

# A Defense of Torture: Separation of Cases, Ticking Time-bombs, and Moral Justification

**Fritz Allhoff**

**Western Michigan University**

**ABSTRACT:** In this paper, I argue for the permissibility of torture in idealized cases by application of separation of cases: if torture is permissible given any of the dominant moral theories (and if one of those is correct), then torture is permissible simpliciter and I can discharge the tricky business of trying to adjudicate among conflicting moral views. To be sure, torture is *not* permissible on all the dominant moral theories as at least Kantianism will prove especially recalcitrant to granting moral license of torture, even in idealized cases. Rather than let the Kantian derail my central argument, I directly argue against Kantianism (and other views with similar commitments) on the grounds that, if they cannot accommodate the intuitions in ticking time-bomb cases, they simply cannot be plausible moral views—these arguments come in both foundationalist and coherentist strains. Finally, I postulate that, even if this paper has dealt with idealized cases, it paves the way for the justification of torture in the real world by removing some candidate theories (e.g., Kantianism) and allowing others that both could and are likely to justify real-world torture.

## §1 SEPARATION OF CASES AND MORAL METHODOLOGY

Separation of cases (or, alternatively, constructive dilemma) is one of the most powerful methodological tools in philosophy. Put formally, it is the deductively valid argument form  $P \vee Q, P \rightarrow R, Q \rightarrow R, \therefore R$ . Put informally, the idea is that, if one of two propositions has to be true, and if each of those propositions entail some third proposition, then that third proposition has to be true. What makes it even more powerful is that the original propositions need not even be atomic, but could rather disjunctively relate even more propositions. By taking this latter approach, we can expand the scope of the argument to accommodate three, four, or many propositions to show that, so long as any of them is true (and insofar as

at least one of them must be), and all entail some other proposition, that other proposition must be true.

Since not all of us like to think in terms of formal logic, let us consider what I take to be one of the most powerful applications of separation of cases, which I attribute to Judith Jarvis Thomson's "A Defense of Abortion."<sup>1</sup> Admittedly, she does not speak in these terms, but this is irrelevant to the structure of her seminal paper. Before the paper, the abortion debate largely centered on whether fetuses were "persons" (roughly, "members of the moral community"): anti-abortionists alleged that fetuses were and should be afforded direct moral status and pro-choice advocates argued that they were not and should not be afforded that status. The debate was quite intractable insofar as the opposing sides disagreed on this fundamental issue. In Thomson's article, she offers, as a methodological concession, that fetuses are persons and then goes on to argue that *even if* they are, abortion is still morally permissible. Whether her arguments work is, of course, debatable, but the *structure* of her argumentation is profound insofar as it renders a highly controversial issue (*viz.*, the personhood of the fetus) *completely irrelevant*. The reason is that, *if* the fetus is not a person, then abortion can be defended by the pro-choice advocate (and perhaps conceded by the anti-abortionist) and, *if* the fetus is a person, then abortion is arguably permissible by her own argumentation. Obviously the fetus either is or is not a person. *Therefore*, abortion is permissible. Hence the constructive dilemma.

I am certainly not appealing to the Thomson paper in virtue of its substantive conclusion (*viz.*, that abortion is permissible), as there are obviously many things that we could object to along the way. Rather, I appeal to it to show the power of separation of cases, particularly as pertains to controversial antecedent claims. There are lots of seemingly intractable disagreements in philosophy, and many of those disagreements make the resolution of certain debates impossible. After all, if the parties to those debates simply have different starting points, nobody will ever be able to convince the other side. While a resolution of the disagreement on starting points would be desirable, it is probably too much to hope for. Nevertheless, separation of cases offers a powerful alternative: if certain conclusions can be shown to follow *regardless* of any plausible starting points, then we can be content in having secured those conclusions despite some agnosticism as to which starting point is the most appropriate.

With this methodological framework in place, the perceptive reader can probably see where I am going. There is an important substantive moral issue on the table: is torture morally permissible?<sup>2</sup> As we try to answer this question, we presumably must invoke some moral theory (perhaps coupled with empirical assumptions). There are lots of moral theories out there, and there are certainly no immediate prospects for determining which one is the right one. More skeptically, it is not even clear that there *is* a right one: moral pluralists and moral particularists have offered substantial arguments against the monism that has historically been so prevalent in moral philosophy. Nonetheless, all but the moral nihilist will admit that *something*, be it monism, pluralism, particularism, etc., has to be right. The possibility of nihilism can be marginalized in this case since, if nihilism is true, then, *a fortiori*, torture is not morally impermissible. If I merely wanted to argue

that torture is not morally impermissible, then I could easily accommodate the nihilist. However, I want to advocate a stronger position: that torture is, in some cases, morally required. This conclusion clearly cannot square with nihilism since the nihilist thinks that predicates such as “is morally required” are never instantiated. Since few of us are nihilists, I propose to ignore this view, though I do take it that the nihilist would (quickly) grant me the less ambitious conclusion.

Aside from nihilism, there are the aforementioned multitude of first-order normative views. Certainly I cannot pick which one of those is right. But, I do have separation of cases available to me. And, if I can show that *any* of those views will give me my desired conclusion, then the first-order debate is *irrelevant*. To be sure, there is a near-infinitude of ethical theories floating around these days, and I cannot discuss them all in this paper. But these ethical theories tend to be variations on central themes, and I can certainly discuss archetypical examples of those themes and hope to accommodate other (close) versions. As mentioned above, ethical theories can typically be classified in three different categories: ethical monism, ethical pluralism, and ethical particularism. Ethical monism is the most dominant tradition in moral philosophy and includes its most-recognized theories: consequentialism, deontology, virtue ethics, and social contract theory. While there are certainly other versions, consequentialism achieves its most auspicious presentation through John Stuart Mill’s hedonic act-utilitarianism.<sup>3</sup> Similarly, Immanuel Kant is usually flaunted as the flag bearer of deontology, though contemporary versions of deontology probably bear less similarity to Kant than contemporary versions of consequentialism do to Mill.<sup>4</sup> Plato’s virtue ethics might have been the first, but Aristotle’s has been the most lasting, though it again has contemporary faces.<sup>5</sup> Social contract theory similarly has a multitude of versions, and there are important differences among them, though we shall survey representative cases.<sup>6</sup> Moral pluralism and moral particularism have been more recent alternatives to the aforementioned monisms; they have been espoused by, for example, W. D. Ross and Jonathan Dancy, respectively.<sup>7</sup>

Returning to the notion of separation of cases, if *one* of these theories is right, and if all sanction the moral permissibility of torture (in some cases), then it follows that torture *is* morally permissible (in those cases). Unfortunately, things will not be quite as easy as this since one of our candidate moral theories will prove especially recalcitrant to endorsements of torture. As might be expected, this is going to be Kantian deontology. And this, of course, could throw a wrench in the whole project: if *one* of our candidate moral theories does *not* secure our conclusion, then separation of cases will fail and the conclusion will not be secured. Yet, I will not despair, and for two reasons. First, if *only* Kantians can object to torture, this really would be a substantive result. To wit, we will have now surveyed the entire landscape of moral philosophy and reduced the opposition to torture to a single (and probably lonely) outpost. But second, and more satisfyingly, I think that we can render the Kantian’s position implausible, thus removing it from the candidate moral theories. And, if we can eliminate its candidacy, then separation of cases is back in business and we can get to work securing our desired conclusion. But more on that later.

