient better off, but to me it seems more reasonable to say that the fact that the surgery was indicated from an *ex ante* point of view combines with the high transaction costs to legitimate the intervention.

45. This point is of considerable practical importance. One would worry greatly about empowering government officials with the ability to violate deontological constraints based upon theories that people somehow would have waived them at some earlier stage. The idea of "inalienable" rights is probably in many cases a practical expression of this fear.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

Of course, one could simply bite the bullet and say that certain types of actions are deontologically prohibited even if everyone would benefit *ex ante* or *ex post* from allowing them. An odd feature of this view, at least as Zamir and Medina present it, is that their deontologically constrained actions seem to be chosen precisely because of the harm that they inflict on designated individuals.[46] Their list does not include any arguably victimless crimes such as sodomy or blasphemy or dishonesty to self. The paradox is to explain how one could put actions on the list because of their intended harmful consequences for the victims and not want to take them off the list if it becomes clear to the actor and victim and part of their intent that the putative victim actually has rational reason to want this category of action to be allowed.

IV. CONCLUSION: A WORRY ABOUT DEONTOLOGY IN PRACTICE

This review focuses primarily on difficulties of threshold deontology at the level of theory, that is, whether it can be accurately expressed in mathematical terms, whether it applies equally to the state and to private actors, and whether its strictures yield to efficiency considerations. However, one of the principal concerns about threshold deontology as a set of workable philosophical norms is the potential for abuse of its concepts in practice.

Zamir and Medina rather confidently assert various policy conclusions that they believe follow from deontology, such as freedom of speech and antidiscrimination law. But it is not evident precisely how these conclusions are derived. The authors place significant stress on individual autonomy, but it can certainly be argued that many types of speech, such as libel, hate speech, and subversive advocacy, are threats to the autonomy or free choices of their victims and themselves should be deontologically constrained. With regard to discrimination, it seems equally arguable, at least for people of a certain political persuasion, that the ability to choose one's business associates is part of human autonomy.[47]

Consequentialist ideals such as utilitarianism are of course subject to abuse as well. Since utility is hard to measure, the analyst has a fair amount of freedom to represent her policy views as those that are consequentially best. Still, utilitarianism at least has a pull in the direction of counting

46. *See* ZAMIR & MEDINA, *supra* note 3, at 42 ("The central deontological constraint is against harming other people.").

47. This is not merely a matter of the action's being constrained even in the presence of autonomy benefits nor of reducing the number of constraint violations. The question is: Which action is constrained in the first instance? Is the speaker herself deontologically constrained from speaking in a manner that imposes harm? Is the government deontologically constrained from interfering with private choices of contractual associates? Perhaps there is a right answer to these questions, but the present issue is how easily the norm can be abused to rationalize policy conclusions reached on other grounds.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press

all preferences more or less equally, which gives less room for the analyst's preferences to exert overwhelming trumping power over those of her fellow citizens. Deontology does not seem to have a similar mechanism for disciplining the conclusions of the policy analyst as to when an action abridges autonomy or dignity sufficiently to invoke a deontological protection. Similarly, threshold deontology leaves a great deal of room for the analyst's preemptive conclusion that the threshold of consequences has or has not been exceeded in the particular case. A particular worry is that this power to determine when constraints have been overcome will tempt commentators into insulating policies they believe are consequentially justified from legitimate consequentialist challenges.

Perhaps one benefit of Zamir and Medina's mathematical approach to threshold deontology would be to introduce some regularity to this process and limit the potential for arbitrary assertions by the analyst that her preferred policy is immune to consequentialist challenges. If a particular quantitative negative weight were chosen for intentional abridgements of autonomy, for example, that weight would then have relevance to many policy judgments other than the specific one at issue, and the analyst would have to be careful about gerrymandering the weight to reach a treasured policy conclusion for the specific case at hand. It remains to be seen whether ideals such as autonomy or individual dignity are capable of being assigned consistent and plausible numerical weights that can keep threshold deontology from becoming a way of rationalizing rather than of reaching policy conclusions.

https://doi.org/10.1017/S1352325211000176 Published online by Cambridge University Press