

The Inefficacy Objection to Deontology: What it is, Why it is Important, and How to Respond to it*

Mark Bryant Budolfson
University of Vermont

Version 2.6

ABSTRACT

What you do as a single individual often seems to make no difference, especially in a large market based society. At first glance, this may seem to raise a challenge only for traditional consequentialist theories. However, I argue that inefficacy also raises important problems for deontological theories, and ultimately shows that existing deontological theories are unable to offer a plausible account of what individuals are required to do in a wide range of collective action situations that are common in modern society. In light of this, I identify a previously unnoted kind of deontological reason that can explain the ethics of not making a difference, where this new kind of reason depends on *the degree to which harm is essential to an act*. I argue that an appeal to such a notion is well motivated, firmly grounded in common sense, exists in the law and is necessary there as well. It has the potential to help with other classic challenges for deontological theory as well.

KEYWORDS

deontology, normative ethics, collective action, inefficacy, difference making, complicity, Kantian ethics, harm, consumer ethics, animal ethics, supply chain ethics

* Thanks to Christian Barry, Heather Berginc, Brian Berkey, Tom Blackson, Cheshire Calhoun, Eamonn Callan, Richard Yetter Chappell, Stew Cohen, Christian Coons, Terence Cuneo, John Devlin, Tyler Doggett, Jeff Downard, David Faraci, Ada Fee, Brian Fiala, Bob Goodin, Chris Griffin, Liz Harman, LP Hodgson, Ryan Jenkins, Alex King, Victor Kumar, Melissa Lane, RJ Leland, Eden Lin, Joel MacClellan, Sarah McGrath, Brad McHose, Tristram McPherson, Eliot Michaelson, Julia Nefsky, Alastair Norcross, Howard Nye, Govind Persad, Ángel Pinillos, David Plunkett, Joe Rachiele, Rob Reich, Ryan Robinson, Gideon Rosen, George Rudebusch, Carolina Sartorio, Debra Satz, Dave Schmidtz, Dan Shaha, Sam Shpall, Daniel Silvermint, Peter Singer, Michael Smith, Patrick Taylor Smith, John Thrasher, Ian Vandeventer, Chad Van Schoelandt, Alan Wertheimer, Jane Willenbring, Jack Woods, two anonymous referees, and audiences at the University of Arizona Center for the Philosophy of Freedom, Bowling Green, Stanford Center for Ethics in Society, University of Vermont, and the Colorado State Animal Ethics Conference for helpful discussions.

I. INTRODUCTION: THE INEFFICACY OBJECTION

In a large market-based society the actions of a single individual often seem to make no difference. This is especially true in connection with questions that are of core interest to ethics and social philosophy about what individuals have reason to do in collective action situations – for example, whether an individual is required to vote, whether it is wrong to consume products that are produced in morally objectionable ways, and many other examples. In these cases, it matters greatly what we together do, but yet the choices of a single individual can seem to make no difference to the welfare of others and also seems to involve no violation of rights. In light of this, it is worth considering the merits of the ‘inefficacy objection’ that in such cases, ethical theories cannot explain why individuals still sometimes have strong reasons to ‘cooperate’, even if doing so makes no positive difference.

For an example that is much-discussed in the recent literature, consider an individual’s decision whether to consume a good that is produced in a way that involves serious rights violations or imposition of unjustifiable suffering. According to utilitarians like Peter Singer, consuming such a product is wrong because it has unacceptable (expected) consequences on balance for welfare. However, even if we agree (as we often should) with Singer’s idea that the magnitude of suffering on the production side of the equation greatly outweigh the comparatively unimportant pleasures on the consumption side, the inefficacy objection is that his conclusion about the welfare effects of consumption does not follow and, upon careful reflection, turns out to be false. That is because, according to the key premise of the inefficacy objection, whether an individual consumes such products cannot really be expected to have any effect on amount of suffering or the violation of rights that lies behind them; at the same time,

an individual's decision to consume such products does have a positive effect for that individual. As a result, the inefficacy objection is that Singer's true premises, together with the actual empirical facts about the workings of the marketplace, entail that the expected effect of an individual's decision to forego ethically objectionable products is actually *worse* than the expected effect of consuming them, in contrast to what Singer's reasoning assumes.

In this paper, I set aside the question of whether the inefficacy objection is correct about the expected consequences of a single individual's choices, as that issue has been debated at length in a number of recent papers.¹ Instead, I focus on important questions for deontological theories that arise on the assumption that the premise of inefficacy is correct. I explain, first, why inefficacy is then an important objection not only to consequentialist theories, but also to deontological theories, by arguing that existing deontological theories are then unable to offer a plausible account of what individuals are required to do in a wide range of collective action situations. In light of this, I argue that we should then endorse a new source of deontological reasons, which is then needed to explain the cases at issue. This new kind of reason depends on *the degree to which harm is essential to an act*, which I analyze at a structural level and explain in more detail below. I note that analogous appeals to *degree of essentiality of harm* exist in the law and appear necessary there as well, and fit well with common sense.

II. WHY THE INEFFICACY OBJECTION RAISES A DIFFICULT PROBLEM FOR NORMATIVE THEORIES, AND WHY FAMILIAR APPEALS TO CONSTRAINTS AGAINST HARMING OTHERS AND THE LIKE DO NOT PROVIDE AN ADEQUATE SOLUTION

¹ For recent arguments that the inefficacy objection cannot be disarmed in the way many consequentialists suggest, see Budolfson 2018 and Nefsky 2011. These papers respond to classic arguments that consequentialism can successfully reply to the inefficacy objection, such as those offered in Singer 1980, Parfit 1984, and, following Singer and Parfit, Norcross 2004 and Kagan 2011.

Perhaps the most common reply to the inefficacy objection is to dismiss it on the grounds that it merely rehearses the familiar paradox of voting, which asserts that individuals do not have good reason to vote in elections because there is virtually no chance that a single individual's vote will matter. However, this is an inadequate reply for several reasons. To see why, note that insofar as individuals have reasons to vote in elections, there is broad agreement that those reasons arise from one or more of the following considerations:

- the expected welfare consequences of a single vote, given e.g. the probability that an individual's vote will trigger a *dramatic threshold effect*,²
- the fact that voters have a *personal preference* to vote,³
- the fact that voters collectively *cause* the outcome of the election in a particularly ethically important way,⁴
- the fact that voters have other *non-welfare-based* reasons to vote.⁵

But if we accept the premise of the inefficacy objection that a single act of 'cooperation' cannot be expected to make a positive difference, then none of these considerations can explain the

² For example, see Barry 1978, pg. 39, and Parfit 1984, pp. 72-74; compare also Lawford Smith 2015. For a classic objection to Barry and Parfit, and an argument that the inefficacy objection is correct in the case of voting, see Brennan and Lomasky 1993, chapter 4.

³ For example, see Mueller 2003, pg. 306.

⁴ For example, see Goldman 1999, and Moller 2005.

⁵ For a summary of such ideas, see Brennan 2012. See also Barry and Øverland 2016, Driver 2016, Nefsky 2018.

range of cases at issue in a straightforward way, even from a deontological point of view. The point of this section is to provide an initial argument for this claim – that these considerations and existing deontological views cannot explain these cases in any straightforward way – which will then motivate consideration of a number of other possible deontological explanations, which I will argue also turn out not to work, which will then motivate my positive account further below of the novel way in which I propose that deontologists should explain these cases.