## §2 TICKING TIME-BOMB CASES: WHAT DO THEY SHOW?

By now, most of us are familiar with the ticking time-bomb cases, though I think that the implications and significance of these cases have been poorly understood. The conditions can be variously formulated, but they are usually rough homologies of the following: Imagine that a terrorist has placed a bomb in some public place such that, if detonated, many lives will be lost. Further imagine that law enforcement has just apprehended a terrorist and has epistemic certainty that: the terrorist is responsible for the bomb; the terrorist knows the details that could lead to its defusing; absent the terrorist's confession, there is no other way to defuse the bomb; if the information is provided, defusing is certain; absent torture, the terrorist will not divulge the relevant details; and, given torture, he will divulge those (and only those—this is relevant to prevent the possibility of misinformation) details. What should law enforcement do? Is torture permissible? Impermissible? Obligatory?

Presumably everyone thinks that torture is permitted in this case, perhaps excepting those pesky Kantians. However, we might still ask, *so what?* In other words, even if there is a (near) universal intuition that torture is permissible in *this* case, what could *it* possibly have to do with more general questions as to the permissibility of torture? To be sure, it is a highly abstracted and idealized case. As I have represented it, it is *so* idealized (particularly the stipulated epistemic certainty) that we might reasonably think that its conditions will never be met in the actual world. So what is the point in talking about it? What is its purpose?

First, let us talk about what its purpose is *not*, which is to provide necessary conditions for the implementation of torture. After being presented with this case, some people comment that the intuition that we should torture in this case is all well and good, but that real-world torture will never be permissible because these conditions will not be met. Similarly, I have heard that these cases are offering “requirements” for the application of torture (e.g., epistemic certainty, etc.) that will never be actually met.<sup>8</sup> But these comments clearly misunderstand the point of the case. The upshot of this case is supposed to be that it is permissible that we torture the terrorist. And, if it were to be interpreted as *merely* providing necessary conditions for torture (as opposed to ones that were jointly sufficient), then this conclusion would never logically follow.

Instead, the case should be read as providing *sufficient* conditions for the application of torture: *if* we have the terrorist in custody, *if* we know *P*, *if* we can  $\phi$ , etc. Whether all of these antecedents are ever instantiated in the real world is irrelevant to the permissibility in *this* case. And, since the conditions are meant to be sufficient, rather than necessary, we have *no information* (yet) as to whether torture is permissible in the real world when one or more of the conditions is not met. These claims just follow from the logic of conditionals. More formally, we could interpret this case as suggesting, for example,  $(P \wedge Q) \wedge R \dots \rightarrow S$ ;  $\neg P$ , has *no* implications for the truth of *S*. Less formally, this case “shows” that torture is permitted in highly idealized cases and is *completely* silent about whether torture would be permitted in other (less idealized) cases.

We now have two options: we can accept the intuition from this case as morally informative or else we can reject it as morally irrelevant. Let us pursue the first strategy for now and defer the second until later (*viz.*, §4). If we accept this intuition as morally informative, then we presumably must adopt a moral theory that can accommodate it. Hence the return to separation of cases. But, before we proceed, let us stop to ask an important methodological question: should moral theories be held hostage to intuitions from highly idealized cases? I think that the answer is yes, though some people will disagree, so this question merits some discussion. First, we could ask what prospective deficiencies this case might have for being morally informative. The critic will probably want to argue that it does not reflect the real world, is too fanciful, etc. Whether it *actually* reflects the real world is irrelevant, but we might concede a relevant concern is whether it *could* (i.e., counterfactually) represent the real world. The reason is that *merely* the fact that it is counterfactual cannot serve as evidence against it: any thought experiment involving, say, human reproductive cloning is “imaginary,” though we might reasonably expect some of them to be reified in the near future. So the critic is, I think, going to have to concede that actualization is morally irrelevant. Well, what about likelihood? Look, he might continue, maybe your cloning example will become actualized in a few years, but this ticking time-bomb case will *never* be actualized. Again, however, I fail to see how this objection can bear much critical weight. Insofar as the critic must already allow counterfactual considerations as morally relevant, we can push against the plausibility of this likelihood requirement. Strictly speaking, *every* state of affairs is infinitely unlikely since there are a finite number of ways that state of affairs could have been realized and an infinite number of ways in which it could have failed to have been realized. Maybe the critic now wants to regroup and say something to the effect of ticking time-bomb cases inhabiting *distant* possible worlds (because so many things would have to change in the actual world to instantiate a ticking time-bomb case) and then go on to say his moral theory must only accommodate close possible worlds. But here I fail to see how any non-arbitrary principle could be proffered to distinguish the close possible worlds he wants to accommodate from the distant ones he wants to ignore; other than mere stipulation there is no principled way to draw this line. Even if the issue here is one of metaphysical vagueness, the critic still needs to persuade us that this case is reasonably far away as to not count, and it is hard for me to see how this argument would proceed without being *ad hoc*.

I suspect that the critic is now finding me quite obnoxious, though these are meant to be seriously theoretical points. If the critic wants to liberate his theory from the ticking time-bomb case intuition (while accepting its legitimacy in that case), he is going to have to tell a story as to why *that* intuition can be cavalierly discarded. And, as I have argued, invocations neither of ‘imaginary’ nor ‘unlikely’ will be very profitable. While these are negative arguments against the discharging of the intuition, I think that positive arguments are on offer in favor of its retention. Moral *theory* is logically distinct from moral *practice* (even if the former is supposed to inform the latter). There are various merits that a theory might have: simplicity, parsimony, explanatory power, etc. And the accommodation of our intuitions (even in extreme cases) is an important asset a theory might have. Whether any healthy

person has ever actually faced prospective organ harvesting to save five ailing patients need not have any pragmatic implications, but it can certainly have theoretical import insofar as it gives deontologists a reason to reject consequentialism. Furthermore, nothing (pragmatically) important hangs on accepting the intuition in the ticking time-bomb case. Certainly this intuition alone is not going to give license to the so-called “stress and duress” tactics currently on offer at Guantánamo, so the moral theorist need not (necessarily) worry about accommodating it.

The acceptance of the intuition, however, is not without consequence. In the next section, I will return to the previously promised separation of cases and argue that the accommodation of this intuition commits nearly all moral theorists to the acceptance of torture (in at least some cases). Whether that acceptance will give rise to, say, Guantánamo-like practices is debatable, but then the debate will become an empirical one (i.e., the details will matter) as opposed to a theoretical one. I have not forgotten that some people will want to resist this ride and throw the intuition out of the window, but we will get to them soon enough.

### §3 EXAMINING THE THEORIES AND THEIR COMMITMENTS

Let us now look at several candidate moral theories and see where they would stand in regard to the moral permissibility of torture in ticking time-bomb cases. For now, let us consider: utilitarianism, non-Kantian deontology, virtue theory, social contract theory, ethical pluralism, and ethical particularism. These comments will necessarily be brief, but my aim is to show that any of these views could, in theory, license torture; if the reader is already willing to grant me this, he may jump to §4.

#### §3.1 Utilitarianism

These points are almost so obvious as not to need mention, but we can proceed quickly. Utilitarianism is a moral theory which holds that the right actions are the ones that maximize total aggregate happiness, so torture will be morally required if and only if it maximizes total aggregate happiness. We can effect this result through stipulation alone, as we have come pretty close to doing in the ticking time-bomb case. If the critic wants to argue that the case, as stated, will not *necessarily* maximize happiness, we could quickly adjust the case. To be sure, torture offers tremendous potential for disutility. Some obvious examples are: the effects on the terrorist, the effects on the interrogator, the potential that the practice of torture could increase the incidence of terrorism, the potential that populations would suffer psychological discomfort knowing that terrorism is practiced, and so on.<sup>9</sup> If these effects and potentials could become manifest as unhappiness, which they surely could, then this weakens the utilitarian case against torture. But these effects and potentialities are irrelevant to our purpose as we can merely stipulate our way around them. In other words, if we need to, we can simply adjust the case such as to limit these disutilities or else to have them eclipsed by greater and countervailing utilities. A critic might object that such adjustments are somehow constitutive of cheating, but this surely misses the point, which is that utilitarianism can *obviously* justify torture in at least some idealized cases.

