

## CHAPTER SIXTEEN

# Ticking Time-Bombs and Torture<sup>1</sup>

*Fritz Allhoff*

In thinking about the moral status of interrogational torture, we are commonly asked to imagine exceptional cases wherein such torture is necessary to save some significant number of lives. These cases are generally referred to as ticking time-bomb cases and invite us to think of the relationship they bear to terrorism: some terrorist has planted a bomb in a crowded metropolitan center that will kill many noncombatants unless the terrorist is tortured. But terrorism need not have anything to do with these cases nor, really, do bombs. Rather, what matters is that there is some threat to many people that can only be avoided – and which, in most formulations, certainly will be – through the torture of someone already in custody. The “ticking time-bomb” locution is therefore somewhat narrow, but not in any drastically misleading way. Furthermore, most of the contexts worth considering – by which I mean the real-world ones most closely approximating these hypothetical constructs – probably will be those involving terrorists and weapons, if not necessarily bombs. And, for purposes of engagement, there is merit in following the standard usage.

### **Origins of Ticking Time-Bomb Cases**

Let us start by considering the origins of ticking time-bomb cases; in doing so, we will also start to understand the methodology. In terms of the philosophical literature, an early formulation owes to a seminal essay by Henry Shue:

[S]uppose a fanatic, perfectly willing to die rather than collaborate in the thwarting of his own scheme, has set a hidden nuclear device to explode in the heart of Paris. There is no time to evacuate the innocent people or even the movable art treasures – the only hope of preventing tragedy is to torture the perpetrator, find the device, and deactivate it. (Shue, 1978, p. 141)

*Contemporary Debates in Applied Ethics*, Second Edition. Edited by Andrew I. Cohen and Christopher Heath Wellman.

© 2014 John Wiley & Sons, Inc. Published 2014 by John Wiley & Sons, Inc.

But not only have these cases appeared in academic journals, they have also crossed over to popular media outlets and, thereafter, into public consciousness. For example, consider the following, which comes from an essay Michael Levin wrote in *Newsweek*:

Suppose a terrorist has hidden an atomic bomb on Manhattan Island which will detonate at noon . . . Suppose, further, that he is caught at 10 a.m . . . but preferring death to failure, won't disclose where the bomb is. What do we do? If we follow due process, wait for his lawyer, arraign him, millions of people will die. If the only way to save those lives is to subject the terrorist to the most excruciating possible pain, what grounds can there be for not doing so? I suggest that there are none. (Levin, 1982)

Long before these formulations of ticking time-bomb cases in the second half of the twentieth century, Jeremy Bentham wrote on the morality of torture almost two hundred years earlier, dating from the late 1770s (Bentham, 1973, pp. 56–62, 63–70).<sup>2</sup> Bentham probably had the first formulation of a case that looked anything like a ticking time-bomb case, though this came later – in 1804 – and was not otherwise attached to a systematic treatment of torture. Consider what he wrote:

Suppose an occasion, to arise, in which a suspicion is entertained, as strong as that which would be received as a sufficient ground for arrest and commitment as for felony – a suspicion that at this very time a considerable number of individuals are actually suffering, by illegal violence inflictions equal in intensity to those which if inflicted by the hand of justice, would universally be spoken of under the name of torture. For the purpose of rescuing from torture these hundred innocents, should any scruple be made of applying equal or superior torture, to extract the requisite information from the mouth of one criminal, who having it in his power to make known the place where at this time the enormity was practicing or about to be practiced, should refuse to do so? To say nothing of wisdom, could any pretense be made so much as to the praise of blind and vulgar humanity, by the man who to save one criminal, should determine to abandon [one hundred] innocent persons to the same fate? (Twining and Twining, 1973, p. 347)

The principal difference between Bentham's case and the others previously presented is simply whether the harm that the torture aims to dispel is already active (*viz.*, the current torture of innocents) or else prospective (*viz.*, the future explosion of a bomb). Morally, there need not be any difference between these cases: what matters is whether the torture is necessary to prevent the harm. If that harm is temporally distant, then that would undermine the need to torture in so far as there might be other – and less morally offensive – ways to dispel it. But so long as the torture is necessary, then whether the harm is ongoing, imminent, or even temporally distant is irrelevant. We will return to this below but, for now, the point is merely that Bentham's case is structurally similar to the others.

Aside from this more casual presentation of a single case, Bentham also offered a more extended treatment of torture, as recorded in two manuscript fragments. It is worth considering these fragments for at least three reasons, only the first of which is historical. The second is more philosophical in that Bentham starts to elucidate some of the key logical elements of ticking time-bomb thinking, even if that discussion floats free of a particular ticking time-bomb case. And the third bears on the relationship between ticking time-bomb methodology and utilitarianism, a relationship that is more

complicated than usually acknowledged. With these three reasons in mind, let us now look at some of Bentham's writings on torture.