To begin, note that by hypothesis theorists cannot appeal to dramatic threshold effects and expected consequences more generally, because we are assuming that the key premise of the inefficacy objection is correct, which simply claims that such an explanation will not work. Furthermore, theorists also cannot appeal to personal preferences, because many individuals do not have a strong enough personal preference, e.g., not to consume objectionable products and, more generally, do not have a preference to engage in ‘cooperative’ acts in the cases at issue here when it is costly to themselves to do so. In addition, while an appeal to a personal preference for voting (or expression via voting) is common and appropriate in economic discussions of voting where the issue is the *descriptive* challenge of explaining how voting could be *rational*, the issue here, in contrast, is *normative*, and so an appeal to personal preference or expressive value seems beside the point.⁶ The issue here is that supposing individuals do not happen to have personal preferences to avoid objectionable products, etc., what *ethical reasons* do they nonetheless have, and what exactly is their source?

In answering this question, a further challenge for a deontological explanation is that unlike cases where others are cooperating to successfully provide a public good where one might be said to owe it to the mass of cooperators to contribute oneself, the cases at issue here are typically marked by a general background of noncooperation and the further lack of successful

⁶ See Mueller 2003. Descriptive accounts of voting are similarly not on point for the *normative* question of what ethical reasons individuals have to vote.

provision of even basic regulation and other related public goods. So, the solution to the issues here typically requires new political action that is not even nearly being provided, and that is of a very different kind than merely increasing the number of individuals ‘voting with their dollars’. For these reasons, in the cases at issue here individuals do not seem to have significant reasons that follow from familiar anti-freeriding principles, since in the cases here the essence of the problem is that there is *not* successful pro-social cooperation, and so no successful cooperation they could be claimed to be freeriding on.

In addition, individuals in the collective action situations at issue – such as consumers of objectionable products – typically do not *cause harm* in anything like the way that is central to traditional deontological prohibitions on causing harm. For example, it is those who *produce* objectionable goods (rather than *consumers*) who are harming others in canonical ways that are familiar from deontological constraints against assault and the like. Perhaps more to the point, individual consumers are also not a partial cause of harm in the same ethically important way that individual bad voters are a partial cause of harm if they vote into power a very harmful official in an election. That is because there is an ethically relevant difference in the kind of causal connection at issue between voting cases and the collective action cases at issue here. To illustrate the difference, consider the causal connection between individual consumption of factory-farmed meat and the harms that lie behind it. A relevant fact is that in Australia, New Zealand, and many other large nations consumers have the same kind of individual meat consumption behavior as in the United States, but in the former nations there is not the same kind of harm that lies behind those decisions, because farm animals are generally treated humanely, at least in comparison to the sort of inhumane treatment that is common in the United States. The explanation of this difference despite nearly identical consumer behavior is that the *primary cause* of the horrible mistreatment of animals in the United States is not consumer behavior on the ‘demand side’, but is instead the decisions of producers on the ‘supply

side', as well as the perverse incentives created by governmental policies. This is in contrast to voting cases in which individual voters are collectively the *primary cause* of election outcomes, and thus could be claimed to be individually partial primary causes of those outcomes. As a result, it is incorrect to claim that animal consumption causes animals to be mistreated rather than treated humanely *in the same way in ethically relevant respects* that voting causes one candidate to win rather than another, or in the same way in ethically relevant respects that an individual causes harm to others in canonical cases of impermissible harm. In contrast to these canonical cases of being a *primary cause* of harm, individual decisions to consume factory farmed meat are *not part of the primary cause* from an ethical point of view of the harms that lie behind them. Similar remarks apply to most other cases at issue here. This illustrates why it is implausible to claim that consuming products typically *causes* the harm that lies behind them *in a way that is analogous in ethically relevant respects* to the way that voting causes a particular candidate to win rather than another, or in a way that is analogous to the way that an assault harms the victim.

As a result, deontologists cannot explain the collective action cases at issue here by simple appeal to the canonical constraint against harming others, because it is not clear why one should think that the canonical constraint against harming others is violated in the cases at issue even if we assume that such a constraint is violated by individuals who harm as part of a collective action that is the primary cause of harm. At the same time, further below, I offer a substantive theory that deontologists can use to argue that the constraint against harm is indeed violated in these cases; again, the point of this section is merely that in the absence of such a further substantive account, a simple gesture toward familiar deontological principles, including constraints against harming others, is inadequate for the reasons given here.

To further support this point about the inadequacy of a straightforward appeal to a constraint against doing harm, note that if the key premise of the inefficacy objection is correct,

in the cases at issue a single individual's 'non-cooperative' act also does not *make a difference to the harm* suffered by anyone, and it does not even *benefit those who do cause such harm*. As a result, deontological theories cannot respond to the inefficacy objection by appeal to side constraints against adding to the harm of others, or even by appeal to an alleged constraint against benefiting those who are harming others.⁷ The claim that an individual's act does not benefit those who cause harm in the cases at issue may seem surprising, but the argument is that although a single individual's consumption of, for example, animal products might have a very small but genuine effect on the revenues of *retailers* such as supermarkets and restaurants, at the same time if it is assumed that it does not make a difference to the number of animals produced (as the inefficacy objection claims), then by the same reasoning it should not make a difference to the profits of the factory farm *producers* who actually cause harm. In other words, if the inefficacy objection is correct that when supply chains are long and complex an individual's consumption cannot be expected to make a difference to *the quantity produced*, then a single individual's consumption also cannot be expected to make a difference to *the profits of producers* for similar reasons. For this reason, consuming ethically objectionable products generally does not benefit those who cause the relevant harm if the inefficacy objection is correct.

III. WHY FAMILIAR APPEALS TO COMPLICITY DO NOT PROVIDE AN ADEQUATE SOLUTION

The preceding observations undermine the common claim – familiar from deontological approaches to ethics – that individual actions that collectively lead to ethically undesirable

⁷ I defend the claim that mere participation in consumer society does not violate a canonical constraint against unjustly harming others in more detail against a number of further objections in [Budolfson unpublished a.](#)

outcomes – such as purchasing ethically objectionable products – are wrong on the grounds that one either violates deontological constraints in a straightforward way, or else *supports* objectionable institutions – for example, by *voting with one’s dollars* in a way that *benefits* firms. The arguments above suggest that these attempts at straightforward deontological explanations do not work, and motivate the positive view I will offer further below of a novel way of understanding the source of deontological reasons in these and other cases.

But some authors appeal to what they take to be a more general deontological constraint against *complicity* in harm according to which benefitting or supporting wrongdoers is not necessary for impermissible complicity. For example, Christopher Kutz seems to have in mind a more general prohibition on such complicity with his slogan “no participation without implication”.⁸ In light of this, some deontologists will insist that even granting all of the arguments in the previous section, there are nonetheless decisive reasons not to consume ethically objectionable products that arise indirectly from the fact that such consumer actions *involve* consumers in a system that causes harm or violates significant constraints, even if consumers do not thereby violate canonical constraints themselves.