This is the only conclusion that I am trying to establish at present, and I take it to be so straightforward that I will move on.<sup>10</sup>

### §3.2 Non-Kantian Deontology

While we will discuss Kant in §4, it would be a mistake to think that all deontologies will necessarily oppose torture. To be sure, Kant's will, but his deontology is only one species of the family of moral views, and other deontologies could surely have alternative commitments. Most fundamentally, deontology is an approach to moral philosophy which is predicated upon the moral primacy of rights and duties (etymologically, 'deontology' derives from the Greek '*deon*,' which means obligation). In this sense, we could even cast utilitarianism as a form of deontology since the former holds that we have a duty (or are obligated) to maximize happiness. However, such an assignment would clearly not cohere with the theoretical commitments of most deontologists, so we might not want to take it too seriously. Instead, let us consider the notions most fundamental to deontology: duties and rights. Now, imagine that these somehow come into conflict with each other. For example (and simplicity), imagine that A, B, and C all have a claim-right to their lives. Further imagine that A is getting ready to shoot B and C and that these executions can only be prevented if D shoots A. What should D do?

The deontologist can now say one of two things. On the one hand, she could say that D cannot shoot A because rights are inviolable or, to use some famous locutions, because rights are trumps<sup>11</sup> or offer side constraints.<sup>12</sup> On these lines, we may not violate some right *even if* such a violation will engender fewer overall rights violations. Alternatively, the deontologist might adopt some sort of aggregative approach which would hold that the right actions are the ones that either maximize or minimize whatever features she takes to be morally relevant (e.g., rights-preservations or rights-violations, respectively). For this latter sort of deontologist, D should shoot A: if he does, there is one rights-violation (*viz.*, A's) and, if he does not, there are two (*viz.*, C's and D's). Rights-violations are minimized through the shooting, thus making it morally required. (A similar story could be told for maximizing rights-preservations.)

Which is the more plausible view? I have to think that the aggregative approach makes a lot more sense. Ronald Dworkin has argued that in cases of rights conflict, we should look not at the explicit formulation of the right but rather to the values that suggested the right in the first place.<sup>13</sup> So, if individuals have a right to life, it is because life itself is something that is valuable and worth preserving. Given a conflict then, where the violation of one person's right (to life, let's say), could prevent the violation of five other person's right to life, the values that suggested the right to life would suggest violating the one in order to prevent violation of the five.<sup>14</sup> By considering why we would endorse rights in the first place (because we value the objects of those rights), it seems permissible to act such that the underlying values (and their associative objects) are preserved to the highest degree possible. This approach is sometimes (mis)labeled "rights-based utilitarianism,"<sup>15</sup> and has some contemporary defenders.<sup>16</sup>

Nevertheless, some people have objected to it, most notably Robert Nozick. As I implied above, someone who really took rights seriously might think that

some aggregation procedure is more attractive than rights fetishism, but Nozick does have a concern some might recognize as legitimate; he asks: “why . . . hold that some persons have to bear some costs that benefit other persons more, for the sake of the overall social good?”<sup>17</sup> Instead of the previous example, imagine the following: A is preparing to execute B and C but will refrain if B executes D (where D is some innocent third-party). The aggregative deontologist would presumably have to say that B should execute D since then we only have one rights violation (*viz.*, D’s) instead of two (*viz.*, B’s and C’s). If we agree with Nozick that this commitment is problematic, then that would be a strike against the theory.

However, there are at least two available responses. First, in a footnote, Nozick expresses concern with the application of his theory to “cases of catastrophic moral horror,”<sup>18</sup> which are precisely what the ticking time-bomb case is trying to model. So it is not immediately clear that *even Nozick* would push his “rights as side constraints” conception into the current arena. Second, Nozick’s objection to aggregative deontology is effectively that innocent third parties might absorb all the burdens such that unrelated parties derive all the benefits. But this concern is obviously rendered moot in the ticking time-bomb cases since the terrorist is, *ex hypothesi*, not an innocent third party: he planted the bomb and is sitting on the information that could be used to disarm it.

So, now let me try to tie together various points that I have tried to make in this section. First, I proposed that deontology could diverge from Kant’s formulation of it; the deontologist is solely committed to the moral priority of rights and duties and need not endorse Kant’s version. Second, I proposed that, when rights come into conflict, the deontologist will have to make a decision: she can either adopt an aggregative approach or else she can reject such an approach. I then argued that the principles that would suggest the moral significance of rights would favor the aggregative approach. In regards to Nozick’s potential criticisms of this approach (*i.e.*, if he wanted to maintain them in the face of catastrophic moral horror), I argued that they could be allayed in the time-bomb case under discussion.

While the conclusion should be looming, allow me to make it explicit. In the ticking time-bomb case, we might reasonably expect torture to minimize overall rights violations. As in §3.1, we might adjust the details of the example to guarantee this result, but I take it to already be implied. If the deontologist is of the aggregative sort, this means that torture will be morally required on her view. If the deontologist is of the non-aggregative sort because of Nozick-like worries, torture will still be required since those worries do not apply in this case. Either way, torture is morally permissible.

The critic might object to the argumentation of this section in two ways. First, she could claim that I have considered a peculiar version of deontology and not the one that she wanted. To be sure, other versions of deontology (*viz.*, Kant’s) will *not* give the conclusion that I have drawn, though we will return to them later (see §4). But I claim that I have been fair: I have tried to abstract a *general deontology* from many of the extant versions, and I have tried to show what commitments this general form will have. *Obviously* not every deontology will have them, particularly insofar as some deontology could even be created *ad hoc* around



the duty not to torture! I do, however, hope to have addressed the general spirit of this family of views, if not all the particular versions. Second, the critic might allege that there are criticisms against the aggregative view other than Nozick's which would not have been defused through my argumentation. However, I am not sure what those other criticisms might be, even though I have tried to think of them. I certainly invite further challenge on this matter.

### §3.3 Virtue Theory

Virtue theory is a tricky one to deal with. One reason is that it is not even directly concerned with how people should *act*, but rather with how people should *be*. Of course the virtuous person will tend to act virtuously, but the focus of the theory is on the character traits that would lead to the virtuous acts rather than on the acts themselves. For these reasons, virtue theory has been criticized as not being "action guiding" because, unlike consequentialists or deontologists, the virtue theorist cannot offer a principle of right action. Or, if he could, it would be so formal as to be of very limited pragmatic use. For example, the virtue theorist might say that we should "act as the virtuous person would" or, more simply, that we should "act virtuously." And surely, as the theory goes, this is the right answer. But this is not terribly *helpful* insofar as such a principle offers far less direction than Mill's Greatest Happiness Principle or Kant's Categorical Imperative (particularly for those of us without access to a virtuous person). Of course, the virtue theorist will want to say that the focus of moral theory (on action as opposed to character) is inappropriate. Maybe this is true and maybe it is not, but it is hard to see how that suggestion is supposed to be very useful in moral practice.

Regardless, we shall try to do the best that we can. In deciding whether to torture the terrorist in the ticking time-bomb case, we must solve the riddle of what the virtuous person would do. Since this is a thought experiment and we cannot see what a virtuous person actually would do, we must imagine what a virtuous person *would* do (were one to exist and to be confronted with this dilemma). Unfortunately, we cannot be much more scientific than imagining the best person that we know (or, even better, that we can conceive of) and plugging him or her into the thought experiment, though this is all that the theory offers us. What would the virtuous person do? Would he torture or not?

First, let us think about what the virtues are. Plato postulated wisdom, courage, temperance, and justice.<sup>19</sup> Aristotle kept these (albeit with slight taxonomic emendations) and added generosity, pride, good temper, truthfulness, modesty, wittiness, friendliness, modesty, and righteous indignation.<sup>20</sup> Which are relevant to our case? Or, in other words, if the virtuous person has all of these character traits, which will influence his decision whether to torture? Probably more than one will come into play. For example, torturing someone is not very nice and is therefore demonstrative of the vice of unpleasantness (which opposes friendliness). Conversely, letting many people die preventable deaths is not very friendly. If the virtuous person wanted to *maximize* his display of friendliness, should he not torture? Similar conflicts probably exist with other virtues.