First, Bentham asserts that torture may be applied in two cases. "The first is where the thing which a Man is required to do being a thing which the public has an interest in his doing, is a thing which for a certainty is in his power to do" (Twining and Twining, 1973, p. 312). And he continues that torture is otherwise permissible:

where a man is required what probably though not certainly is in his power to do; and for the not doing of which it is possible that he may suffer, although he be innocent; but which the public has so great an interest in his doing that the danger of what may ensue from his not doing it is a greater danger than even that of an innocent person's suffering the greatest degree of pain that can be suffered by Torture, of the kind and in the quantity permitted to be employed. (Twining and Twining, 1973, pp. 312–313)

Then Bentham asks: "Are there in practice any cases that can be ranked under this head? If there be any, it is plain that there can be very few (Twining and Twining, 1973, p. 313). That Bentham was reserved about the extent to which torture can be justified is noteworthy: being a utilitarian hardly commits one to the promiscuous use of torture, as there are myriad utilitarian reasons to oppose it (Twining and Twining, 1973, pp. 348–350). We will discuss some of these below, but I want to get early purchase on the concept of exceptional – as opposed to normalized – torture.

After these introductory remarks, Bentham goes on to offer a series of moral rules that have to be satisfied for the legitimate application of torture. While the details of those rules need not concern us here, they are precisely the sorts of principles that undergird contemporary ticking time-bomb cases. For example, torture should not be applied without (near-) certainty that the would-be tortured has the relevant knowledge (Rule 1); that torture is only appropriate as a last resort in "cases which admit of no delay" (Rule 3); that minimal means should always be preferred to extreme ones (Rule 4); that the prospective benefits are greater than the prospective costs (Rules 5 and 7), and so on (Twining and Twining, 1973, pp. 312–315).

Importantly, many of these rules were effectively codified in Bentham's hedonic calculus, published shortly thereafter (Bentham, 2007, ch. 4). Bentham predicated his utilitarianism on seven factors, all of which are at least implicitly manifest in the ticking time-bomb cases: intensity, duration, certainty (or uncertainty), propinquity (or remoteness), fecundity, purity, and extent. While intensity and duration are rarely emphasized in the cases, they certainly *could* be, and such invocations would seemingly only make the cases more compelling: imagine that the terrorist need only be subject to a "comparatively minor and brief" form of torture to disclose the location of the bomb. (Note that this is not to suggest that torture could ever be minor – which some might argue to be incoherent – but rather that it most certainly comes in degrees and could be *comparatively* minor.)

The other features, though, are at least near-explicit in the cases. Certainty is perhaps the most conspicuous feature of ticking time-bomb cases: *everything* is certain. It is certain that the detainee is a terrorist. It is certain that he has information regarding the location of the bomb. It is certain that the torture will produce the information. It is certain that the information will lead to the timely deactivation of the bomb. And many critics of the cases promptly seize upon all of these, which undoubtedly

represents a departure from (at least almost all) actual cases (Shue, 2006, pp. 231–239). I will return to this criticism below.

Next come fecundity and purity, which are opposite sides of same coin: when we torture, we will get even more good things (fecundity) and these good things will not be offset by any bad things (purity). This fecundity is thereafter magnified by the invocation of extent, the last of Bentham's elements, which holds that, not only will a single life be saved through the torture, but rather *a lot* of lives will be saved (recall Levin's "millions"). The purity condition is fulfilled in so far as no bad consequences – aside from the pain and suffering of the tortured – are postulated. And, while it is open for the critic to say that such cases do not preclude such consequences, it is equally open to the proponent to merely issue such a stipulation, at least at this stage of the dialectic.

Critics nevertheless do complain about ticking time-bomb cases precisely on the issue of purity. They assert that torture would have to be institutionalized (Davis, 2005, pp. 161–178); including the implementation of training programs for the torturers (Wolfendale, 2006, pp. 269–287); that such institutionalization portends harms for liberal democracies; that our torturing our enemies makes it more likely that our enemies will torture us (Arrigo, 2004, pp. 4–6); that torture makes it more likely for us to perpetuate other wrongs (Fiala, 2005, pp. 127–142); and so on. We will return to these issues below but, for now, I just want to mention them.