However, if nothing more is said to fill out the detail of this theory, then such a simple appeal to the notion of complicity is implausible for a number of reasons. For example, such a simple appeal overgeneralizes and implies that you are almost never permitted to consume anything at all. For example, petroleum companies routinely violate significant constraints,⁹ and almost every possible consumption activity involves one with and supports such companies to at

⁸ Kutz 2000, pg. 122. See also Schwartz 2010, chapter 4 for discussion of complicity theories. Tom Regan’s main argument against consuming animal products might also be interpreted as implicitly relying on a very general constraint against complicity in the violation of rights: “Since [animal agribusiness] routinely violates the rights of these animals...it is wrong to purchase its products” (Regan 2004, p. 351). See also McPherson 2016 for more clear articulation of this idea.

⁹ For examples of routine abuses, see Maass 2009, especially chapters 2, 3, and 4.

least as great an extent (if not much greater) than, e.g., buying animals at supermarkets and restaurants depends on and supports factory farms.¹⁰ As a result, if such a simple notion of complicity really did give rise to strong reasons not to consume products, then it would also be wrong not only to consume petroleum products because of the oil industry's complicity in serious harm, but it would also be wrong to consume almost everything else as well, because almost everything depends on petroleum products via dependence on transportation companies, which turn a blind eye to oil companies' abuses that are known to benefit transportation companies in the form of lower fuel costs.

This reveals that almost every consumption activity is complicit in harm in the sense that it depends on and supports companies that violate important constraints to a similar extent as, e.g., our consuming factory-farmed meat does. But despite all of this, our considered judgment is that it is nonetheless permissible to consume *some* everyday products, even if we judge that some others, such as perhaps factory farmed meat, are impermissible to consume. As a result, a simple general appeal to complicity overgeneralizes in an implausible way, and cannot do the ethical work that is needed of distinguishing and explaining the difference between highly objectionable and relatively unobjectionable actions that are in some natural way connect us to wrongdoing.¹¹ With similar thoughts and a rich analysis of their own, Chiara Lepora and Bob Goodin summarize the upshot for simple appeals to complicity:

¹⁰ Note especially that individual purchases of gasoline for personal use are often permissible even though we thereby purchase gasoline directly from the petroleum companies themselves, or are at least only one step in the supply chain removed from such companies – and in the gasoline supply chain there is much less of a buffer caused by waste and inefficiency than in the supply chain for, e.g., animal products.

¹¹ This is not to say that such considerations of complicity do not generate *any* reason for individuals to forego consuming petroleum products and engaging in other such acts on the basis of complicity – the point is merely that the reasons not to engage in acts simply on the basis of such a thin notion of complicity must be more easily *outweighed* by one's own interests, because they are much weaker than the reasons that arise when it is *more essential* to consumption and other acts that harm or violation of rights must lie behind them (as is argued in more detail in the next section).

‘Complicity,’ in the way that catch-all term is commonly deployed in ordinary discourse, is morally pretty uninformative. It amounts to little more than empty name-calling. Differentiating in a more fine-grained manner between various more precise ways of being ‘complicit’ is morally more revealing. Morally, there is a big difference between collaborating in wrongdoing yourself and condoning others’ wrongdoing, between conspiring to do wrong and conniving in others’ wrongdoing by turning a blind eye to it.¹²

In the current context, this point can be made even sharper, by noting that in light of the preceding considerations, by participating in consumer society each of us is complicit in harm only in a sense that is more indirect, diffuse, and thus less objectionable than any of the senses of complicity that Lepora and Goodin note here.

In a later section I develop a view that provides a principled distinction of the sort that Lepora and Goodin note is needed between more objectionable and more innocuous forms of complicity. Before developing that positive view, in the next two sections I first consider other avenues that may initially seem promising to deontologists in light of the inefficacy objection.

IV. WHY FAMILIAR APPEALS TO UNIVERSALIZABILITY DO NOT PROVIDE AN ADEQUATE SOLUTION

¹² Lepora and Goodin 2011, pg. 53. See also Lepora and Goodin 2013.

It may still seem that a straightforward response is open to deontological defenders of *universalizability theory* – i.e., rule consequentialism, cooperative consequentialism, Kantian ethics, and so on. However, initial appearances to the contrary notwithstanding, such theories do not provide a successful response to the inefficacy objection.

Turning first to theories that are explicitly Kantian in order to show that they do not provide a successful response to the inefficacy objection, it is useful to begin by noting how an ideal Kantian might make ethical decisions, in order to then argue that such an agent would act on consumption maxims that pass the Kantian universalizability test, but yet do not clearly counsel against e.g. consuming ethically objectionable products, and do not clearly counsel against other forms of ‘non-cooperation’ at issue here.

With that in mind, ideal Kantian ethical reasoning might be modeled as deriving maxims that apply in specific situations from a general 'way of life maxim' that captures the general goal of respecting others with appropriate respect for oneself, such as "I will to respect the agency and interests of humanity [and perhaps also the interests of other sentient creatures] and, subject to that constraint, promote my own projects and interests". With this ideal way of life maxim in hand, an ideal Kantian would then derive specific maxims for action in specific circumstances. In the cases of e.g. consumption at issue, the more specific maxim that follows from this general maxim appears to be something like: "In connection with consumer goods, I will: that I not interfere in the agency of others or violate the familiar constraint against harming others, that I support just institutions, regulations, and other forms of social progress that would eliminate current forms of harm and deprivation of agency and that would promote our collective interests, and subject to those constraints, that I promote my own projects and interests". This maxim passes all four steps of the universalizability test described in connection

with Kant's first formulation of the categorical imperative, and thus actions based on this maxim are permissible on the standard interpretation of the Kantian view.¹³

However, a person who acts on such a maxim may consume ethically objectionable products and act 'non-cooperatively' in many of the cases at issue here, as in those cases by doing so one would not obviously violate any of the constraints mentioned in the first part of the maxim, partly because doing so does not violate the familiar constraint against interfering with or harming others (as argued above), and partly because doing so is consistent with also supporting proper regulation and other forms of social progress to redress the relevant harms in the background. Thus, absent further argument, there is no clear Kantian response to the inefficacy objection based on the categorical imperative, since there is no clear reason for condemning most of the actions at issue on the basis of the categorical imperative. Further below, I offer a substantive theory that Kantians can endorse according to which constraints against interference or harm are indeed violated in the cases at issue; the current point is merely that in the absence of such a further substantive account, a simple appeal to familiar Kantian notions is inadequate for the reasons given here.

¹³ Here I follow the interpretation of Kant as providing a decision procedure via the first formulation of the categorical imperative for determining whether an act is permissible. For example, following John Rawls and Onora O'Neill, Robert Johnson summarizes this decision procedure as: "First, formulate a maxim that enshrines your reason for acting as you propose. Second, recast that maxim as a universal law of nature governing all rational agents, and so as holding that all must, by natural law, act as you yourself propose to act in these circumstances. Third, consider whether your maxim is even conceivable in a world governed by this law of nature. If it is, then, fourth, ask yourself whether you would, or could, rationally *will* to act on your maxim in such a world. If you could, then your action is morally permissible. If your maxim fails the third step, you have a 'perfect' duty admitting "of no exception in favor of inclination" to refrain from acting on it. If your maxim fails the fourth step, you have an 'imperfect' duty requiring you to pursue a policy that can admit of such exceptions. If your maxim passes all four steps, only then is acting on it morally permissible" (Johnson 2008 [2004]). For an alternative Kantian approach to these issues, see Wallace forthcoming.