Maybe the most obvious virtue to invoke is justice. Which conception should we use? Plato proposes that justice is a harmony of the tripartite soul, which is

to be governed by the rational element. Aristotle invents a taxonomy of justices, differentiating the general from the particular and subdividing the particular into distributive justice and rectificatory justice. I do not see how either Plato's or Aristotle's conceptions are of much use in settling the question before us. If we jump ahead a couple of millennia, there might be some promise in the Scottish Enlightenment. Francis Hutcheson, for example, thought that justice consisted in being motivated by universal benevolence.<sup>21</sup> If this is right, we might expect the just person to perform torture (for reasons similar to those expressed in §3.1 above). David Hume criticizes Hutcheson's account and provides his own, which is roughly (and perhaps controversially) that justice is an artificial virtue ultimately motivated by self-interest.<sup>22</sup> On Hume's conception, we could refine the example such that the virtuous person was one of the people who would perish in the bomb detonation (e.g., we could stipulate that the bomb was at an unknown location within the building and the terrorist was willing to become a martyr). Then torture would be in his self-interest, so justice would require that he perform it.

But look, maybe we are making this too hard by trying to determine which conception of justice is the proper one here (or even by assuming that justice is necessarily the "dominant relevant virtue"—whatever that would mean). Rather, let us retreat to the original question: what would the virtuous person do? Let us take this hypothetical individual, unencumbered with potentially idiosyncratic conceptions of the virtues, and ask whether he would perform torture. If this person does not perform torture, many people will die. And the torture of the terrorist will alleviate all of these deaths. We can even further stipulate that the necessary torture will be comparatively minor. I do not see any plausible line that could hold that our virtuous person would not engage in this torturing. To be sure, he will be evincing at least some vice through the torture; virtue ethicists sometimes refer to such situations as "tragic dilemmas" (i.e., ones in which vicious behavior will result *regardless* of what the virtuous person chooses).<sup>23</sup> We might reasonably suppose that, in such cases, the virtuous person should pursue the least vicious path available, and we might further reasonably suppose that it would be less vicious to perform a single act of torture than to allow many deaths.

To be sure, this conclusion follows from a high generic virtue ethic, but I take that to be part of its strength. In other words, because it is not following from any idiosyncratic conceptions of individual virtues, I think that it has more power. So I think that this is the right conclusion to draw: the virtue ethicist should allow that the virtuous person should torture in the ticking time-bomb case, though we can perhaps acknowledge that such a case is a tragic dilemma and complete virtue cannot be displayed regardless of what the virtuous person does.

### §3.4 Social Contract Theory

Social contract theory has perhaps gained more adherents in recent years and maybe even now rivals the traditional three normative theories (*viz.*, consequentialism, deontology, and virtue ethics). Despite the recent appeal of the works of John Rawls, David Gauthier, Tim Scanlon, etc., social contract theory has a rich tradition with roots in Thomas Hobbes, John Locke, David Hume, Jean-Jacques Rousseau, etc. What stance would social contract theorists take on torture? To be

sure, it would depend on the version of the theory that one adopted, and there are substantial differences among the different views.

Nonetheless, social contracts all share some set of central features, of which the most notable might be: "The idea that morality is deeply implicated in the very notion of agreement, and vice versa, so that whether an action is right or wrong must depend on whether the act accords with or violates principles that are, or would be, the object of suitable agreement between equals."<sup>24</sup> Following a distinction proposed by Darwall, we might broadly categorize social contracts into two broad families, contractarianism and contractualism, depending on how they understand the notion of 'equality': for the contractarian, the equality of the contracting parties is "merely *de facto* and their choice of principles rationally self-interested"; this can be contrasted with contractualism which proceeds from "an ideal of *reasonable* reciprocity or fairness between *moral equals*."<sup>25</sup> Hobbes offers the classical statement of the contractarians notion of *de facto* equality:

Nature hath made men so equal in the faculties of body and mind as that, though there be found one man sometimes manifestly stronger in body or of quicker mind than another, yet when all is reckoned together the difference between man and man is not so considerable as that one man can thereupon claim to himself any benefit to which another may not pretend as well as he. For as to the strength of body, the weakest has strength enough to kill the strongest, either by secret machination or by confederacy with others that are in the same danger with himself.<sup>26</sup>

Alternately, contractualists postulate some moral norms which exist antecedent to the contract; these often have to do with mutual respect, justice, fairness, etc., and Hobbes explicitly denies such values in his State of Nature.<sup>27</sup>

Since these two families of views have such different commitments, we should look at them individually. Start with contractarianism, in which the parties to the contract will be motivated by nothing other than self-interest and there are no antecedent restrictions in their pursuit of this goal (save for similar motivations by everyone else). On this scheme, there are various ways that we might approach the torture issue. First, I know that I am not a terrorist, and I know that I might stand to suffer by the terrorist's bomb if torture were not exacted upon him. Therefore, it would be in my self-interest to torture the terrorist in the ticking time-bomb case, and I would push for such provisions in the covenant. Since, as the example has been stipulated, there are more would-be sufferers from the terrorist's bomb (*viz.*, lots) than there are terrorists (*viz.*, one), the popular will would certainly condone torture in these idealized cases. Even if, at the time of legislation, I did not *know* that I would not become a terrorist, any sort of reasonable probabilistic projection would sanction my endorsement of torture in the cases under consideration.<sup>29</sup>

This line of reasoning would not be popular among everyone, particularly the contractualists. For one, it permits tyranny of the majority, which contractualists might deem *unfair* (a word lacking in the pre-contract vocabulary of the contractarian). For another, even if some contractarianisms lacked this feature, the contractualist might nevertheless object to the dominance that self-interest plays in the theory, whether by imparting some other normative end or else placing some restriction upon pursuit of this one. A standard contractualist move is to introduce some further motivation (e.g., justification of one's reasons to others; cf. Scanlon) or

else some further construct, such as Rawls's veil of ignorance and its subsequent principles of justice. Given the popularity of Rawls's version, let us focus on it as a plausible contractualist view. The veil of ignorance deprives the contracting parties of any idiosyncratic features about themselves that they might use to differentially bias the outcomes of the bargaining process in their favor. For example, if I did not know whether I were male or female, Rawls thinks that I would be unwilling to legislate in a way that might benefit one group at the expense of the other. The reason, according to Rawls, is the "maximin rule" by which I should try to maximize the position of the worst well-off because I would not want to occupy that position.<sup>30</sup> From behind the veil of ignorance, Rawls thinks that we will adopt two principles of justice which will restrict all other substantive legislations we might make for our society. The first principle is that "each person is to have an equal right to the most extensive scheme of basic liberties compatible with a similar scheme of liberties for others" and the second is that "social and economic inequalities are to be arranged so that they are both (a) reasonably expected to be to everyone's advantage, and (b) attached to positions and offices open to all."<sup>30</sup>