Regardless, the proponent of ticking time-bomb methodology is still free to say: "look, that just *is not how the case goes!*"; the critic cannot load conditions into the case that are patently excluded by presupposition. To do so is simply to change the case and to ask a different question altogether, and precisely not the one that we currently care about. Rather, the question at hand is whether torture is permissible given features either stated or implied in the ticking time-bomb cases, and this is a question on which moral philosophy owes an answer. Following that inquiry, we can then think about what implications it has vis-à-vis (real world) cases that relax some of the idealizations and abstractions. But, as the dialectic goes, the aforementioned complaints are completely irrelevant. Let us now move past these methodological preliminaries and to the normative upshot.

## Utilitarian Views on Torture

Of all the major moral theories, utilitarianism probably offers the most direct justification of interrogational torture. In the next section, "Deontological Views on Torture", I will argue that rights-based theories need not be opposed to torture, but making that case plausible certainly requires more work than it does in the case of utilitarianism. For example, we might think that people – terrorists or otherwise – have a right against being tortured, and being able to take rights seriously while still being able to license torture therefore presents an obvious obstacle. For the utilitarian, though, no such obstacle exists: so long as the hedonic calculus comes out right, torture is readily justified. Opponents of torture therefore can go either of two ways. First, they can reject utilitarianism, whether for more general reasons or specifically because it permits torture. In doing so, they must adopt some sort of moral theory in its place; this could either be a rights-based approach (see following section) or something else. Second.

they can deny that utilitarianism would justify torture in practice, even if it could in theory. This line has the merit of being responsive to utilitarian commitments and, therefore, is able to say something to the utilitarian; the first line simply does not engage the position. For this reason, I have a lot of respect for critics who advance utilitarian arguments against torture in so far as those arguments are responsive to the dialectic: a convincing response to the utilitarian should not be that his theory is wrong – presumably he has already thought about this – but rather that his theoretical commitments are different than he might have suspected (Arrigo, 2004).

As mentioned in the previous section (Origins of Ticking Time-Bomb Cases), the basis of a utilitarian argument for torture is straightforward: it could prevent the deaths of many noncombatants. Their noncombatant status does not matter from the utilitarian perspective, so long as the aggregate value of those otherwise forfeited lives is sufficiently high to cover the costs of torture. For the utilitarian, torturing one terrorist to save one thousand other terrorists could even be justified, at least if the thousand to be saved were not going to sow too much suffering on the world. It is worth noting that the utilitarian calculation here is about preventing a utility loss rather than effecting a utility gain: it is not that torture brings more utility into the world, but rather prevents its longitudinal exit. This distinction, though, is not of much use to the utilitarian, who only cares about maximizing total aggregate happiness: whether that quantity turns out to be positive or negative is beside the point, so long as it is maximal. In other words, torture could manifest a utility loss yet still be justified so long as that loss is less than the alternatives that would be realized without torture.

Again reprising an earlier argument, there are several desiderata that the utilitarian would require torture to satisfy. First, torture should be the least harmful remedy applied, and, similarly, some insufferable form of torture should not be deployed when a lesser one would elicit the valuable information. If the information that would save lives could be gleaned through some less offensive means (e.g., simple questions), then those means should be pursued. There is little reason to torture before otherwise asking the location of the bomb. Furthermore, there should be some expectation that the torture will be efficacious, be it: against someone who we can reasonably expect to have the intelligence; against someone with the appropriate vulnerabilities; not against someone who will hopelessly deceive us with misinformation; effected on a timetable commensurable with the extant threat, and so on. There is no doubt that, in the real world, all of these requirements can get messy but, in the land of ticking time-bombs – where we currently are – they are straightforwardly stipulated.

While these stipulations are often challenged, such challenges radically misunderstand the state of play. For example, consider an ill-named paper by Vittorio Bufacchi and Jean Maria Arrigo, "Torture, Terrorism and the State: A Refutation of the Ticking Time-Bomb Argument", which is anything but (Bufacchi and Arrigo, 2006, pp. 355–373). Bufacchi and Arrigo think that the ticking time-bomb argument takes its premises to be that a terrorist is captured and that, if he is tortured, he will reveal information regarding the location of a bomb. From this, the conclusions are meant to be that torture is permissible, that information about the bomb is retrieved, and that lives are saved (Bufacchi and Arrigo, 2006, p. 360). But then they argue that this formulation has suppressed premises (e.g., that it is almost certain that the terrorist has information about the bomb) and, more to the point, that "all the premises in the argument are contentious from an empirical point of view": intelligence is never infallible,

torture is not guaranteed to work, torture is not efficacious in short time periods, misinformation is revealed under torture, and so on (Bufacchi and Arrigo, 2006, pp. 361–362). Ultimately, they maintain that the ticking time-bomb argument fails by its own utilitarian lights. And this sort of strategy has been repeated elsewhere by others, be it about the institutional costs of torture (Davis, 2005, pp. 161–178) – including the costs training programs for torturers (Wolfendale, 2006, pp. 269–287) – or various other negative consequences (Brecher, 2007, chs 2–3). Assuming that all of these objections can be developed in utilitarian currency, where does the utilitarian argument for torture stand?