On a related note, deontological theorists also cannot appeal to the doctrine of double effect, because for almost every non-pathological person who is aware of the harm that lies behind many of the products we consume, that harm is a foreseen but *unintended* side effect of our consumption, and so is not objectionable on grounds of double effect – and this is true whether or not the inefficacy objection is sound. So, the doctrine of double effect is particularly ill-suited to explain why it is objectionable to consume products that are produced in objectionable ways.

Setting aside Kantianism and double effect, other universalizability theories seem subject to counterexamples insofar they are interpreted as giving any kind of straightforward response to the inefficacy objection. To see why, consider examples like the following:

Stampede Case

We find ourselves in an enormous stampede. Unless everyone stops stampeding, it is clear that an increasing number of innocent people will be seriously harmed and killed. However, it is also clear that there is virtually no chance that everyone will stop stampeding soon, and so anyone who stops stampeding will almost certainly be seriously harmed or killed in a way that does no good for anyone else.

Charity Case

You have a choice between giving to Charity A and Charity B. Charity A does a lot of good per dollar on average, but you know that your donation to A would not result in any more good being done, because it would amount to no more than insignificant digits in A's enormous budget. In contrast, Charity B does slightly less good per dollar on average,

but you know that your contribution will result in your funds being transferred directly into the bank account of several desperately poor people who would not otherwise receive money if you did not contribute to B. So, the positive difference you would make by giving to Charity B is very large, whereas the positive difference you would make by giving to Charity A is zero.

These cases are analogous to real-world collective action problems in which ethical actors are ‘thrown into’ a non-ideal situation in a way that makes action required in their circumstances that is in some clear sense non-universalizable (as in the Stampede Case), or where the average effect of individual action of a particular type is very different than the actual (marginal) effect of one additional action of that type, (as in the Charity Case). In the Stampede Case, the thing to do is to continue stampeding despite the fact that universalizability principles imply that one should stop stampeding because everyone [in the same relevant circumstances] [acting on a maxim that implies] continuing stampeding would be dramatically worse than everyone [in the same relevant circumstances] [acting on a maxim that implies] stopping stampeding. In the Charity Case, the thing to do is to give to charity B, despite the fact that the average effect [the ‘footprint’¹⁴] of giving to charity B is worse than the average effect of giving to charity A [in the relevant circumstances], and despite the fact that some universalizability principles imply that one should give to charity A because everyone [always acting on a maxim that implies] giving to charity A would be significantly better than everyone [in the same relevant circumstances] [always acting on a maxim that implies] performing any of the relevant alternative acts. At the

¹⁴ The *footprint* of an act of a particular type is simply the *average individual effect* of all actual acts of that type. The most familiar use of such footprint reasoning in consumer ethics is in connection with the carbon footprint of various goods.

very least, these are the considered judgments that most of us have about what there is reason to do in these examples.

This shows that a dilemma exists in connection with universalizability theories and collective action problems: the theories could be understood in a way that provides a response to the inefficacy objection, but in that case they deliver (what most would judge to be) mistaken verdicts on many cases involving general non-cooperation such as the Stampede Case and Charity Case above; alternatively, they can be understood in a way that does not deliver these counterintuitive implications, but then they do not generate any clear response to the inefficacy objection. The upshot is that universalizability theories do not give rise to any clearly satisfactory response to the inefficacy objection. Furthermore, universalizability theories are often intended by their adherents to be understood as not implying these counterintuitive verdicts, and thus do not generate any clear response to the inefficacy objection even by those adherents' own lights.¹⁵

V. WHY FAMILIAR APPEALS TO SYMBOLIC VALUE AND VIRTUE THEORETIC NOTIONS DO NOT PROVIDE AN ADEQUATE SOLUTION

In light of all of these problems for straightforward attempts at a deontological reply to the inefficacy objection, and given the assumption here that a consequentialist reply is also unavailable, it is worth noting that virtue theoretic notions such as 'symbolic value'¹⁶ also cannot

¹⁵ Similar remarks apply regarding rule utilitarianism, cooperative consequentialism, and causal consequentialism in connection with the inefficacy objection and the Stampede Case and Charity Case. I discuss universalizability theories in more detail in [Budolfson unpublished b](#) and [Budolfson unpublished c](#).

¹⁶ For relevant discussion of this kind of reason see Hill 1979, Hill 1983, Appiah 1986. For further discussion of symbolic value and the ethics of collective action, see Nefsky 2018.

be imported into deontology in any straightforward way to solve the problem. That is because even if it is granted that one might have *sufficient* reason deriving from symbolic value or some other source of virtue to perform the symbolic act of, say, refusing to consume objectionable products even though it is costly to oneself to do so, that does not make it plausible that such an act is *required*, because it is still much more plausible to think that such an act would be *supererogatory* given the general defeasibility of such distinctively virtue-theoretic reasons. In other words, although it is plausible that there are genuine reasons for action that arise from considerations of symbolic value, a general fact about such distinctively virtue-theoretic reasons is that they are not generally of sufficient weight *when the costs of pro-social action are high* to make that action required, as opposed to merely rational and supererogatory. Furthermore, as noted above, participating in almost every normal activity in contemporary society would make us complicit in harm in some important sense, and so an appeal to symbolic value risks overgeneralizing implying that almost every normal activity is impermissible because of the symbolic disvalue that it expresses. Furthermore, as Julia Nefsky has argued, such a simplistic appeals to symbolic value often assumes an implausible view of what is communicated by the individual actions at issue – when in fact there is often no good reason to assume that non-cooperative acts that are known to occur in a collectively non-ideal context communicate endorsement of a collectively non-ideal outcome.¹⁷

Taken together, the preceding sections show that even if we assume that it is wrong not to vote in typical elections, that does not imply that there is any plausible reply to the inefficacy objection to deontology. Assuming that the key premise of the inefficacy objection is correct, a satisfactory deontological account must therefore look for a previously unappreciated source of reasons for action.

¹⁷ Nefsky 2018.

VI. THE SOLUTION: ESSENTIALITY OF HARM AS THE MOST OBJECTIONABLE FORM OF COMPLICITY IN HARM

What then is the source of the strong reasons we have not to be complicit in harm in some cases? What explains the difference between those cases and others in which complicity is much more innocuous? More generally, how can deontology capture all of the reasons for action one has in collective action situations, especially when one's actions don't make a difference? After all, most people would agree that even when what one does really makes no difference, violates no rights, and does no harm, there are still *some particularly objectionable ways* of being complicit in harm, even though the discussion above indicates that not all ways of being complicit are particularly objectionable.

If an appealing account can be identified of what explains the difference between cases of objectionable vs relatively unobjectionable complicity in harm, and if that account delivers a compelling and satisfying package of theory and verdicts on cases in our considered judgment, then it will amount to an appealing account of what exactly deontology should say about the inefficacy objection and the ethics of collective action.

