

## Empirical Objections to Torture: A Critical Reply

FRITZ ALLHOFF

Department of Philosophy, Western Michigan University, Kalamazoo, Michigan, USA; and Centre for Applied Philosophy and Public Ethics, Kingston, Australian Capital Territory, Australia

*Those who support torture in ticking-time-bomb cases are often criticized as failing to consider empirical objections to torture; however, torture's critics often wield this charge uncritically, doing little more than throwing out platitudes without considering the role of those platitudes in the dialectic. I agree with the critics that more empirical engagement is owed than is typically on offer, but deny that such engagement vindicates their position. This article therefore considers various stock objections to the actual use of torture, while ultimately arguing that those objections fail to undermine the use of torture in exceptional cases. In particular, we will consider the efficacy and reliability of torture, the institutional requirements for torture, the nefarious spread of torture, and whether there are better alternatives to torture. In each of these discussions, let us frame them against the associated contentions made by critics regarding the inapplicability of ticking-time-bomb cases to the real world. The last two sections of the article consider where the burden of proof falls in this debate and, in particular, whether the proponent of torture needs real-world ticking-time-bomb cases to defend exceptional torture, as well as what such cases might be.*

**Keywords** ethics, philosophy, terrorism, ticking bombs, torture

### Torture Doesn't Work

One of the central assumptions of ticking-time-bomb methodology is that torture will be applied to a guilty terrorist who, once tortured, will readily divulge information that can be used to disarm a lethal threat in which he is complicit. Regarding the efficacy and reliability of torture, this assumption is collectively predicated on the following: that we have apprehended the responsible terrorist; that torture will produce valuable information; and, conversely, that torture will not produce confusing misinformation. In other words, there are various ways in which ticking-time-bomb assumptions can go wrong. First, we might have the wrong person, whether an innocent person altogether or else a person without actionable intelligence in the relevant regard (even if a terrorist overall); a related worry is that we would keep torturing someone once he has already given us all the information he had such that this torture was superfluous. Second, the terrorist, once

Fritz Allhoff is an Associate Professor in the Department of Philosophy, Western Michigan University, and a Senior Research Fellow at the Centre for Applied Philosophy and Public Ethics, Australia.

Address correspondence to Fritz Allhoff, Department of Philosophy, Western Michigan University, Kalamazoo, MI 49008, USA. E-mail: fritz.allhoff@wmich.edu

tortured, might not reveal any valuable intelligence. Or worse, third, he might actually degrade the status of our countermeasures by giving us false intelligence, thus leading us to use our resources less well than we would have done had we not tortured at all.

Since critics have pressed all these issues, serious discussion is warranted; and I will turn to that shortly. But, first, it is important to clarify what the state of the dialectic is meant to be at this stage, particularly insofar as at least some critics have lost the thread. Granting that any of the three concerns raised in the previous paragraph are reasonable, what is supposed to follow? Or, to put it another way, when people say that “torture doesn’t work,” what is that supposed to mean? Assuming that the metric for whether torture works is whether it elicits valuable intelligence, an implausible position is that torture *never* works. I will return to this in §6, but it has to be the case that at least one person has, under torture, disclosed lifesaving information; whether the torture was necessary for that information is, for now, irrelevant. It also has to be the case that torture does not *always* work as there are surely cases where the victim of torture has failed to disclose valuable information. If the claim is that “torture never works,” that can be rejected out of hand, but it is a straw man to saddle the defender of torture with the alternative claim that “torture always works” since nobody would say that anyway. Let us therefore agree that torture sometimes works, which seems the most reasonable empirical position.

What implications does that have for the permissibility of torture? In other words, if torture is not always efficacious, do we have less reason to endorse it? And the answer, of course, is yes: as the efficacy of torture goes down, the moral benefits thereof fall. Consider, for simplicity, a variant on the standard ticking-time-bomb case:

*Red Sweatshirt.* Our intelligence reveals that a terrorist has just set up a bomb in a crowded building and has exited wearing a red sweatshirt. Law enforcement sets up a perimeter and starts to canvas the area; two men in red sweatshirts are apprehended, both of whom deny any knowledge of terroristic activity. Run the rest of the story as in standard ticking-time-bomb cases, the adjustment being that the bomb can be disarmed only if both men, one of whom is innocent, are tortured.