As discussed elsewhere, I look at ticking time-bomb cases not as arguments, but as cases (Allhoff, 2012, ch. 5). The upshot of these cases is some sort of moral judgment about the permissibility of torture, and those judgments somehow figure into our moral methodology through, for instance, a reflective equilibrium in which those judgments align with our moral principles (Rawls, 1999, pp. 42–45). We can make ticking time-bomb cases figure into some argument, but that argument looks something like this: (1) in all cases – and all else being equal – if we can choose a lesser harm to a greater one, we should; (2) in ticking time-bomb cases, torture is the lesser harm; therefore (3) we should torture in those cases. I take the first premise to be self-evident. The second is the key kernel encoded into ticking time-bomb cases. The conclusion deductively follows.

Bufacchi and Arrigo's formulation seems clunky by comparison, though maybe they aim to break out my second premise. Regardless, their strategy is just to deny some presuppositions of the cases on empirical grounds. But this move misses the point of the dialectic since the realm of discourse at this stage is non-empirical; rather, it is about the hypothetical cases. Their response is a red herring in so far as they change the question from being about ticking time-bomb cases to being about the world. Ultimately, their position has to do with torture policy, but ticking time-bomb cases are not about torture policy; they are about one-off applications of torture (Wisnewski, 2009, pp. 205–209).

What matters to the utilitarian is how the hedonic calculus plays out in the cases under consideration. For the purposes of this chapter, we are still concerned with ticking time-bomb cases, and those straightforwardly license the moral permissibility of torture under a utilitarian approach. Elsewhere, I afford empirical considerations more discussion (Allhoff, 2012, ch. 7). Note, though, that opponents of torture offer an important concession when they appeal to empirical issues. Their empirical turn presumably allows that torture *could* be justified; otherwise, there is no need – save, perhaps, dialectical expedience – to make that move. If the allowance that torture could be justified is nevertheless tightly followed by the contention that it *would* not be justified, then we just have to look at the cases and see how it comes out. By even getting to the cases, though, the opponent to torture has lost ground.

## Deontological Views on Torture

Many people want little to do with utilitarianism. This is a theory that supposedly lets the sadists abuse their victims, cares more about total aggregates than fair distributions, and so on. For our purposes, though, it is this first concern that matters in so far as

utilitarianism makes no provisions for deontological constraints, such as the right not to be tortured. Does everyone have such a right? Is it absolute?

To get the proper point on the first question, consider this: does the terrorist responsible for threatening many lives and unwilling to end that threat have a right against torture? I am not sure. Certainly there are things that we can do to forfeit our rights. A negligent parent forfeits his right of custodianship to his child. A drunk driver forfeits his right to operate a motor vehicle. A murderer forfeits his right to freedom when incarcerated. Perhaps, then, a terrorist can forfeit her right against torture. This suggests one way we can accept a right against torture but deny that it prohibits torture in all circumstances: we allow that people can forfeit the right. Call this the *forfeiture strategy*.

Another rights-based defense of torture denies that the right entails absolute prohibitions. Many rights prohibit people from doing things to the bearers of rights. Property rights, for instance, forbid use without permission. But in emergencies we might say these rights are permissibly infringed. Perhaps the right against torture is like that: people always have this right, but infringing it might be justifiable nevertheless. Call this the *justified infringement strategy*.

The justified infringement strategy argues against a stronger position than the forfeiture strategy. The latter simply sets aside the right against torture. But the justified infringement strategy grants both that rights against torture are universal and that they always give reasons for people not to violate them. It simply appeals to weightier reasons. This is the approach I will pursue.

The terrorist's rights are not the only rights in play. The victims of the bomb stand to have their rights disrespected as well (e.g., their rights to life). Someone's rights are going to be infringed: either the terrorist's right against torture or else the noncombatants' respective rights to life. Therefore, just trotting out the terrorist's right against torture misses the central point of ticking time-bomb cases, even from a rights-based perspective.