The proposal here is to ground the needed deontological explanation on *the degree to which it is essential to an act that (unjust) harm or the violation of basic rights lies behind it*. The idea is that there is a continuum between, at one end, metaphysical necessity of harm or rights violations lying behind an act, to, at the other end, no connection at all to harm or rights violations via an act, with many intermediate degrees and dimensions along which it is more or less essential to an act that harm lies behind it – and as harm becomes more essential to the success of an act, the strength of the deontological reasons one has not to perform that act

increases, even if it is held fixed that no such act does any harm itself. I explain the details of this proposal in the next section, and in the rest of this section focus on the intuition for the idea.

To illustrate the idea, consider a can of vegetables sold at a supermarket that is produced in a normal way. Although substantial use of fertilizers and petroleum products might actually lie behind those vegetables, and thus those vegetables might have a surprisingly high footprint of harm, it is *highly inessential* to consuming those vegetables that such harm occurs in the background, because, for example, a simple change in regulations could almost immediately remove most of those harms while leaving the nature of the product and its creation the same. More generally, there is nothing in the nature, actual consumption, or any other aspect of the ethically relevant *modal neighborhood* most closely associated with consuming those vegetables that necessitates harm or the violation of significant constraints. The suggestion here is that this explains why you do not have strong reasons not to consume a can of vegetables, and do not have strong reasons not to perform many other actions that do not essentially require harm, even when you know that performing those actions makes you complicit in harm in an indirect sense.¹⁸

Such an appeal to the degree of essentiality of harm and rights violations also exists in the law. In *New York v. Ferber* and in *Ashcroft v. Free Speech Coalition* the US Supreme Court held that a compelling state interest exists to prohibit the promotion and consumption of child

¹⁸ So, as a further illustration of the basic idea: if an evil person with enormous power in politics and the marketplace causes the harm footprint of each unit of food that we consume to be 25% worse by bribing regulators to allow him to cause this harm, but yet your choices as an individual make no difference to the amount of this harm that lies in the background, then the correct conclusion to draw from a deontological perspective is not that your food choices have become 25% more ethically objectionable – on the contrary, although you now have *some* additional reason other things equal not to consume those products, only the actions of the evil player and the corrupt regulators are seriously wrong, and they carry nearly all of the weight of additional badness that has been added to the system. In contrast, if consumers develop a taste for unnecessary objects for which it is *essential* to those goods being produced that innocent persons suffer, then the actions of both consumers and producers are seriously wrong and both bear the weight of the additional badness that is created.

pornography insofar as that pornography is “intrinsically related to the sexual abuse of children”.¹⁹ This can be seen as an appeal to the degree of essentiality of harm, and as the basis for distinguishing between, on the one hand, objectionable child pornography that may be constitutionally prohibited by legislation, and, on the other hand, constitutionally protected and arguably unobjectionable ‘non-pornography’ that nonetheless involves depictions of children engaging in sexual acts – for example, depictions in mainstream films and textbooks, as well as *virtual* depictions of child acts (e.g. via realistic digital animation) that would be prohibited were they to involve actual child performers, but which are protected because it not essential to such materials that child abuse lies behind them.²⁰ The suggestion here is that this distinction tracks fundamental deontological reasons, and thus that an analogous explanation should be endorsed in ethical theory – for example, for why individuals have decisive reasons not to consume child pornography even when an individual act of consuming that product need not have any negative effect on anyone, as in a case where it is freely downloaded from a site in the privacy of a person’s home in a way that doesn’t create any incentive for additional production of that material.

¹⁹ *New York v. Ferber*, majority opinion by Justice White.

²⁰ There are a number of other analogous uses in the law of the notion of degree of essentiality. For a related discussion of other dimensions of complicity, that is especially richly developed in connection with human rights law, see Lepora and Goodin 2013. For examples tied to law and the marketplace, World Trade Organization (WTO) decisions increasingly turn on how essential it is to a product or process that violation of fundamental societal values lies behind (given that domestic regulations to protect such values are permissible under the WTO, and can then be imposed on imports under terms that are “no less favorable” (in connection with GATT Article III(4)). For example, one specific issue is whether blanket prohibition of foie gras imports are permissible on the grounds of domestic humane treatment requirements, given that it is debatable how essential it is to foie gras *as it is defined in the relevant regulations* that it is raised inhumanely. In the future similar issues may arise regarding how essential it is to products or processes that harm to human, animal, or plant life or health lie behind them (in connection with GATT Article XX(B)). See WTO disputes *Canada Seals*, *California Foie Gras*, and *California Low-Carbon Fuel*.

To further illustrate the proposal here in the context of the marketplace, suppose you learn that the avocados that you consume are distributed via a supply chain that is largely controlled by criminals who routinely commit human rights atrocities – as is allegedly true for many avocados consumed in North America.²¹ Although that harm is serious, it is also true that it is in no way essential to the avocados themselves, or to the avocado growing process itself. In such a case, buying the avocados need not be *seriously wrong*, even if there are defeasible reasons that tell against consuming the avocados. This is in contrast to a similar but crucially different case involving the same harm footprint in which it is *more essential* to avocados that such harm lies behind them, perhaps because (as a fantastic example) in this second imaginary case they are of an engineered variety that must be watered with the blood of healthy people in order to thrive.

As another example, imagine a cream to treat a skin condition that is produced by extracting a substance from the bodies of innocent human beings who are painfully killed in the process, where this substance cannot be procured in any other way. Although the process is as horrible as it sounds, suppose the amount of the substance obtained from one innocent person is enough for about 100,000 bottles of the cream; at the same time, other creams are made in a normal way that does not harm anyone, and they are almost as effective at treating the skin condition. However, the factories that produce this more normally produced treatment require more inputs than their horrible competitor, and behind those inputs lies the misconduct of their suppliers and the petroleum companies that provide the fuel to transport their goods and run those factories. As a result, both kinds of skin cream have the same footprint of harm and rights violations once the violations of suppliers and petroleum companies are taken into account,

²¹ See De Cordoba 2014. This case (and the frequently discussed case of marijuana) are arguably more complicated if it is true that the basic production system would be very substantially changed if the background harm by others in the supply chain changed. See the discussion of essentiality of origins below.

although the footprint of the operations under their direct control differs dramatically. In light of all these facts and the assumption that it will not make any difference to human suffering whether you personally consume either of these products, it would still be worse to consume the first kind of skin cream than the second. If this is not immediately clear, imagine that, because of a serious skin condition, you must use one of these two kinds of treatments. Is it a matter of moral indifference which you chose? The considered judgment of most people is that in such a case one must choose the second rather than the first.²²

As another example, consider the difference between consuming ordinary goods like watches that were merely produced by the Nazi regime (without slave labor, etc.), and consuming soap and other products that were made directly from the bodies of those killed by the Nazi regime. In both cases, consuming the products now would have no important effect on others, and could not, of course, *support* the Nazi regime. However, there are still strong reasons not to consume products that were made in a ruthless way from the bodies of innocent people at the direction of an evil regime, whereas there are not strong reasons not to consume products like watches that were manufactured in an ordinary way merely at the direction of such a regime. The distinction above provides a straightforward explanation of the difference, because it is fairly essential to the actual production of the soap that humans are used in unacceptable ways that violate significant constraints, whereas it is fairly inessential to the actual production of the watches that anyone is used in unacceptable ways.