Is this torture permissible? Surely, it is worse to torture a guilty person and an innocent one than to torture only a guilty one, but I maintain that this torture could still be justified if there are enough people at risk in the building.

It is worth noting that the above hypothetical is synchronic rather than diachronic, but nothing changes—for me at least—if the considerations are diachronic. In other words, we were meant to imagine torturing two people at the same time in order to dispel some single threat (at that time). Change the case to one in which there are two threats (at different times) and we apprehend one suspect for each threat. Our intelligence is hardly infallible, and it is certain that one of these people is innocent; torture against the innocent one will not “work,” and one of those threats will, therefore, be actualized. Should we torture in both cases if such torture ensures the elimination of only one threat? As I have argued elsewhere, I think that the answer is yes, so long as the other moral considerations are of sufficient magnitude.<sup>1</sup>

Regardless, my contention is that the fact that torture does not always elicit valuable information does not rule out the possibility of justified torture. So what

is supposed to be going on when critics allege that torture does not work? As above, nobody said that it always worked: the contention can be only that torture could work often enough that the benefits still outweigh the costs. Ineffective torture increases the costs and moderates the benefits, but, if the benefits are high enough, those costs are still covered. At any rate, this discussion has been broadly methodological as it is not obvious what the (sometimes) inefficacy of torture really could mean for the moral debate. If the point is just meant to be that the moral calculus is more complicated than ticking-time-bomb cases suggest, then I welcome it. If it is somehow supposed to be that torture is, therefore, unjustified, this does not follow. In either case, let us now move on to some of the specifics.

Darius Rejali poses the key question as this: Can torture be precise, scientific, and professionally administered and yield accurate information in a timely manner? He then breaks out this question as comprising the following component questions: Can torture be scientific? Can one produce pain in a controlled manner? Does technology help torturers in this respect? Can pain be administered respectfully and professionally? Can interrogators separate deceptive from accurate information when it is given to them? How accurately do cooperative prisoners remember information after torture? Does this investigative method yield better results than others normally at an army's disposal? If not, does this investigative method yield better results under conditions of constrained time? In the following discussion, he contends that the answers to each question "give no comfort to advocates of torture, no matter how they qualify the questions. Apologists often assume that torture works, and all that is left is the moral justification. If torture does not work, then their apology is irrelevant. Deciding whether one *ought* or *ought not* to drive a car is a pointless debate if the car has no gas."<sup>2</sup>

I refer the interested reader to Rejali's ensuing discussion, all of which is carefully researched and quite compelling. My strategy is not to rebut any particular points that he makes but rather to deny that his conclusion follows from his premises. Before doing that, it will do some good to at least get a sense of some of the discussion on offer just so that we have a shared starting point. First, Rejali contends that pain is a highly complex phenomenon and implies that academic discussions of torture fail to understand these complexities.<sup>3</sup> For example, he denies that pain is a simple metric that can just be "increased" as the torturer needs: bodies are desensitized to pain; increases in pain stimulus portend no straightforward relationship to the pain actually experienced by the victim; humans differ unpredictably in their ability to endure extreme pain; there are various pain thresholds, all of which bear complicated relationships to each other (e.g., of sensation, of recognition as pain, of tolerance, of endurance); and so on. His ultimate conclusion is that:

... the notion of a science of torture rests on simple folklore about pain. This folklore teaches that all people avoid pain and seek pleasure, more injury produces more pain, and so it is simply a matter of calibrating the quantity of pain for each individual. These views do not make any sense in torture. The people interrogators most want to question are also those most likely to embrace and resist pain. More injury often produces less pain, especially over the course of an interrogation. Pain is not an undifferentiated sensation that is amenable to a scale. If there turns out to be a science of torture one day, it will look nothing like common folklore imagines it.<sup>4</sup>

Again, let us grant everything that he has said. Torture is messy, and we cannot calculate *ex ante* how much pain will be needed to break some particular terrorist, much less how to deploy that quantum of pain. Surely, torture would be an easier affair with such knowledge, but it hardly follows that the lack thereof makes torture impermissible. Torture engenders pain, and—especially given time constraints—pain could be the most efficacious way to bring about the disclosure of lifesaving information. Rejali seems to think that a science of torture is necessary for its moral justification, but I simply disagree. All that matters is that some terrorist will not surrender the information absent the threat of torture coupled with the corollary that torture might lead to the release of that information; surely both these premises are straightforwardly true. As the associated probabilities fall (cf. “might”), then torture becomes less advisable. But, given the potential value of the intelligence, I submit that they do not even need to be that high in the first place.