Imagine that A is getting ready to murder B and C and that these murders can only be prevented if D shoots A. What should D do? The deontologist can say one of two things. On the one hand, he could say that D cannot shoot A because rights are, to use some famous locutions, "trumps" (Dworkin, 1978, p. ix) or "side constraints" (Nozick, 1974, pp. 28–33). On this view, we may not disrespect some right *even if* doing so 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 may shoot A: if he does, there would be one rights-infringement (*viz.*, A's) and, if he does not, there would be two (*viz.*, B's and C's). Rights-violations are minimized through the shooting, thus making it morally permissible.

Which is the more plausible view? I think that the aggregative approach makes more sense. Ronald Dworkin has argued that in cases of rights conflict, we should look to the values that suggested the right in the first place (Dworkin, 1978, p. 191). So, if individuals have a right to life, it is because life itself is something that is valuable and worth preserving. Given a conflict, for instance, where the disrespecting one person's right to life could prevent the violation of five other persons' rights to life, the values that led to the creation of the right to life would suggest infringing the one in order to prevent

violating of the five. We endorse rights in the first place because we value the objects of those rights. Thus it would be permissible to act such that the underlying values and their associated objects are preserved to the highest degree possible. As I suggested above, someone who took rights seriously might think that some aggregation procedure is more attractive than viewing rights as necessarily absolute. Nozick does, however, raise a legitimate concern against aggregation. 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?" (Nozick, 1974, p. 33).

One possible response is to say that the ordinary implications of rights do not apply in emergencies. Indeed, in a footnote, Nozick expresses concern with the application of his theory to "cases of catastrophic moral horror" (Nozick, 1974, p. 30). These are precisely what ticking time-bomb cases model. So, ultimately, it is unclear whether even Nozick would object to rights violations in these cases.

Let us now integrate various threads from this section. First, ticking time-bomb cases turn on conflicts of rights and not merely the rights of the terrorist. Second, torture can minimize overall rights violations in ticking time-bomb cases. Third, we can acknowledge that not all rights violations are of equal moral significance. We might even suppose for the sake of argument that the right against torture counts more than the right to life. While I suspect that this latter supposition is false, the rights of a substantial number of people who would otherwise die militates in favor of torture. To put it another way, even if the right against torture is five times more morally valuable than the right to life, the rights of the thousands that the terrorist threatens swamp his right against torture. On a straightforward rights aggregation, torture comes out to be morally permissible. And even if we were to acknowledge that there could be problems with such an approach in general – such as those suggested by Nozick – they might not apply to ticking time-bomb cases in particular. Therefore, they fail to establish the absolute moral impermissibility of torture.

## **Rejecting Torture: Absolutism-in-Principle and Absolutism-in-Practice**

The two preceding sections evaluated the utilitarian case for torture, as well as a particular rights-based approach. Elsewhere, I have argued that a range of other moral theories (e.g., virtue theory, social contract theory, etc.) can license torture in exceptional circumstances (Allhoff, 2005, pp. 243–264). However, critics may argue that I have not grappled with absolutist moral prohibitions.

Let us say that absolutism with regards to torture can come in two different forms: absolutism-in-principle and absolutism-in-practice (henceforth let us designate these a-principle and a-practice, respectively). The difference between these two is modal: a-principle holds that torture never *could* be justified whereas a-practice holds that it could, but never *would*, be justified. In other words, a-practice allows that torture could be justified, but denies that whatever circumstances are sufficient for this justification – perhaps including those of ticking time-bomb cases – will be manifest in the real world. A-principle, by contrast, denies that such circumstances are possible at all.

First, it bears emphasizing that comparatively few people explicitly defend a-principle; even those who do defend absolutism with regards to torture overwhelmingly

advance arguments defending a-practice (Mayfeld, 2008; Tindale, 2005). The limited advocacy for a-principle is not surprising given that it is a very extreme position, effectively the most extreme that one could possibly advocate. To wit, a-principle postulates some class of actions that, regardless of circumstances – whether real or imagined – is always impermissible. No view could be more extreme in so far as it is logically impossible to allow torture in any fewer cases. Regardless, some people do defend a view like this. For example, consider Kim Lane Scheppele: “I do in fact believe that torture is always and absolutely wrong, given the position we should accord to human dignity, even of terrorists” (Scheppele, 2005, p. 287). Or Ben Juratowitch: “torture is so barbaric that the right to be free from it is never defeasible. However desperate the countervailing circumstances, torture is always wrong and should never occur” (Juratowitch, 2008, p. 81). Anyone of a sufficient Kantian bent would also defend a-principle since interrogational torture treats its victim as a means to our end, operates on a maxim that generates a conflict in will, fails to respect autonomy, and so on; any act with these commitments is, to use a Kantian term, categorically impermissible.