²² A further issue I will not pursue here is how genuine consequentialist reasons interact with and trade off against the deontological reasons I am defending here. For example, consider a modification of the skin treatment cases: as before, there is no marginal effect of a single individual consuming the bad treatment, but this time suppose that consuming the second normally produced treatment does incrementally increase the amount of harm and other bad consequences – perhaps equal to $\frac{1}{2}$ of its harm footprint. How should we think about the weight of all things considered reason in this new case, assuming it is extremely important that an individual use one of these two treatments? My inclination is to think that the deontological reason not to consume the bad cream still outweighs the consequentialist reason to consume the bad cream, but further discussion would be needed, and the opposite view may appeal to some.

Again, this notion also offers better explanations than the theoretical alternatives, such as an explanation in terms of a simple appeal to complicity in harm. For example, in the case just described both the soap and the watches were – let us suppose, for the sake of a clear example – produced by the Nazi regime itself, which means that the producer of both was equally complicit in violating constraints – because it was, after all, the very same producer in both cases – and a consumer is equally complicit in whatever sense is claimed to be relevant, since (it is easy to suppose) there is the same connection in all relevant respects between the consumption and the producer in both cases. Nonetheless, there is an important difference in the permissibility of consuming those goods. Furthermore, as noted above, if a simple appeal to complicity really did give rise to strong reasons not to consume products, then it would be wrong to consume petroleum products because of the oil industry’s complicity in serious harm, and it would be wrong to consume almost everything else as well, because transportation companies are connected to the oil companies’ abuses that are known to benefit transportation companies in the form of lower fuel costs. Again, despite all of this complicity in harm, it is nonetheless permissible to consume many such products, and the best explanation appears to be the one offered here that it is *far from essential* to acts of consumption of such things that suffering is caused and significant constraints are violated in the background.²³

VII. MODAL ANALYSIS OF THE ESSENTIALITY OF HARM, AND ITS POTENTIAL AS A UNIFIED SOLUTION TO OTHER CHALLENGES FOR DEONTOLOGICAL THEORIES

²³ One way of thinking about an upshot of all of this for deontological practical ethics is that in evaluating the reasons an individual has for and against various acts, the harm footprint of each act should be discounted by the degree of essentiality of that harm to that act. This grants the idea that harm footprints generally result in reasons for action, and that the size of the harm footprint always matters, but avoids the implausible implication that nothing else matters under conditions of inefficacy. Instead, it endorses the compelling idea that the extent to which acts are connected to particular harms is highly variable in an ethically important way, and provides a framework for further analysis of those degrees of connection.

The philosophical assumption behind an appeal to the *degree of essentiality of harm* is that it is possible to define an ethically and metaphysically respectable metric that measures the extent to which a departure from actuality is required in order for the same act to occur without the harm or the violation of rights that actually lies behind it – the greater the departure that is needed, the *more essential* it is to that act that harm or the violation of rights lies behind it, and the stronger the reasons there are not to perform that act. In supposing that such a metric can be defined, the proposal in this section is to follow the approach associated with David Lewis and Robert Stalnaker in supposing that analogous metrics can be defined using modal notions, which has led to great progress in analyzing counterfactuals and other fundamental philosophical notions.²⁴

In addition to the analogy to the analysis of counterfactuals, another useful analogy is from epistemology: after Gettier showed that justification, truth, and belief were insufficient to draw the distinctions that needed to be drawn in the domain of knowledge,²⁵ a kind of reliability was suggested as the crucial additional property relevant to knowledge, and much progress was then made over several decades by using similar modal notions to home in on what exact notion of reliability might be the crucial one relevant to knowledge.²⁶ With this example in mind, the structure of the current paper is analogous, as it argues that existing deontological notions are insufficient to draw the distinctions that need to be drawn in the domain of ethics, and then suggests that the degree of essentiality of harm is the crucial additional property that explains the fundamental facts to be analyzed, and then provides some indication of how we can begin to

²⁴ For example, Stalnaker 1968, Lewis 1973, Lewis 1979.

²⁵ Gettier 1963.

²⁶ See e.g. Nozick 1981, Lewis 1996, and Williamson 2000 for increasingly attractive modal accounts of knowledge.

home in on the exact notion of degree of essentiality that is relevant to ethical reasons by using modal notions.

To put the idea another way, the discussion in previous sections suggests that we need to make careful and subtle distinctions in order to distinguish the types of complicity in harm that are particularly objectionable from those that are not; the suggestion here is that we can ground the needed distinctions in the structural idea of the degree of essentiality of harm, and that modal notions provide useful tools for further clarifying and ultimately working toward an ideally precise analysis of the idea.

So, although I focus in this paper on the structural features of the analysis of the degree of essentiality of harm that is relevant to ethical discourse rather than the precise details of the analysis, some initial remarks are desirable to illustrate how modal notions provide further initial illumination of the general idea of the essentiality of harm, and to provide some initial useful substance, which can then be refined by future work that aims to more precisely home in on a full substantive analysis. As in the original presentation of the basic structure of the Lewis/Stalnaker account of counterfactuals, the main point here is not to defend a complete and maximally precise analysis of how exactly the relevant orderings should be made, but rather to note the virtues and explanatory power of an analysis with this basic structure, which postulates a continuum of greater and lesser departures from actuality that represent how essential it is to an act that harm lies behind it.²⁷

²⁷ Chiara Lepora and Bob Goodin have independently arrived at a similar view, which endorses something like the degree of essentiality of harm or violations of rights along many dimensions, and provides a rich analysis in connection with human rights law and other “conceptual cousins” of complicity. See Lepora and Goodin 2011 and Lepora and Goodin 2013. Shepski 2013 and Cuneo 2016 have independently developed somewhat related thoughts about supply chain ethics and animal consumption respectively, both of which could be analyzed using the modal apparatus introduced here, although those authors do not pursue such an analysis. For application of a methodologically similar semantic and metaphysical analysis to core social and political issues, see Goldman 1972 and Goldman 1974.

In articulating the modal relation that constitutes the degree of essentiality of harm relevant to ethical discourse, presumably some aspects of actuality should be held more fixed as in Lewis's substantive account of counterfactuals, and some kinds of departures from actuality should be weighted as more important than others – and weightings should reflect our considered judgment of what is ethically important in particular cases. So, in evaluating the degree of essentiality of harm that lies behind a particular act, presumably the actual physical laws governing cause and effect should be held fixed, as well as (less stringently) facts about technological possibility; then departures from actuality that are more 'closely connected' to that act itself should be weighted as more important than departures from actuality that are not at all 'central' to that act or that occur 'far in the background'.²⁸ Then, the further from actuality one has to go to find cases where such an act is successful but harms do not lie behind it, the more essential it is to that act that harm lies behind it, and thus the stronger the reasons one has of this type to avoid that act.