Rejali’s third and fourth questions are less central than the first two. His discussion of technology effectively just reiterates his point that torture is not scientific, which we have already discussed. If technology is just meant to offer new and creative ways to inflict pain (e.g., magnetos and stun guns), then it does not change the basic structure of the torture debate. Some technologies, however, really do have the potential to be transformational. Consider, for example, functional magnetic resonance imaging (fMRI); fMRI might be usable as a polygraph to determine whether those under torture are providing misinformation, either to mislead or else just for reprieve from interrogation. Although a research agenda whose goal is to determine whether the ordeal of torture compromises the efficacy of these diagnostics would hardly get approval from institutional review boards, we can surely suspect that various government agencies are quite interested in the answer.

Can pain be administered professionally and respectfully? With many lives hanging in the balance, I am not sure that this question should be taking priority. Nevertheless, Rejali raises at least two important points that deserve mention: torture leads to brutality and deskilling.<sup>5</sup> The key here is *leads to*, which identifies a different focus in Rejali’s book than mine. In particular, he is deeply concerned with pervasive, institutionalized torture and spends much of the book discussing problems thereof. By contrast, my focus is on exceptional cases rather than normalized ones. This bears notice, though I will talk about it more in §2 and §5.<sup>6</sup> For me, torture does not “lead to” anything substantial because the cases are too infrequent. For example, Rejali worries that rivalries between interrogators will lead each of them to try to outdo the other or else that coercive interrogation undermines other professional skills (e.g., why do the fingerprinting and other processing when you have a bat?).<sup>7</sup> But imagine that torture is rare; for the sake of imagination, suppose that the U.S. government were to deploy it against ten suspected terrorists per year. If Rejali’s worry that police forces will forget how to fingerprint because they just go ahead with torture is taken seriously—and I doubt that it could be—then it is a nonstarter because the (in)frequency of torture cannot get the worry off the ground. An analogous point could be made about brutality or any of the other systematic concerns he has in the associated discussion.

The fifth and sixth questions again get to the nuts and bolts about the efficacy of torture: Can interrogators spot the truth, and how well do tortured victims remember relevant information? For torture to be effective, we want the victims to give accurate information, and we also want the interrogators to be able to

recognize misinformation. Rejali argues that interrogators barely do better than chance in identifying lies and, in some cases, do worse.<sup>8</sup> But what is supposed to follow from this? Imagine two cases, one in which a terrorist is not tortured and chooses not to disclose the truth and one in which the terrorist is tortured and reveals either good or bad intelligence, though we do not know which. The problem with Rejali's argument is the baseline: it is bad that we might get misinformation, but that does not tell the whole story. Rather, it is better that we *might* be getting actionable intelligence than that we surely get no intelligence at all. There could be costs of bad intelligence, such as when we move our security forces to the wrong place, thus wasting resources. But, even if we incorporate those risks into our intelligence calculus, the options are wasting some resources and having a chance at saving lives or else not having a chance at saving those lives at all; the former could easily come out as superior to the latter.<sup>9</sup>

The complementary worry to intelligence officers not recognizing the truth is those under torture being unable to offer it. For example, imagine that we are after the details of some attack, the coordinates of which happen to be 48.41° north latitude and 114.34° west longitude. That is a lot of numbers to remember, and there is no doubt that torture would make it harder to produce them; surely, torture is traumatic, degrades the recall powers of its victims, and so on. But, if the terrorist would have had access to that information in the first place, he surely would have known that the target was Whitefish, Montana, so the exact coordinates are really irrelevant, and the city's name is simple enough to remember. Would he be able to remember whether the bomb was at 123 Central Avenue, 213 Central Avenue, or 312 Central Avenue? Maybe not, and this would leave local law enforcement with some work to do. As it turns out, however, the 100 block is much more crowded than the 300 block, so they might reasonably suspect the ill-intentioned terrorist to have planted the bomb on the former rather than the latter. Nobody suggested that torture was infallible, just that it might, in some cases, be our best option. And, regardless, if the options are the choice between a tortured terrorist scrambling some of the details and letting people die, the former could well be the way to go. Finally, pace Rejali's worries about the science of torture, it is at least the case that some forms of torture will be more or less psychologically debilitating than others or else that reasonable recovery periods will allow for psychological reconstitution. Without presupposing some perfect system in this regard, it can, nevertheless, be recognized that we will torture somewhat better or worse in regard to this metric by recognizing the mental degradation of those under torture.