That said, there are reasons to think that a-principle is implausible. Consider any candidate moral value with which one might construct a moral theory. Suppose there is some act,  $\Phi$ , that (maximally) promotes its pre-eminent value  $V$  in some case  $C$  (hereafter, I use “promotes” to mean “maximally promotes”). Every moral theory other than a-principle treats this supposition in a straightforward way and endorses  $\Phi$ . A-principle, however, needs it to be the case that  $\Phi$  is ruled out a priori (i.e., without consideration of the details of  $C$ ). There are two ways this exclusion could go. First, a-principle could deny that  $\Phi$  promotes  $V$  (in  $C$ ). So the idea would be that, necessarily,  $\Phi$  is incompatible with whatever value undergirds our moral theory. Or, second, a-principle could say that even if  $\Phi$  promotes  $V$  (in  $C$ ),  $\Phi$  is nevertheless impermissible.

This second response borders on the nonsensical in so far as some theory’s commitments preclude action to promote its own pre-eminent value. It even seems contradictory in so far as the theory espouses  $V$  while, at the same time, obstructing its realization (i.e., disavowing  $\Phi$ , which promotes  $V$ ). This is not to object to there being *other* theoretical commitments that  $\Phi$  might compromise, but that just forces us to elaborate the argument. For example, maybe  $\Phi$  promotes  $V_1$  while, at the same time, diminishing  $V_2$ ; imagine that these are both values of some particular a-principle. Now there are theoretical resources to oppose  $\Phi$ , namely its impact on  $V_2$ . But all this shows is that our first example was overly simplistic since it postulated only a single value for the theory. If we allow that  $\Phi$  promotes  $V_1, V_2, \dots, V_n$  (i.e., all the values of a-principle), then it will still be curious that  $\Phi$  is prohibited.

A-practice, unlike a-principle, makes substantive claims about the world. While a-principle holds that there are no circumstances under which torture *could* be justified, a-practice holds that there are no circumstances under which torture *would* be justified; this former claim is true or false *a priori*, whereas the latter is true or false *a posteriori*. In other words, to argue that torture is never actually justified requires some engagement with the world, or at least it does to the extent that a-practice means to be saying anything different from a-principle. To wit, a-principle also holds that torture is never actually justified, though this “actually” is redundant in so far as torture never could have been justified, whether actually (i.e., in the real world) or in imagined cases.

Various people defend a-practice, but a useful formulation is that by Daniel Statman, who argues that “[t]he moral danger of torture is so great, and the moral benefits so

doubtful, that in practice torture should be considered as prohibited absolutely” (quoted in Gross, 2009, p. 146). I choose this formulation among others because it usefully sets his view apart from a-principle in so far as the conclusion pertains to the implementation of torture *in practice*. It is perfectly consistent with Statman’s thesis that there are fantastic cases wherein torture could be justified, but that none of those cases would ever obtain in the real world. Henry Shue writes something similar, slightly taking himself to task for his earlier work: “I now take the most moderate position on torture, the position nearest the middle of the road, feasible *in the real world*: never again” (Shue, 2006, pp. 213–239).<sup>3</sup> Regarding the observation above, it is possible that Statman and Shue actually endorse a-principle, but then why do they talk about “practice” and the “real world” in their papers? Given their choice of language, a-practice is the more natural reading. Regardless, this section is not about either of their views in particular, but rather about a class of views that emphasize practice instead of theory; theirs are merely likely candidates.

I think of a-practice as making empirical claims about the world as opposed to being committed to any particular moral principles more generally. In other words, a-practice does not have any moral commitments other than the one that torture cannot be justified in the real world; such a view tells us nothing about the moral permissibility of, for example, abortion or euthanasia. This is in contrast with the views considered in the two previous sections; these were independent moral theories whose commitments we could evaluate vis-à-vis torture. Utilitarianism, for example, says that we should maximize total aggregate happiness, and then we can ask whether torture (in some particular case) does so. Not so for a-practice, which does not tell us anything other than that torture will never be justified in the real world. Aside from this sole moral commitment of a-practice, it is otherwise compatible with a wide range of moral theories. For example, there are utilitarian defenses of a-practice, and there could be myriad other defenses of it as well: Torture will never promote rights, virtues, the social contract, or whatever else takes moral priority. It is worth noting that this is not meant to be a criticism of a-practice; in fact, I take it to be a strength.