It seems to me that there is no reason, in principle, that torture would necessarily be impermissible on Rawls's view. Remember that these rules are ordered, which is to say that applications of the second principle are restricted by the first principle. Therefore, we need not worry about the difference principle (*viz.*, 2a) if torture is mandated by the first principle. As proposed in §3.2, there will be times when liberties will come into conflict. In the torture cases, someone is going to be deprived of his liberty not to be tortured or else a lot of people are going to die, thus being deprived of various liberties they might have otherwise had. Rawls asks us to adopt the most extensive scheme of basic liberties consistent with similar extension to all, but we must now choose which liberties are going to be prioritized: the liberty against being tortured or the liberty to life (and whatever other liberties this liberty preserves). Given the conflict, we cannot grant all of these liberties to everyone, so we must make choices. Obviously Rawls rejects utilitarianism,<sup>31</sup> but he clearly wants to maximize the number of liberties afforded to the members of society (consistent with the provision of similar liberties to everyone). Certainly I think that it is preferable to extend liberties to continued life (and whatever corollary liberties that liberty would support) to all of society at the expense of the withdrawal of the terrorist's liberty from torture. This strategy need not even be motivated by some sort of rights-utilitarianism (e.g., that we should violate the rights of the terrorist so that fewer rights overall are violated), but rather from a conception in which the liberty to life is, to use Rawls's word, more *extensive* than the liberty from torture. In preserving the liberty to life, we certainly preserve a tremendous number of other liberties, whereas the abrogation of the liberty from torture is more "localized" insofar as it does not deprive the victim of *all* other liberties that s/he might have otherwise had; death is surely maximally deprivational on any liberty calculus. Prioritizing life and its associative liberties need not even violate Rawls's concerns about respecting the "separateness of persons" since we could view this decision not as a sacrifice of the terrorist to the collective, but rather as a conceptual point about what our commitment to maximally extensive liberties requires.

So I think that torture can be justified even on (some) contractualist grounds. Without even looking at any particular contractualist view (e.g., Rawls's), the point could be made more generally that, whatever pre-contractual norms exist, they could invoke terms (e.g., equality, respect, fairness, etc.) that could plausibly be interpreted to require us to save the many over the one given the necessity of such a choice. If so, then we can accommodate contractualism as compatible with the permissibility of torture (in idealized cases). Alternatively, if some versions of contractualism are incompatible with the permissibility of torture (which, admittedly, they might be), then we can extend the arguments of §4 against them.

### §3.5 Moral Pluralism and Moral Particularism

In the past century or so, an increasing number of moral philosophers have objected to monistic theories on the grounds that the ones "thus far proposed [are] all unbearably crude, and [these philosophers] contend that any attempt to reduce morality to a single principle will inevitably leave something out."<sup>32</sup> The alternatives to monism can come in either of two flavors, depending how deep the criticism runs: moral pluralism or moral particularism. For the pluralist, there are an irreducible number of moral concepts that are relevant, though these can perhaps somehow be aggregated into some principle or, at a minimum, they can be identified. For the particularist, morality is irreducibly context-dependent and there are *no* principles (whether monistic or pluralistic) that are finite, substantive, and exceptionless.

The most famous version of pluralism is that of W. D. Ross, who proposed a list of seven *prima facie* duties: duties of fidelity (e.g., promise-keeping, including honesty) and reparation, duties of justice and gratitude, duties of beneficence and self-improvement, and a duty of non-injury.<sup>33</sup> For Ross, morality cannot be reduced to, for example, Mill's happiness or Kant's categorical imperative, yet these theories certainly identify morally relevant features which should be afforded moral weight. Ross thought that we use some intuitive faculty to access appropriate moral reactions to complex cases, though that feature of his account is irrelevant: all that matters, for our purposes, is the idea that there are multiple morally valuable inputs.

Since utilitarianism, (non-Kantian) deontology, and virtue theory could all license the use of torture (see §3.1, §3.2, and §3.3, respectively), then I submit that any pluralistic moral theory which depends on the moral inputs identified by these theories will also find torture morally permissible *a fortiori* (in the same cases). A pluralism would not, of course, have to hold that all and only those deemed morally relevant by those theories were to be preserved: the pluralist could reject some of those features as morally relevant (in which case torture would still be permissible) or he could add new ones (in which case it might not). On this latter path, the permissibility of torture would depend on the details of the theory, though I doubt that many pluralists would espouse moral features that any of the dominant theories in moral theory have failed to accommodate. If so, we would have to evaluate the theory's commitments individually, but it is certainly reasonable to suspect that most pluralisms would be able to find torture morally permissible, at least in our highly idealized cases.

The particularist challenge to monism is even more profound since she even denies that a list of moral duties can be given *at all*; she emphasizes the role that judgment plays in moral evaluation, how judgment dictates which features are morally relevant, and how those features interact with other morally relevant features. For example, the particularist might think that the “principles appealed to by Rossian pluralists, will inevitably lead to error: the features that make something good are just so complicated, the conditions under which an action is right just so variegated, and the properties that make a person virtuous just so nuanced, that the moral realm resists capture by *any* number of finite principles.”<sup>34</sup> Particularism has been espoused by Jonathan Dancy, and has other contemporary defenders.<sup>35</sup>

For the particularist to decide whether torture is morally permissible, it would obviously depend upon the details of the case; there is no way that we can state the details of the case so that we are inextricably given the conclusion that torture is permissible since to do so would be to presuppose the satisfaction of some sort of theoretical criteria which the particularist would eschew. Nevertheless, I think that we can safely depend upon the particularist’s emphasis on judgment or attendance to the salient features of a case to establish that torture could be permissible in idealized cases. Or perhaps it is easier to show that the alternative cannot be ruled out: since the particularist denies principle and emphasizes cases, there is no way that her theory could *necessarily* rule out torture in the way that Kant’s might (by, for example, identifying some necessary feature of torture as inconsistent with his requirements for moral permissibility). If it is not the case that, for the particularist, torture is necessarily impermissible, then torture is possibly permissible; this follows from modal logic. And then, of course, it would come down to the details of the case, which we could stipulate as needed to get a particularist commitment to the permissibility of torture. For our purposes, all we need is that possibility, so we can move on.

#### §4 KANTIANISM AND MORAL JUSTIFICATION

We now come to the promised discussion of the Kantian stance on torture and, perhaps to the surprise of some readers, I will *not* argue that Kantianism can be construed in some way by which torture is morally permissible. Rather, I fully grant that, given a Kantian framework, torture is most definitely impermissible. However, I will use this result to argue *against* the plausibility of Kantianism, rather than to suggest that we have some reason to question the permissibility of torture.

First, consider why the Kantian cannot sanction torture. We can get the impermissibly by any of the formulations of the categorical imperative, but the second one (the Humanity Form) seems to be a popular one in this regard, so we can start with it. According to this formulation, we must always treat humanity, whether ourselves or others, as ends and never merely as means. Torture fairly clearly violates this edict since, through torture, we are treating the terrorist as a means to some end (*viz.*, the diffusion of the terrorist threat) and, furthermore, to some end which he does not share (*i.e.*, we are treating the terrorist *merely* as a means).

Other commentators have gone on to fill out the details of the Kantian opposition to torture,<sup>36</sup> though the previous sketch suffices for my needs. Nevertheless, I am willing to grant many other points that they might want to make. For example, some hold that torture is one of the *worst* moral harms since, in addition to the suffering that it inflicts upon the tortured (which the utilitarian can accommodate as a moral wrong), torture turns the terrorist against himself insofar as it forces him to derail projects (e.g., terrorism) that are deeply important to him; not only *are* his projects being derailed, but he is being forced to play an instrumental role in their destruction. In this sense, torture is a greater moral harm than even murder since even though the terrorist would be unable to personally realize his ends in either case, he is being forced to play an active role in the compromise of his ends in the torture case.<sup>37</sup>

These points are all well and good, and I am even inclined to think that most of them have some purchase. However, I think that, ultimately, none of them is that relevant. So, moving forward, let us *grant* these various points that the Kantian might want to make, including that torture is an especially bad, or even among the worst, of all moral wrongs. Certainly I do not think that torture is *not bad*: it obviously is. Granted, my intuitions tend to be fairly utilitarian, but even the utilitarian is going to say that there is something wrong with torture, though he will go on to say that wrongness can be mitigated by some good that will be produced. The Kantian denies the legitimacy of the utilitarian's maximization/ aggregative approach, but surely both moral theories (or even all moral theories) will acknowledge the *prima facie* wrongness of torture. So the question, which I think ultimately pits Kantianism and other absolutist deontologies against the rest of moral theory, is the following: is it morally permissible to inflict a *prima facie* moral wrong *in order to* bring about a greater moral good (or, conversely, to prevent a greater moral harm)?<sup>38</sup> It seems to me that the answer is obviously yes, and it seems to the Kantian that the answer is obviously no. What do we do now?