Finally, Rejali wonders whether torture is comparatively advisable given alternatives available to the intelligence community, whether in general or when time is a factor. More broadly, he thinks that public cooperation through the recruitment of informants is the most successful way to combat terrorist threats;<sup>10</sup> this is true in virtue of the low costs (e.g., no torture), higher reliability of information, and so on. This is fine as far as it goes, and I certainly agree that we should try to recruit informants as opposed to torturing rapaciously. That said, this line is silent about the cases wherein we do not have any informants and time is too short to find any. In response, Rejali argues that torture would not work under time constraints anyway since "real torture—not the stuff of television—takes days, if not weeks."<sup>11</sup> Stress and duress tactics, for example, take a while to elicit information, should they ultimately be effective at all. But, surely, there are myriad forms of torture that hardly take any time at all (e.g., electroshock, beating), so Rejali's comments

are curious. Regardless, the point needs to be only that torture operates on a faster timetable than the alternatives and that we can reasonably suppose it to be more expedient than building up a network of informants. Torture does not need to be instantly effective for it to be better than the alternatives.

Whether we torture—as opposed to doing something else—depends on the nature of the threat. Focusing on the temporal element, let us say that torture either does or does not stand in the relevant time frame with regard to the threat. For example, following Rejali, suppose that torture (usually) takes days to generate actionable intelligence. If the bomb is going to go off in days (as opposed to hours), then should we torture? A few days could be enough time to give us at least some confidence that torture would work while giving us less confidence that its alternatives would (e.g., building informant networks). Therefore, torture could be the best option. Alternatively, suppose that the bomb will go off in hours rather than days and that torture is less effective on this timetable. Should we torture? Supposing that we lack meaningful alternatives, we are sadly back in the position of torturing such that we might save lives or else not torturing and ensuring that we do not. Assuming that there is some reasonable chance that the torture would work, then I prefer the first option. It is worth emphasizing that nothing anyone has said denies that this possibility is reasonable; Rejali's own writing just suggests the possibility to be far from certain. The probabilities obviously matter, but there is nothing disingenuous in supposing those probabilities to be non-negligible such that torture could still be justified when there are enough lives in the balance.

There are various reasons to be skeptical about the certainty that torture will elicit actionable intelligence. In ticking-time-bomb cases, this certainty is stipulated, but, in the real world, things are much more complicated. Rejali offers an important contribution in articulating various reasons to doubt the efficacy and reliability of actual torture, but my contention is that nothing he says changes the fundamental moral calculus or shows torture to be categorically unjustified. Again, his target is principally that of pervasive, institutionalized torture, but that is simply not the sort of torture that I defend. Even in exceptional applications, I can allow that torture does not always work, but, first, nobody ever said it did, and, second, it does not need to *always* work in order for it to be morally justifiable in *some* cases. Having now discussed this extension of ticking-time-bomb thinking to the real world, let me now turn to institutionalization.

### **Torture Requires Institutions**

A second class of criticisms against the real-world implementation of torture invokes the failure of ticking-time-bomb cases to recognize essential institutional infrastructure. Ticking-time-bomb cases have it that we just “go torture” the terrorist. But how? Where? Who inflicts the torture? Where did this person learn his craft? Is there oversight of the process? If so, by whom? And so on. The complaint is that torture does not just appear *ex nihilo* but rather requires the existence of various institutions. Once we recognize these institutional requirements, then the costs side of the torture ledger starts to look a lot more pronounced. No longer are we simply absorbing the one-off torture of a terrorist, but rather we must acknowledge that such torture is possible only if a wide range of other costs are borne, whether social, economic, or moral.