So let us ask: does the defender of a-practice think that we should torture in ticking time-bomb cases? For such a person, this question is just not interesting since it (allegedly) gains no traction with the world. What matters is not whether we should torture in ticking time-bomb cases, but rather that there are no such cases. Ethics is about action, and we can only act in the world, not in imagined hypotheticals. Elsewhere, I have defended ticking time-bomb methodology against its critics (Allhoff, 2012, ch. 5), but the criticism I now mean to engage is not concerned with the methodology *per se*, but rather tries to render that methodology impotent by appeal to empirical facts.

In response, there are several important points to develop. First, a-practice is only as plausible as its empirical claim that torture will *never* promote our core moral values. It is very hard to show that such a thesis is true. For example, the ivory-billed woodpecker was thought extinct since 1944, only to turn up in an Arkansas swamp in 2005. None of these woodpeckers was seen for over sixty years, and the received view was that they were gone forever. It turns out that the birds were still around and the naysayers were wrong. The point is that it only takes one case to falsify a negative existential claim, whether it be a sighting of a believed-to-be-extinct bird or some constellation of improbable features leading to justifiable torture.

The defender of a-practice is doubtless going to say that this is all hopelessly confused. His claim is not that justifiable torture is as elusive as the ivory-billed woodpecker, but rather that justifiable torture *cannot* exist (in the real world) because the benefits are too low and/or the costs too high. While it is *possible* for a rare bird to be spotted, it is not possible to justify a heinous act of torture. But, as I read this response, it just *denies my premise, namely that the empirical possibility of justified torture cannot be ruled out*. My position clearly has the dialectical advantage in so far as its only claim is that I allow that such and so *might* happen (in the real world) and, if so, torture could be permissible. My opponent has to say that this is false, but such empirical omniscience from a philosophical armchair is untenable.

The better way to go is something quasi-historical: to say that we have never seen a ticking time-bomb case in the real world, have no reason to think that there ever will be one, and so on. Sure, it is *possible* that such a case will arise, though we have no good reason to think that it will and, indeed, many good reasons to think that it will not. On this sort of inductive approach, we have a lot more confidence that we could find an ivory-billed woodpecker than that torture could be justified; there might have been some reason to hold out hope in the former case (e.g., ornithologists have sometimes made mistakes), but there is no such reason in the latter.

Again, these are empirical claims and ones that are rarely supported by any sort of actual empirical evidence. It is probably more fair for the defender of torture to owe us an actual ticking time-bomb case than for the critic to have to show that there have not been any, and I have discussed this elsewhere (Allhoff, 2012, ch. 7). The point at present is merely that the empirical details matter. And, in fact, a-practice rarely takes this challenge very seriously. We hear plausible-sounding claims about institutional requirements, potential for abuse, the ineffectiveness of torture, and so on, but these are almost always offered without any serious empirical engagement. Disparaging remarks about torture in Argentina and Abu Ghraib are also common, yet these are about as far from ticking time-bomb cases as one could get and still be talking about torture. That said, the empirical work is sometimes taken quite seriously, and Darius Rejali's painstaking and magisterial work is the best example, particularly its investigations into past torture regimes (Rejali, 2009). As impressive as Rejali's book is, though, it makes hardly any reference to the relevant philosophical literature. This is not necessarily a criticism of his work – it has different goals and comes from a different disciplinary orientation – the only point is that, in the end, the empirical and the philosophical need to be integrated.

## Conclusion

Let me make a few final remarks about a-practice. I certainly think that this is the best way to go in opposing torture. As indicated in the two sections on deontological and utilitarian views on torture, it is just not very hard to justify torture in theory. Rather, the rub comes when moving to practice, and a-practice is unequivocal in that regard: no torture in practice, ever. Two provisional conclusions bear emphasis. First, there is strong antecedent pressure against the thesis that torture would *never* be justified for the simple reason that we do not know all the scenarios we will ultimately encounter.

It seems to me that the more plausible route is to leave open the possibility of justifiable torture, while being skeptical and protective about its application. A-practice is not content with this less-ambitious approach and, for me at least, therein lies one of its faults. Second, a-practice is highly committed to several empirical assumptions, whether about the non-existence of ticking time-bomb cases, the low benefits of torture, its high costs, and so on. All of these assumptions have to be examined and defended, and in ways that are sufficiently attuned to the appropriate philosophical issues. I submit that this has not been done by opponents of torture, quite probably because such an attempt would founder.<sup>4</sup>

## Notes

- 1 This essay is adapted with permission from Fritz Allhoff (2012), ch. 5, pp. 88–92; ch. 6, pp. 114–139, especially sections 5.1, 6.1–6.5. Reproduced by permission of The University of Chicago Press.
- 2 He also provided one of the earlier characterizations of torture: “Torture, as I understand it, is where a person is made to suffer any violent pain of body in order to compel him to do something or to desist from doing something which done or desisted from the penal application is immediately made to cease” (Twining and Twining, 1973, p. 309).
- 3 In earlier work, after presentation of a ticking time-bomb case, Shue writes: “I can see no way to deny the permissibility of torture in a case *just like this*” (Shue, 1978, p. 141). Shue goes on to issue various disclaimers about whether these cases would actually obtain, but he does not rule out the possibility. The more recent paper is unequivocal in its practical prohibition on torture.
- 4 For more discussion on the empirical issues, however, see Allhoff (2012, ch. 7).