One (ambitious) approach would be to try to adjudicate among the conflicting positions. Since many moral philosophers view this debate as unresolved (or even worse, unresolvable), I doubt I will be able to convince everyone, though I will advance an argument toward this end. Ultimately, I hope to render the Kantian position implausible and, as goes it, so goes the opposition to torture. So, for now, let us assume that at least some moral theory is justified (by which I mean that we have substantial warrant or evidence to believe in its truth). There are two traditional modes from whence this justification would derive: foundationalism and coherentism.<sup>39</sup> What I want to imply here is that, given our commitments in the ticking time-bomb case, Kantianism cannot be the right moral theory. So let's see how this argument would go.

First, I am *assuming* that torture is permissible in the ticking time-bomb case. Since this is a thought experiment, I can stipulate it however I wish; in addition to the features explicated in §2, let me reiterate the features that *lots* of people stand to die (e.g., millions, billions, etc.) and that the necessary torture to elicit the preventative information could be comparatively minor. If anyone wants to disagree with the permissibility of torture in this case, I simply do not know what to do other than throw my hands up in exasperation. Whether such a case

would ever attain is completely irrelevant to the following argumentation. All that matters, for present purposes, is the admission that *if* it were to happen, then torture would be morally permissible; the endorsement of this conditional in no way commits anyone to the acceptance of its antecedent. So, henceforth, I will assume that torture is permissible in the ticking time-bomb case. If someone disagrees with this, my arguments will be ineffectual, and the most I can say in response is that I find their intuitions (or, as likely, “intuitions” as they are probably post-theoretical) to be hopelessly implausible.

For the reader who was willing to accept the argumentation in §3, we have now reduced the opposition of torture to Kantianism (or a cluster of related views, perhaps including some contractualisms or non-aggregative deontologies; henceforth, I will focus on Kantianism, but my arguments will equally apply to any other view with similar commitments). This opposition, however, will only be as strong as the normative view on which it depends, and I will now argue against the plausibility of any normative view with such commitments. And here is where the power of the ticking time-bomb case, hinted at in §2, will (hopefully) become apparent, as I take the implications of that case (*viz.*, the permissibility to torture) to be sufficient to derail entire moral theories. So, now let’s turn to moral justification.

Start with the foundationalist: this is the view that at least one moral claim is justified non-inferentially, as well as the likely corollary that some other moral claim(s) will be justified inferentially from the foundational one(s). There are various ways by which the foundationalist might proceed, but a common one is some variant of ethical intuitionism by which we ascertain some foundational moral claim by intuition alone; from this moral claim, we can derive others.<sup>40</sup> Foundationalism could be “top-down” or “bottom-up”: our foundations could be general moral principles that allow us to deductively derive the morality of individual cases or else our foundations could regard the morality of individual cases from which we could inductively derive general moral principles. Either way, there will be some intuitions that are foundational and which form the basis of all further moral knowledge. While there is some debate as to what features foundational intuitions will have, a standard characterization is that they are “self-evident” (i.e., something like that they are recognized as true once they are cognitively apprehended); this concept certainly warrants further discussion, but I shall not address that here.<sup>41</sup>

To my mind, top-down moral foundationalism is implausible for the reason that no general moral principle is self-evidently true: given the pluralism of intuitions regarding first-order normative views, I think that none is self-evidently true for, if one were, such pluralism would not exist. Alternatively, bottom-up foundationalism strikes me as a lot more viable. For this to work, we would have to have moral knowledge in some instances, and we would then use our knowledge in these instances to inductively choose a moral theory that could accommodate them. If the moral theory cannot accommodate them, then it cannot be (fully) correct. This follows from logic alone: if the theory implies some commitment, and if that commitment is inconsistent with the foundations of our moral system, then the theory must be abandoned or revised.<sup>42</sup>



My next move should be no surprise: I take the moral permissibility of torture in the ticking time-bomb case to be self-evident; anyone who understands the details of the case would, I think, consent to the torture. If Kantianism is true, then the torture would be impermissible. But the torture is *not* impermissible. Therefore, Kantianism is false. Therefore we can remove it from the list of our candidate moral theories. Anyone who wishes to avoid the conclusion of this argument must deny one of its premises (since the argument is valid by *modus tollens*), and I do not see either premise as being challengeable: I have already argued that Kantianism is committed to opposing torture, and I really fail to see how any reasonable person could deny the permissibility of torture in the ticking time-bomb case.

And here is where the ticking time-bomb case, commonly ridiculed as being irrelevant to real-world situations, demonstrates its power: if we are moral foundationalists and accept the permissibility of torture in the ticking time-bomb case, then Kantianism has to be wrong as it cannot accommodate this conclusion. Again, this argument does not license the so-called “stress and duress” tactics at Guantánamo, but it renders the Kantian obstructionist irrelevant since that view is inconsistent with a self-evident moral truth. It seems to me the critic might now say one of two things. First, she might say that she does not take the permissibility of torture in the ticking time-bomb case to be self-evident; against this critic I simply do not know what to say other than that I suspect she is being intellectually disingenuous. Second, she might say that she is not a foundationalist and that the previous arguments are irrelevant. And, of course, she does not have to be a foundationalist since there are other views of justification available. For example, there is coherentism. So let us look at that one.

The coherentist thinks something along the lines that someone’s “moral beliefs are epistemically justified if, and then to extent which, they cohere with the other things she believes.”<sup>43</sup> This view has achieved recent prominence through Rawls’s method of reflective equilibrium, which comes in two flavors: narrow and wide reflective equilibrium. Narrow equilibrium occurs when one has a set of moral principles which coheres with moral judgments that, on reflection, one is willing to endorse and wide equilibrium occurs when these moral principles and judgments cohere with other relevant and non-moral principles (e.g., sociological, psychological, physical, metaphysical, etc.).<sup>44</sup> For our purposes, narrow reflective equilibrium will suffice. While the analogy with science should not be over-emphasized, some general similarities between the moral methodology and Rawls’s narrow reflective equilibrium are apparent.<sup>45</sup> Just as, in science, we can test hypotheses against experimental data, our moral methodology allows us to test our general principles against considered judgments in individual cases. If our scientific hypothesis is inconsistent with the observed data, we can do one of two things: reject the hypothesis as inadequate or else reject the data as non-representative. Similarly in our moral methodology, if our principle fails to cohere with our considered judgments, one of them has to go, and it can be either the principle or the considered judgment; principle holds no primacy over considered judgment. Which one goes depends on which we are more committed to, have more confidence in, has more coherence with other principles/judgments, etc.

So let us apply the method to the issue currently under discussion. We have a principle (*viz.*, Kantianism) and a considered judgment (*viz.*, the permissibility of torture in the ticking time-bomb case), and these two do not cohere. We could seek a wide reflective equilibrium by bringing in non-moral considerations, or we could bring other moral considerations to bear on our narrow reflective equilibrium. But, for simplicity, let us pretend that this principle and considered judgment are all that lie before us and that we must somehow dissolve the tension between them by rejecting one or the other. Which do we reject? Since, for the sake of simplicity, we are not positing other principles and/or judgments which might need to be equilibrated, we merely have to think about which we are more committed to or have more confidence in: Kantianism or the permissibility to torture in the ticking time-bomb case. Which is it?