Central to these criticisms is the contention that the institutional costs exist even if torture is rare; should that not be the case, then the criticisms cannot indict the

position that torture is licensed only in exceptional cases. For example, consider Jean Maria Arrigo, who makes this point explicit:

The use of sophisticated torture techniques by a trained staff entails the problematic institutional arrangement I have laid out: physician assistance; cutting edge, secret biomedical research for torture techniques unknown to the terrorist organization and tailored to the individual captive for swift effect; well-trained torturers, quickly accessible at major locations; pre-arranged permissions from the courts because of the urgency; rejection of independent monitoring due to security issues; and so on. These institutional arrangements will have to be in place, with all their unintended and accumulating consequences, *however rarely* terrorist suspects are tortured.<sup>12</sup>

Let us set aside the issue of training for torturers until later; enough people are worried about these sorts of institutional requirements that dedicated discussion is warranted. In general, however, my response to the institutionalization objections is twofold: first, I deny that we need as much institutional apparatus as critics allege, and, second, even if we did, we can either accept the institutions as a necessary cost or else make do as well as we can without them should those costs be too high. This second response is more broadly theoretical, however, so let us start with Arrigo's arguments. In doing so, I propose to use a paper other than the one from which the quote given above derives; this other paper, coauthored with Vittorio Bufacchi, more neatly designates the institutions allegedly in play.<sup>13</sup> In particular, Arrigo and Bufacchi maintain that torture requires the complicity of the following establishments: medical, scientific, police and military, and legal.

There is little doubt that medical personnel have been used to develop more efficacious interrogation patterns. For example, psychologists at Guantánamo were organized into "behavioral science consultation teams" (BSCTs, or "biscuits") that advised on interrogations. Medical personnel have, presumably, also been used to resuscitate those suffering the ills of hostile interrogations, sometimes such that those interrogations can be thereafter resumed. I have written elsewhere on the moral status of these collaborations,<sup>14</sup> but the present discussion is, instead, simply about institutions rather than medical ethics. Are medical personnel *necessary* for torture? In a straightforward way, the answer is obviously no: the practice of torture does not require—either logically or pragmatically—the presence of medical personnel. The fact that medical personnel have been complicit in torture does not show that the medical establishment is essential for torture any more than the use of a book to prop a door open shows that the book was essential to keep the door from closing.

The relevant issue is not what is required for torture but rather what is required to torture "well." Even granting that medical personnel increase the efficacy of torture, it does not follow that torture would not be reasonably effective without them. Were some suspect to be rendered unconscious by an overzealous interrogator, actionable intelligence would be less likely to be ascertained in an expedient way. While this gives us reason to want some medical personnel on hand to revive the detainee—either for his health or for our ends—it hardly means that an entire medical *establishment* is needed to ensure the efficacy of torture. Imagine, for example, that we have somewhere on the order of ten torture sites set up around the world, and that we want each of those sites to be staffed with a few medical

personnel (physicians, psychologists, etc.). If we assign, say, three medical personnel to each site, then this amounts to only thirty medical personnel *worldwide*. Not much of an institution. Allowing these thirty to collaborate (e.g., through teleconference and/or videoconference) ensures that the medical bases can be adequately covered. Or, to the extent that they are not, more will be said below.

The second institution identified by Bufacchi and Arrigo is scientific: some group needs to develop the techniques that will be used for torture, particularly technological correlates thereof. In §1, for example, I suggested that fMRI might be used to ensure the honesty of the suspect under torture, but there are a range of problems with this proposal. For one, the duress of torture might compromise the integrity of fMRI results, though we would not know whether that is true until some scientist figures it out. And it could be even worse than that insofar as institutional oversight would have to license this research, thus further multiplying the resources that would, ultimately, be required to support torture.

There are a few things to say in response. First, much of the scientific research that ultimately contributes to torture need not originally be undertaken for that reason. Tasers, for example, were first developed for police as a nonlethal way to subdue aggressive suspects; now they provide one sort of off-the-shelf technology that can be used in torture.<sup>15</sup> In this case, it would be inappropriate to account for the associated researchers on torture's institutional ledger since their research had nothing (directly) to do with torture. It is not as if these people would have been off developing cures for malaria, cures that were, ultimately, held up for research into torture. That scenario would reasonably give us pause since limited institutional resources could have been used in some other (and potentially better) way, but the *disanalogy* with the Taser is straightforward insofar as Taser research was meant to serve as a legitimate aid in law enforcement. If that research was later deployed for use in torture, it was done so without any additional cost and, indeed, without any (direct) support from the scientific community.