From this starting point, substantive views can diverge on the substance of how exactly this basic structure should be filled out, just as views diverge in connection with modal analyses of counterfactuals, knowledge, and so on. For example, my own preference is to incorporate Kripkean ideas about the essentiality of origins of objects that are mentioned in the agent's intentions behind an action, but others might disagree when developing a more precise view within the same basic structure.²⁹ A more controversial substantive principle I endorse, building on essentiality of origins and in accord with common sense, is that when goods are produced via particular intentionally constructed production schemes and business plans, the essential parts

²⁸ Lewis 1979. The relevant form of technological possibility should also be thought of as a scalar notion, as some things are (so to speak) more technologically infeasible than others, in the sense that they would require much greater investments in engineering, etc. to realize.

²⁹ For the essentiality of origins argument, see Kripke 1980, pp. 111-114.

of those schemes become more robustly counterfactually tied to those goods in the evaluation of the degree of essentiality of harm. So, the more essential it is to a production scheme that people are used in an impermissible way, the more essential it is to the goods produced by that scheme that harm lies behind them.³⁰ Adding these substantive principles yields the verdicts on specific cases that I assume are correct in this paper.

To illustrate how these details yield an explanation of a classically vexing case, consider the use of medical knowledge that was originally gained from unethical experiments on humans during the Nazi era. Suppose the unethical experiments were conducted in 1942, and (as seems to be the actual situation) that it was not technologically feasible to obtain the medical knowledge using means other than human experimentation in 1942, but it would be easily feasible to obtain that knowledge using other means today. Then my account provides a straightforward account of why there were initially strong ethical reasons against using that knowledge, but also explains why there are no longer strong reasons not to use that knowledge, and explains this in a way that avoids grounding the explanation in a problematic way on mere temporal distance. Instead, on the view defended here the explanation is that it was fairly essential to using that knowledge in 1942 that humans must have their rights violated relative to that point in time, whereas it is not essential at all to using that knowledge in 2018 that humans must have had their rights violated.

³⁰ For example, consider sweatshop clothing. Some sweatshops violate significant constraints by chaining people to their sewing machines and forcing individuals to work in horrible conditions that could never be justified as an essential part of their production scheme and business model, whereas other sweatshops merely pay workers a very low wage and insist that they work very long hours, while still arguably offering such workers better prospects than they would have if the sweatshops did not exist. Harm is much more essential to the former goods produced than the latter. In this way, actual facts about production are arguably relevant to how essential particular harms are to goods, just as facts about the actual situation are relevant to what would have happened if some counterfactual assumption were true on standard accounts of counterfactuals. As this indicates, the ethical issues regarding the consumption of products even of particular kinds often requires a careful case-by-case evaluation depending on the specifics of the case – and, as in the case of sweatshops, often requires a substantive view on the significance of the constraints that are (or are not) violated by various kinds of exploitation.

With the basic structure of the degree of essentiality in hand, further investigation and precisification can proceed (as in discussions of counterfactuals, knowledge, and so on) by further investigation of important test cases, and by investigation of how more general principles and classic puzzles might be addressed with such a framework. For example, it merits further investigation whether the kernel of truth in the principle of double effect can be understood as more fundamentally explained by considerations of the degree of essentiality of harm – and whether the degree of essentiality of harm points the way toward a principled distinction between the ‘core’ intuitive motivation for double effect reasoning, while avoiding some of its troubling implications. As another example, it merits consideration whether an appeal to the degree of essentiality of harm can allow further progress on classic puzzles such as the trolley problem, the consideration of which might at the same time shed further light on how exactly the details of the current proposal ought to be filled out.³¹ In addition, the theory offered here has the virtue of giving deontological theories greater resources to explain plausible *differences in the strength of reasons* that individuals have in many cases that might otherwise be difficult to explain by appeal to traditional deontological principles alone.

For current purposes the most important point is that the degree of essentiality of harm is a notion that can be cashed out in terms that are metaphysically respectable, in addition to being embedded in common sense and the law. As a result, an appeal to the degree of essentiality is a good candidate for the best explanation of the cases at issue, especially in light of the failure of more familiar ethical notions to adequately explain these cases.

But why would morality make such a distinction between degrees of essentiality of harm? If there were no good answer to this question, then there would be some reason for skepticism about whether a principle that invoked the degree of essentiality of harm was a

³¹ For example, the structure of the current proposal could be taken to provide a more general theoretical motivation for a view such as that defended in Quong 2009.

genuine moral principle, even if it seemed to correctly capture our initial intuitive judgments about cases, and even if it could be given a satisfying modal analysis. However, on further reflection it is not entirely surprising that morality would make such a distinction, as it seems to be a consequence of the compelling deontological idea that moral reasons are sensitive to agent-centered facts about what an individual agent does *herself* and what is under her *control*. If it is highly inessential to an action that harm lies behind it, but at the same time background actions by others that one cannot change and that are far removed from any direct connection to any action of one's own would give that option a high harm footprint, the current proposal is that morality does not assign that as much negative weight as if a similarly high degree of harm would be directly caused by one's action or fairly essential to an act. On the theory defended here, this is a fundamental way that morality accounts for facts about how closely harms and other effects are tied to the different actions of different individuals.

VIII. CONCLUSION: INTEGRATING THE DEGREE OF ESSENTIALITY OF HARM INTO DEONTOLOGICAL THEORY

One way of integrating reasons based on the degree of essentiality of harm into deontological theory would be to integrate them into a rights-based view. On such a view it is uncontroversial that even when our individual decisions make no difference, others nonetheless have a claim against us that we not engage in activities that harm them unjustly and violate their basic rights. To integrate reasons based on the degree of essentiality of harm, a theorist can add that in such cases of inefficacy others *also* have a claim against us that we not engage in activities for which it is highly essential for the success of our activities that they are harmed unjustly or that their rights are violated. These claims can then be seen as a previously untheorized part of the bundle of basic rights possessed by all persons, and thus as a previously unarticulated part of the rights

protected by a fundamental deontological constraint against harming other and violating their basic rights. So, on this way of merging the proposal above with rights-based deontology, when e.g. one consumes a product such that it is highly essential to it that the basic rights of others must be violated in order for it to be produced, one thereby violates the rights of those others in virtue of violating this kind of previously unnoticed claim that they have against one. On this view, rights-based deontology is ultimately correct, but it is incomplete in the form it has been traditionally articulated, and cannot provide a plausible account of all cases if it is not combined with the proposal above in this way.³² There are also a number of alternative ways of incorporating reasons based on the degree of essentiality of harm into deontological theories – for example, a non-rights-based deontological view is also attractive on which this new kind of reason is merely one among many kinds of deontological reasons that trade off against each other and against welfarist reasons in a way that is not easily conceptualized using familiar notions of rights.