Speaking for myself, I tend to be a moral agnostic, by which I mean that I suspend judgment on which first-order normative theory I would otherwise endorse: they all have their merits and demerits and the adjudication among the conflicting theories is, in my opinion, so confused and unconvincing that agnosticism seems to me the only appropriate stance (if not skepticism that there *is* a correct theory). While I certainly do not expect all moral philosophers to adopt such a view, I suspect that most would be willing to admit that reasonable people can hold alternative views from the ones that they themselves hold and that there are reasonable arguments for those positions as well as reasonable arguments against their own. In other words, there is at least the possibility of debate as to which moral theory is the correct one. Alternatively, I completely deny that there is even the *possibility* of debate as to whether we should torture in the ticking time-bomb case, particularly given the liberties that we can take with the stipulations. Granting these points, we would necessarily have more confidence in *this* considered judgment than we would in *any* moral theory (including Kantianism). Therefore, our moral methodology should preserve the considered judgment as against whatever principle with which it conflicts, meaning that Kantianism is not the correct moral theory, nor any other that would oppose this considered judgment. This argument does not, of course, claim to determine which normative theory is the *correct* one, but rather which ones are *incorrect*. There is a long list of moral theories discussed in §3, and I think that all of them are plausible. However, through the method of reflective equilibrium, I deny the plausibility of any normative theory which cannot accommodate our considered judgment in the ticking time-bomb case, and Kantianism is one such casualty.

## §5 CONCLUDING REMARKS

Before concluding this paper, I would like to make a couple of remarks about the pragmatic implications of the arguments therein. Many of those who discuss the morality of torture are concerned with whether torture would ever be permissible in the real world, even if it could be philosophically justified in principle. As a moral philosopher, the question of real-world justifiability is really not one on which I have much to say; I am more interested in the theoretical arguments that can be rendered on either side of this debate. This is not, of course,

to say, that I do not care whether unjustifiable torture takes place in the world. To the contrary, I care very much, but I do not find myself equipped to make a contribution to that debate apart from being able to offer some sort of theoretical framework which would then require empirical inputs.

Nevertheless, I do not share the skepticism of some who think that torture will never be justified. For example, Michael Davis argues that those who depend on empirical claims for the moral permissibility of torture (e.g., that is maximizes happiness) “never show that [those] claims are true or even probable.”<sup>46</sup> Admittedly, I have no idea whether such empirical claims would be true (or even probable), but this is completely irrelevant for what I want to argue, namely that *if* they are true, then torture is morally permissible. Davis’s appeal to “practical moral absoluteness” (i.e., the idea it might be morally justifiable to violate absolutes in principle if not in practice) circumvents a lot of the important questions of moral philosophy, but I even disagree with his pragmatic pessimism and suggest to replace it with agnosticism.

Admittedly, this has been an ambitious paper which has sought to accomplish a lot and, despite its length, is probably underdeveloped in some places and overextended in others. Nevertheless, its central argumentative strategy should be readily apparent. In §1, I proposed to proceed by separation of cases wherein I would consider multiple moral theories and attempt to show that they all led to the same conclusion; despite agnosticism as to the appropriate moral theory, the conclusion would nevertheless be secure. In §2, I discussed the ticking time-bomb case that would form the basis for the considerations of these moral theories and, in §3, I showed how these theories would handle this case. In §4, I considered the one dominant moral theory, Kantianism, that would prove recalcitrant to accommodating my desired conclusion, and used this failure of accommodation, coupled with discussions of moral justification, to argue against the plausibility of the theory. Once Kantianism (and its ilk) are removed from the list of candidate moral theories, the application of separation of cases can be revisited and the desired conclusion, namely that torture is morally permissible in some cases, can be secured. The extent and characteristics of those cases will depend on which normative theory we adopt, but I suspect that there will be real-world instances wherein torture is morally justifiable and, at a minimum, I hope to have made a contribution to the theoretical project.

## Endnotes

In the writing of this paper, I benefited from discussions with and comments from several people; I would like to give special thanks to Chris Buford, Michael Davis, Jeanette Kennett, Seumas Miller, Thomas Pogge, and Anand Vaidya.

1. Judith Jarvis Thomson, “A Defense of Abortion,” *Philosophy and Public Affairs* 1.1 (1971): 47–66.

2. For the purposes of this paper, I shall not be intimately concerned with the nature or definition of ‘torture’; I propose to leave this concept unanalyzed and propose that my arguments will accommodate any plausible conceptions. If the reader is interested in this debate, I suggest, in this volume, Michael Davis’s “The Moral Justifiability of Torture and other Cruel, Inhuman, or Degrading Treatment” (esp. §1) and Seamus Miller’s “Is Torture Morally Justifiable” (again, esp. §1).

3. John Stuart Mill, *Utilitarianism* (Indianapolis, IN: Hackett Publishing Co., [1861] 1979). For more classical and contemporary discussion, see *Consequentialism*, ed. Stephen Darwall (Malden, MA: Blackwell Publishing, 2002).

4. Immanuel Kant, *Groundwork of the Metaphysics of Morals*, trans. H. J. Paton (New York: Harper & Row, 1964). For more classical and contemporary discussion, *Deontology*, ed. Stephen Darwall (Malden, MA: Blackwell Publishers, 2002).

5. Plato's virtue ethics appears throughout his writing, see, for example: *Republic*, trans. G. M. A. Grube, rev. C. D. C. Reeve (Indianapolis, IN: Hackett Publishing Co., 1992a), esp. 327a1–480a12; "Laches" in *Laches and Charmides*, trans. Rosamond Kent Sprague (Indianapolis, IN: Hackett Publishing Co., 1992b), 178a1–201c5; "Euthyphro" in *Five Dialogues*, trans. G. M. A. Grube (Indianapolis, IN: Hackett Publishing Co., 1981), 2a1–16a5. For Aristotle's virtue ethics, see his *Nicomachean Ethics*, trans. Martin Ostwald (Englewood Cliffs, NJ: Prentice Hall, 1962), esp. 1094a1–1109b28 and 1119b20–1138b14. The strongest contemporary formulation of virtue ethics is Rosalind Hursthouse's *On Virtue Ethics* (New York: Oxford University Press, 1999).

6. Social contracts received early formulations from Thomas Hobbes, John Locke, David Hume, Jean-Jacques Rousseau, etc., and more contemporary formulations from John Rawls, David Gauthier, T. M. Scanlon, etc. Most of these views are presented in *Contractarianism/Contractualism*, ed. Stephen Darwall (Malden, MA: Blackwell Publishing, 2003).

7. See W. D. Ross, *The Right and the Good* (Indianapolis, IN: Hackett, [1930] 1988), esp. chap. 2 and Jonathan Dancy, *Ethics Without Principles* (Oxford: Oxford University Press, 2004). It could be argued that virtue theory is also a version of pluralism since the virtue theorist (usually) holds that there are a number of virtues which have moral importance and which are not reducible to each other. This issue depends on what it means for there to be "one" morally relevant feature (i.e., whether 'virtue' simpliciter can count as monistic or whether the virtues would have to be individuated to tell the right story), but this debate is not relevant for present purposes and I will leave virtue theory as one of the three dominant monistic traditions, though I realize that this interpretation could be challenged.

8. For example, see John Kleinig, "Ticking Bombs and Torture Warrants," *Deakin Law Review* 10.2 (2005): In press.

9. For further discussion of utilitarian considerations against torture, see Jean Maria Arrigo, "A Utilitarian Argument against Torture," *Science and Engineering Ethics* 10.3 (2004): 1–30 and Matthew Wynia, "Consequentialism and Harsh Interrogations," *American Journal of Bioethics* 5.1 (2005): 4–6.