The second and third responses are ones that will reverberate throughout this discussion. The former holds that the institutional costs could be justified, in which case the objection just has us adjust our balance sheets without changing our conclusions. So, while the Taser research might be freely appropriated for torture, it could be hard to determine whether fMRI results would break down given the duress of torture unless we specifically assigned researchers to study that (i.e., to directly involve scientists in torture research). If the advantages of this sort of research are high enough, then the associated costs could be justified. If the advantages are more moderate, then we might just have to forsake them given the costs. More on this below.

Next, Bufacchi and Arrigo point to police and military infrastructures, all of which would have to be reconfigured to support torture; there could also be further psychological effects on those demographics who directly participate in torture. Throughout their article, the authors conflate the notions of *exceptional torture* and *normalized torture*, though that conflation bears particular emphasis in this regard. For example, they write: "Police departments have struggled for decades with the overwhelming bad consequences of coercive interrogation."<sup>16</sup> But I am not proposing that police departments torture or even that the entire military establishment be involved. Surely there are high institutional costs on either of these proposals, which is why nobody has ever made them. Again, we are considering torture under rare and exceptional cases, not as some sort of quotidian practice in every precinct or on every military base.



Something should be said more in the way of specifics, however, and this would be a good time to offer some discussion in that regard. Remember that Arrigo's earlier work postulated the need for torturers to be "quickly accessible at major locations."<sup>17</sup> This sounds right; note that such a commitment is a far cry from having torture extend throughout all levels of law enforcement and military. As suggested above, we might countenance ten U.S. torture sites throughout the world. If three of these were domestic—one on each coast plus one around St. Louis—then no apprehended suspect would be more than two to three hours by plane from the nearest site. The other seven sites could be strategically placed around the world in ways sensitive to security needs (e.g., a couple in the Middle East, none in Oceania). If the local circumstances do not provide for such a site—which we both could and have run as black sites (e.g., secret prisons)—then arrangements could be made on carriers in international waters. Such sites would be fully secured, and the U.S. military would stand to defend them. All told, this hardly adds up to the extensive institutionalization that critics decry: there are a small number of sites with correspondingly limited personnel requirements.

Bufacchi and Arrigo go on to discuss legal institutions, but I will return to that topic in the next section and, therefore, defer it for now. However, before moving on to more general remarks regarding institutional requirements for torture, the issue of torturer training should be addressed; this is one of the most common institutional requirements lamented by critics of torture. For example, consider Jessica Wolfendale, who contends: "The scope and training to produce the torturer needed in the ticking bomb scenario raises serious questions about the legitimacy of these kinds of arguments for the use of torture. . . . [T]his training cannot be neatly contained within the parameters of the ticking bomb scenario because permitting torture in these cases requires *already* having a well-established training regime for torturers."<sup>18</sup> Similar is Henry Shue: "Torture is not for amateurs—successful torturers need to be real 'pros,' and no one becomes a 'pro' overnight. At a minimum, one must practice—perhaps do research, be mentored by the still more experienced. In short, torture needs a bureaucracy, with apprentices and experts, of the kind that torture has in fact always had. . . . Torture is an institution."<sup>19</sup>

It is most certainly false that torturers *have to be* "real 'pros'"; the history of torture is rife with examples of anything but. Torturers are often far from professional, particularly in any robust sense of the word: they do not have professional associations, promulgated codes of ethics, or any of the other hallmarks of professionalism. The only way that this claim could be plausible is if it were read to mean that the best, most idealized torturers were ones who had a lot of practice, training, and mentorship. But maybe we simply cannot have torturers of that sort. Given the rarity with which I would propose torture, a lot of these training opportunities will simply not be available, though there might be approximations thereof (e.g., recordings, simulations, etc.).

My real objection, however, has to do with this dichotomy between trained and untrained torturers, which I take to be spurious. Wolfendale, for example, writes: "It would not do to take an ordinary soldier and make him torture a terrorist suspect at the last minute. One has only to look at the incompetence of the guards at Abu Ghraib (they took *photos*) to see the danger of allowing mere amateurs to torture prisoners." But this saddles her opposition with the completely implausible position that ordinary soldiers or prison guards are all that the intelligence community has access to without creating a cadre of professional and dedicated torturers. Of course,

we should not have privates or reservists (cf. Lynndie England) doing this sort of work, but nobody said that we should. Ironically, Wolfendale acknowledges this on the next page: “In the real world, most torturers are soldiers or military policemen who have been trained in elite units.”<sup>20</sup> This certainly sounds more promising and is a far cry from Abu Ghraib guards. Torturers should, therefore, be drawn, not from the general enlisted population, but rather from special forces like Delta Force, Green Berets, or Navy SEALs. These special forces are trained in relevant ways, including interrogation and interrogation resistance; Navy SEALs are even subjected to simulated drowning (cf. waterboarding) as part of their training. In other words, these units know what a hostile interrogation looks like, so they would not be starting from scratch if imported into such settings.