## References

- Allhoff, F. (2005) A defense of torture: separation of cases, ticking time-bombs, and moral justification. *The International Journal of Applied Philosophy* 19 (2): 243–264.
- Allhoff, F. (2012) *Terrorism, Ticking Time-Bombs, and Torture: A Philosophical Analysis*. London: University Of Chicago Press.
- Arrigo, J.M. (2004) A utilitarian argument against torture. *Science and Engineering Ethics* 10 (3): 543–572.
- Bentham, J. (1973) Of torture. Reprinted in W.L. Twining and P.E. Twining, Bentham on torture. *The Northern Ireland Legal Quarterly* 24: 307–356.
- Bentham, J. (2007 [1789]) *An Introduction to the Principles of Morals and Legislation*. Mineola, NY: Dover Publications.
- Brecher, B. (2007) *Torture and the Ticking Bomb*. Malden, MA: Blackwell Publishing.
- Bufacchi, V. and Arrigo, J.M. (2006) Torture, terrorism and the state: a refutation of the ticking-bomb argument. *Journal of Applied Philosophy* 23 (3): 355–373.
- Davis, M. (2005) The moral justifiability of torture and other cruel, inhuman, or degrading treatment. *The International Journal of Applied Philosophy* 19 (2): 161–178.
- Dworkin, R. (1978) *Taking Rights Seriously*. Cambridge, MA: Harvard University Press.
- Fiala, A. (2005) A critique of exceptions: torture, terrorism, and the lesser evil argument. *The International Journal of Applied Philosophy* 20 (1): 127–142.
- Gross, M. (2009) *Moral Dilemmas of Modern War: Torture, Assassination, and Blackmail in an Age of Asymmetric Conflict*. Cambridge: Cambridge University Press.

- Juratowitch, B. (2008) Torture is always wrong. *Public Affairs Quarterly* 22 (2): 81–90.
- Levin, M. (1982) The case for torture. *Newsweek*, February 7, 1982, p. 7.
- Nozick, R. (1974) *Anarchy, State, and Utopia*. New York: Basic Books.
- Rawls, J. (1999) *A Theory of Justice*. Revised edn. Cambridge, MA: Belknap Press of Harvard University Press.
- Rejali, D. (2009) *Torture and Democracy*. Princeton, NJ: Princeton University Press.
- Scheppele, K.L. (2005) Hypothetical torture in the War on Terrorism. *Journal of National Security Law and Policy* 1: 285.
- Shue, H. (1978) Torture. *Philosophy and Public Affairs* 7: 124–143.
- Shue, H. (2006) Torture in dreamland: disposing of the ticking bomb. *Case Western Reserve Journal of International Law* 37 (2 & 3): 231–239.
- Tindale, C.W. (2005) Tragic choices: reaffirming absolutes in the torture debate. *The International Journal of Applied Philosophy* 19 (2): 209–222.
- Twining, W.L. and Twining, P.E. (1973) Bentham on torture. *The Northern Ireland Legal Quarterly* 24: 307–356.
- Wisniewski, J.J. (2009) Hearing a still-ticking bomb argument: a reply to Bufacchi and Arrigo. *Journal of Applied Philosophy* 26 (2): 205–209.
- Wolfendale, J. (2006) Training torturers: a critique of the ticking-bomb argument. *Social Theory and Practice* 32 (2): 269–287.

## Further Reading

- Allhoff, F. (2009) The war on terror and the ethics of exceptionalism. *Journal of Military Ethics* 8 (4): 265–288.
- Kershner, S. (2005) For interrogational torture. *The International Journal of Applied Philosophy* 19 (2): 223–241.
- Luban, D. (2005) Liberalism, torture, and the ticking bomb. *Virginia Law Review* 91 (6): 1425–1461.
- Mayfeld, J. (2008) In defense of the absolute prohibition on torture. *Public Affairs Quarterly* 22 (2): 109–128.
- Sen, A. (1988) Rights and agency. In *Consequentialism and its Critics*, ed. S. Scheffler, pp. 187–223. Oxford: Oxford University Press.