10. It might be noticed that I have only considered act-, as opposed to rule-, utilitarianism. Of the rule-utilitarian, we might ask whether his rules ever have exceptions. If they do, rule-utilitarianism most likely collapses to act-utilitarianism insofar as we would expect the motivation for admitting the exceptions to be their contributions to aggregate happiness. In this case, the comments of this section would apply equally well to the exception-allowing rule-utilitarian. Alternatively, the rule-utilitarian might deny the existence of any exceptions to his rules. In this case, his view starts to resemble Kantian absolutism, at least structurally if not in content. The arguments made against Kant in §4 will apply equally well against the rule-utilitarian who denies exceptions to his rules; see that discussion if this is your preferred moral view.

11. Ronald Dworkin, *Taking Rights Seriously* (Cambridge, MA: Harvard University Press, 1977), ix.

12. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), 28–33.

13. Dworkin, *Taking Rights Seriously*, 191.

14. We should observe that this approach will yield different results than utilitarianism. For example, consider the "spare parts surgeon" case, which is supposed to inveigh against utilitarianism: imagine that we can kill one innocent person to provide life-saving

organ transplants to five people who will otherwise die. We could make enough further stipulations to commit the utilitarian to the harvesting, though the aggregating deontologist need not have this commitment. The reason is that the harvesting violates one right (*viz.*, that of the harvested), but not harvesting violates zero rights (since the dying patients do not have a right to the organs of the healthy person). Since the harvesting does not minimize overall rights violations, it is therefore impermissible; this is the conclusion many deontologists want.

15. This view is patently not a form of utilitarianism since considerations of *utility* are not deemed morally primary. It is a form of consequentialism, however.

16. See, for example, Amartya Sen, "Rights and Agency," *Consequentialism and Its Critics*, ed. Samuel Scheffler (Oxford: Oxford University Press, 1988), 187–223.

17. Nozick, *Anarchy, State, and Utopia*, 33.

18. Nozick, *Anarchy, State, and Utopia*, 30.

19. See, for example, Plato, *Republic*, 441c3–445e2.

20. Aristotle, *Nicomachean Ethics*, 1119b20–1128b36.

21. Francis Hutcheson, "Treatise II," in *An Inquiry into the Original of Our Ideas of Beauty and Virtue* (Indianapolis, IN: Liberty Fund, 2004), 83–197.

22. David Hume, *A Treatise of Human Nature* (2nd ed.), ed. Paul H. Nidditch (Oxford: Oxford University Press, 1978), III.II.i.

23. Hursthouse, *On Virtue Ethics*, 63–87

24. Darwall, *Contractarianism/Contractualism*, 1.

25. Darwall, *Contractarianism/Contractualism*, 1. This notion of "moral equality" is a tricky one, though I will return to it briefly in §4. While Darwall does not provide an account in his (2003), he is currently working on an account that holds moral equality is "equality in a kind of dignity of authority (to make claims and demands of one another—in particular to command a kind of respect I have called 'recognition respect')" (personal communication). See his *Second-Person Standpoint: Morality and Accountability* (Cambridge, MA: Harvard University Press, forthcoming). See also Elizabeth Anderson, "What Is the Point of Equality?" *Ethics* 109.2 (1999): 287–337.

26. Thomas Hobbes, *Leviathan*, ed. Edwin Curly (Indianapolis: Hackett Publishing Co., [1668] 1994), 74.

27. He writes: "To this war of every man against every man, this also is consequent; that nothing can be unjust. The notions of right and wrong, justice and injustice, have there no place. Where there is no common power, there is no law; where no law, no injustice." Hobbes (1994), 78.

28. For example, John Harsanyi's "impartial observer argument" holds that we should make social decisions as if we had an equal chance of being anyone in the society. Since our case stipulates that there are a lot more non-terrorists who may lose their lives than terrorists who may suffer torture, the rational choice (on his view at least—Rawls's maximin principle would yield the contrary result) would be to support the legislation of torture; this support would, on average (even to near-unity if the number of potential victims were high enough), be in the self-interest of the legislator. See John Harsanyi, "Cardinal Utility in Welfare Economics and the Theory of Risk Taking," *Journal of Political Economy* 61 (1953): 434–5 and "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons in Utility," *Journal of Political Economy* 63 (1955): 309–21. For the alternative view, see John Rawls, *A Theory of Justice* rev. ed. (Cambridge, MA: Belknap Press, 1999), 132–6.

29. Rawls, *A Theory of Justice*, 132–3.

30. Rawls, *A Theory of Justice*, 53.

31. Rawls, *A Theory of Justice*, §28.

32. Selim Barker, "Particular Reasons" (unpublished), 2. Presented at the Research School for Social Sciences of The Australian National University, June 16, 2005.

33. Ross, *The Right and the Good*, 21. Quoted. in Robert Audi, "Intuitionism, Pluralism, and the Foundations of Ethics," in *Moral Knowledge?: New Readings in Moral Epistemology*, ed. Sinnott-Armstrong and Timmons (Oxford: Oxford University Press, 1996), 103.

34. Barker, "Particular Reasons," 3.

35. See, for example, Jonathan Dancy, "The Particularist's Progress," in *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little (Oxford: Clarendon Press, 2000), 130–56 and Dancy (2004). Hooker and Little (2000) has a good collection of critical debate on moral particularism.

36. See, for example, Kleinig, "Ticking Bombs and Torture Warrants," and Davis "The Moral Justifiability of Torture" (*International Journal of Applied Philosophy*, this issue), 161–78.

37. While this sounds plausible as a minor point, I am not willing to concede that torture is, all things considered, worse than death, though I am willing to grant some *prima facie* worseness that it has in that comparison. It seems to me that there are enough other countervailing features (e.g., the benefits of a continued life—note that we could make these particularly rich by stipulation of our examples) that the *prima facie* and all things considered judgments are likely to diverge.

38. This is certainly a critical question in the torture debate, and one whose importance is recognized by both sides. For an opposing view regarding my forthcoming remarks, see Davis (this issue), esp. §3.

39. There are other theories as to what constitutes justification (e.g., contextualism, pragmatism, contract-based approaches, etc.), but I will herein limit my remarks to what I take to be the two most dominant theories (thought note that §3.4 partially addresses the contract-based approach to justification). Also note that the distinction between foundationalism and coherentism has been challenged; see R. M. Hare, "Foundationalism and Coherentism in Ethics," in Timmons and Sinnott-Armstrong (1996), 190–9.

40. See, for example, Robert Audi "Intuitionism, Pluralism, and the Foundations of Ethics." For more detailed discussion, see *Ethical Intuitionism: Re-Evaluations*, ed. Philip Lake-Stratton (Oxford: Oxford University Press, 2003).

41. For more discussion on this topic, see Philip Lake-Stratton, *Ethical Intuitionism: Re-Evaluations*.

42. Some critics might think that no moral theory could adequately capture all of the instances in which we think that we have foundational knowledge and that I am therefore postulating an unfairly high criterion for theory acceptance. Rather, maybe the idea is that we should take the theory that maximally accommodates our other commitments and one theory might be the best even if it cannot accommodate *all* of those commitments. In response to this line, I have a few remarks. First, this point seems to miss the power of *modus tollens*: if the theory implies some commitment, and that commitment is false (in virtue of foundational knowledge we have of its negation), then the theory is *false*. The way to escape this conclusion is not to debate the logic, but rather to adopt some alternative scheme of justification (such as coherence) which can admit of degrees (e.g., coherentists speak of "maximal" coherence).

43. Geoffrey Sayre-McCord, "Coherentist Epistemology and Moral Theory," in Timmons and Sinnott-Armstrong (1996), 178.

44. See Rawls *A Theory of Justice*, (Cambridge, MA: Belnap Press, 1999) 18–9, 42–5 and Sayre-McCord, "Coherentist Epistemology and Moral Theory," 142.

45. John Martin Fischer and Mark Ravizza, eds., *Ethics: Problems and Principles* (Belmont, CA: Wadsworth Publishing, 1992), 12–6.

46. Davis (this issue), 170.