None of this is to say that a wholly dedicated class of professional torturers would not be more effective than whatever other group we co-opted into that role. Rather, the point is that the distinction between “trained” and “untrained” torturers is a nonstarter insofar as our torturers would not be randomly selected from a nearby street corner; rather, they would be highly trained individuals who already possess an allied skill set, extreme mental fortitude, and a dedication to their craft and country. Psychological profiling could further identify the over- and under-zealous, thus giving us a starting group. Couple its potential with on-the-job training and experienced personnel already in the field, and I am far from skeptical about the prospects for effective torturers. To return to the original objection, the point is simply that we do not need dedicated torturer training programs or even a significant number of torturers. Some of these critics seem to envision some sort of “torture college” with things like students, faculty, academic advisers, secretarial staff, administrative support, and so on, but something far more minimal would go a long way toward effective, if not optimal, torture.

More generally, I am skeptical as to how much training is actually required for torture. In some sense, the answer could pretty clearly be none. For example, we could just hand an interrogator a baseball bat and see what happens; the chance of getting actionable intelligence from someone who has it is surely greater than zero. The point is not that this is what we should do—to be clear, I do not propose it—but rather that the issue is how *effective* we want our torture to be. Institutions might make torture more effective, but none is essential for torture. And this is the central point missed by those who raise institutional objections against torture, whether in regard to the training of torturers or otherwise.

Consider how institutional investment in torture would improve outcomes (namely, the retrieval of actionable intelligence). Or, before doing that, consider how any sort of investment would lead to improved outcomes. Roughly, there is a diminishing return on each subsequent unit of investment: the first unit of investment will lead to a large improvement, and each unit of investment thereafter engenders smaller gains. And the aggregated investments do not portend infinite advance of the desired project but rather asymptotically approach some maximum. There are probably irregular processes that eschew this simple model, such as when our first hour on some task gets us nowhere and then we suddenly see the light during the second hour. Nevertheless, as a general rule of investment and results, something like this is approximately right.

I submit that something similar would go on with institutional investment in torture insofar as more investment leads to better outcomes, but with diminishing results. Consider the abstract metrics of institutional investment and the efficacy

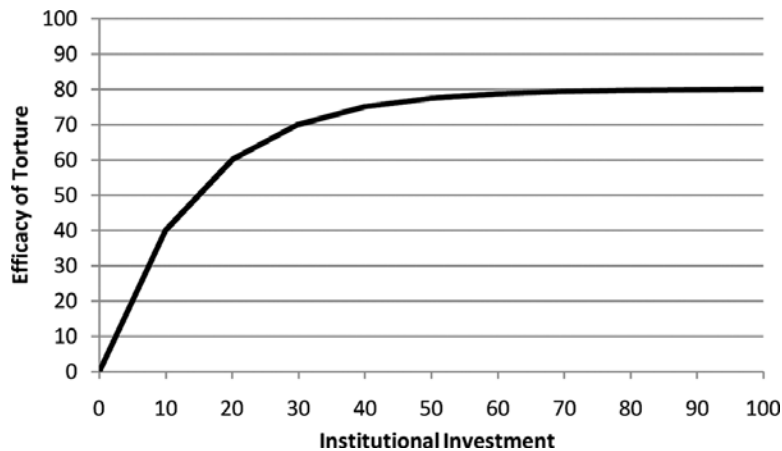


Figure 1. Efficacy of torture versus institutional investment.

of torture such that higher numbers indicate greater investment and better outcomes (i.e., more actionable intelligence). If we plot that relationship on a generic (and non-percentage-based) scale, it would look something like Figure 1. The critics considered so far in this section allege that torture requires institutions, but my response has continually been that, while (maximally) effective torture might, torture *simpliciter* certainly does not; this chart makes that point visually. Our solo torturer with the baseball bat would be on the left side, and the sorts of institutionalized torture envisioned by Bufacchi, Arrigo, Wolfendale, and Shue would be on the right side. There is no reason to move toward their sort of institutionalization if the outcomes fall off as this chart holds simply because such investments are paying such a poor return *vis-à-vis* outcomes.