In sum, the phenomenon of inefficacy raises important problems for deontological theories, and ultimately shows that existing deontological theories are unable to offer a plausible account of what individuals are required to do in a wide range of collective action situations that

³² If we add to this the more controversial view that some crude subset of these claims (such as a claim against others not to cause one serious harm for no important reason) are possessed by non-human animals (where perhaps they have varying stringency as a function of whatever ultimately explains the general differences in moral status between humans and non-human animals), then a deontological view emerges of the wrongness of consuming some animal products. In support of this, it could be further argued that it is no less plausible to think that many animals have a claim against us not to be seriously harmed for no good reason than it is to think that very young, uncared-about children have such claims – and, as reflection on test cases reveals, it is hard to deny that such children have such claims. Again, unless we believe that such children have such claims, it is hard to see how we could *correctly* explain why it is wrong for an individual to watch freely-downloadable child pornography in the privacy of his home in a case where such an individual's actions are known to make no difference to anything outside his home, including to child abuse. For somewhat related discussion, see Feinberg 1974 and Tooley 1972. Here again, it is worth noting that the doctrine of double effect cannot explain these cases despite possible initial appearances to the contrary, because many actual viewers of child pornography do not have the intention that children be harmed, much less that additional children be harmed.

are common in modern society. In light of this, there is good reason to endorse a previously unnoted source of deontological reason that depends on *the degree to which harm or the violation of rights is essential to the success of an act*. Appeal to such a notion is well motivated, grounded in common sense, familiar from the law, metaphysically coherent, and offers the potential for further progress in resolving a number of other classic challenges for deontology. Going forward, the exact nature and implications of this notion can be further clarified using the modal tools outlined above, and it can be used to explore the potential for more fundamental explanations of otherwise disunified phenomena and classic puzzles within deontological theory.

REFERENCES

Appiah, Kwame Anthony. 1986. "Racism and Moral Pollution". *Philosophical Forum* 18:185-202.

Barry, Brian. 1978. "Comment", in Stanley Benn ed., *Political Participation*. Canberra: ANU Press.

Barry, Christian and Øverland, Gerhard. 2016. *Responding to Global Poverty: Harm, Responsibility and Agency*. Cambridge: Cambridge UP.

Brennan, Geoffrey, and Lomasky, Loren. 1993. *Democracy and Decision*. New York: Cambridge UP.

Brennan, Jason. 2012. *The Ethics of Voting*. Princeton: Princeton UP.

Budolfson, Mark. 2018. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response", *Philosophical Studies*. <https://doi.org/10.1007/s11098-018-1087-6>.

Budolfson, Mark. unpublished a. "The Modern Predicament: Harm, Complicity, Compensation, Moral Debt".

Budolfson, Mark. unpublished b. "Collective Action, Climate Change, and the Ethical Significance of Futility".

Budolfson, Mark. unpublished c. "Why Morality and Other Forms of Normativity are Sometimes Directly Collectively Self-Defeating".

Cuneo, Terence. 2016. "Conscientious Omnivorism", in Chignell, Cuneo, and Halteman eds. *Philosophy Comes to Dinner*, New York: Routledge.

De Cordoba, Jose. 2014. "The Violent Gang Wars Behind Your Super Bowl Guacamole," *Wall Street Journal*, January 31, 2014.

Driver, Julia. 2016. "Individual Consumption and Moral Complicity", in Bramble and Fischer eds. *The Moral Complexities of Eating Meat*. Oxford UP.

Feinberg, Joel. 1974. "The Rights of Animals and Future Generations", in Willian Blackstone ed. *Philosophy and Environmental Crisis*, University of Georgia Press.

Gettier, Edmund. 1973. "Is Justified True Belief Knowledge?", *Analysis* 23: 121-123.

Goldman, Alvin. 1974. "On the measurement of power", *The Journal of Philosophy* 71: 231-252.

Goldman, Alvin. 1972. "Toward a Theory of Social Power", *Philosophical Studies* 23, 221-268.

- Goldman, Alvin. 1999. "Why Citizens Should Vote: A Causal Responsibility Approach", *Social Philosophy and Policy* 16, 201-217.
- Hill, Thomas E. 1979. "Symbolic Protest and Calculated Silence". *Philosophy & Public Affairs* 9: 83-102.
- Hill, Thomas E.. 1983. "Ideals of Human Excellence and Preserving Natural Environments", *Environmental Ethics*, 5: 211-24.
- Johnson, Robert. 2008 [2004]. "Kant's Moral Philosophy", *Stanford Encyclopedia of Philosophy*.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs*, 39, 105-41.
- Kripke, Saul. 1980 (1972). *Naming and Necessity*. Harvard: Harvard UP.
- Kutz, Christopher. 2000. *Complicity*. Cambridge: Cambridge UP.
- Lawford-Smith, Holly. 2015. "Unethical Consumption & Obligations to Signal", *Ethics and International Affairs* 29:315-330.
- Lepora, Chiara, and Goodin, Robert. 2013. *On Complicity and Compromise*. Oxford: Oxford UP.
- Lepora, Chiara, and Goodin, Robert. 2011. "Complicity and its Conceptual Cousins", online at http://philrsss.anu.edu.au/sites/default/files/Lepora-Goodin_Complicity%20to%20ANU%20web%2025,12,11.pdf
- Lewis, David. 1973. *Counterfactuals*. Blackwell.
- Lewis, David. 1979. "Counterfactual Dependence and Time's Arrow", *Nous* 13: 455-476.
- Lewis, David. 1996. "Elusive Knowledge", *Australasian Journal of Philosophy* 74:549 – 567.

- Maass, Peter. 2009. *Crude World: The Violent Twilight of Oil*. New York: Vintage.
- McPherson, Tristram. 2016. "Why I am a Vegan", in Chignell, Cuneo, and Halteman eds. *Philosophy Comes to Dinner*. New York: Routledge.
- Moller, Dan. 2005. "The Morality of Overdetermination: Wrongdoing When it Makes No Difference", in *Abortion, Killing, and Overdetermination – Three Essays*, PhD dissertation, Princeton University.
- Mueller, Dennis. 2003. *Public Choice III*. Cambridge: Cambridge UP.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan", *Philosophy & Public Affairs* 39, 364-95.
- Nefsky, Julia. 2016. "Fairness, Participation and the Real Problem of Collective Harm", *Oxford Studies in Normative Ethics* 5, Oxford UP.
- Nefsky, Julia. 2018. "Consumer Choice and Collective Impact", in Barnhill, Budolfson, Doggett eds. *Oxford Handbook of Food Ethics*. Oxford UP.
- Norcross, Alastair. 2004. "Puppies, Pigs, and People", *Philosophical Perspectives* 18, 229-45.
- Nozick, Robert. 1981. *Philosophical Explanations*. Harvard UP.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford UP.
- Quong, Jonathan. 2009. "Killing in Self Defense", *Ethics* 119:507-537.
- Regan, Tom. 2004. *The Case For Animal Rights*, updated edition. Berkeley: University of California Press.
- Singer, Peter. 1980. "Utilitarianism and Vegetarianism", *Philosophy & Public Affairs* 9, 325-37.

Schwartz, David. 2010. *Consuming Choices: Ethics in a Global Consumer Age*. New York: Rowman & Littlefield.

Shepski, Lee. 2013. "Going the (Ethical) Distance", *Journal of Business Ethics* 116: 393.

Stalnaker, Robert. 1968. "A Theory of Conditionals", *American Philosophical Quarterly* 47: 98-112.

Thomson, Judith Jarvis. 1985. "The Trolley Problem", *The Yale Law Journal*, 94, 1395-1415.

Tooley, Michael. 1972. "Abortion and Infanticide", *Philosophy & Public Affairs* 2:37-65.

Wallace, Kathleen. forthcoming. "A Kantian Perspective on Individual Responsibility for Sustainability", in Schockley and Light eds. *Ethics in the Anthropocene*. MIT Press.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford UP.