Of course, the details matter, and I offer the figure to make a conceptual point rather than a practical one. For example, if the payoffs looked like this, I might not propose institutionalization past the “10” level since the institutional costs have to be reflected in the ledgers. If those costs are high enough, then we simply have to do without the associated institutions since our moral calculus cannot cover them. This is not to say that we will be diminished to interrogators with bats, but the torture college might well be off the table. What ultimately matters is that we not institutionalize torture beyond what our moral calculus supports; similarly, if torture becomes less effective in the process, then our moral calculus has to recognize that as well. The key conclusion, however, is that this less effective torture might still be preferable to overwhelming institutional costs being incurred. As I have argued throughout this section, however, I am more sanguine about institutionally minimal torture than others.

### The Nefarious Spread of Torture

Ticking-time-bomb cases ask us to consider torture in isolation; our consent to torture in those cases has no implications for the practice of torture in other cases. Critics complain that this sort of isolation does not exist in the real world and that the legitimization of torture establishes a dangerous precedent. Once we are willing to torture in some cases, we will be willing to torture in others.<sup>21</sup> Eventually, those cases might bear little resemblance to ticking-time-bomb cases, so we best never embark on this slippery slope. Parallel to this expansion of torture, there are worries

about how torture brutalizes the torturer or otherwise affects the psychologies of those involved with torture. As these people gain exposure to torture, they might become desensitized toward it, which could lead to causing more, or less discriminating, pain. Or maybe the normalization of torture will lead to the normalization of other immoral acts.<sup>22</sup> All of these concerns are effectively that the allowance of torture in isolated cases will have negative downstream consequences. The most substantial worry is the one mentioned first, that torture in exceptional cases threatens its normalization. Shue makes this point eloquently:

Rousseau says at one point that pure democracy is a system of government suitable only for angels—ordinary mortals cannot handle it. . . . As devilish as terroristic torture is, in a sense it too may be a technique only for angels; perhaps only angels could use it within the only constraints which would make it permissible and, then, lay it aside.<sup>23</sup>

So imagine that ticking-time-bomb cases show us that torture is permissible in cases that satisfy various conditions,  $C_1, C_2, C_3, \dots, C_n$ . Further imagine that we meet a case wherein all but one of those conditions is satisfied (e.g.,  $\sim C_1, C_2, C_3, \dots, C_n$ ). One worry is that we deem the latter case to be “close enough” to the former and go ahead and torture anyway. Another worry is that we simply cannot help ourselves, such are the pleasures of torture. In other words, the first worry portends a moral progression and the second a psychological one. Let us take these in order.

Regarding the moral progression, it obviously matters which conditions are (not) satisfied. As I have argued elsewhere, ticking-time-bomb cases are not meant to put necessary conditions on the application of torture; rather, those conditions are meant to be sufficient.<sup>24</sup> If torture is justified under conditions  $C_1, C_2, C_3, \dots, C_n$ , then it does not follow that torture is not justified under  $\sim C_1, C_2, C_3, \dots, C_n$ . Assuming for the sake of argument that torture is justified in ticking-time-bomb cases, the formal worry cannot be that such cases would lead to unjustified torture in the real world since it would still be an open question whether torture could be justified in cases sufficiently resembling—yet nevertheless different from—ticking-time-bomb cases.

What matters is simply that lesser harms can be effectuated through torture than through its absence. For this reason, there is nothing inherently wrong with the spread of torture, regardless of how many conditions are relaxed from ticking-time-bomb cases. Can torture be justified if the conditions of the case are not identical to those of ticking-time-bomb cases? Per the red sweatshirt case in §1, I think that the answer is yes, though I also suspect that the relations need to be reasonably close. To put it another way, nothing mostly or completely different from a ticking-time-bomb case has much going for it *vis-à-vis* the justification of torture.

Returning to the critic, what is wrong with the spread of torture? Most simply, nothing: the problem could be only with the spread of *unjustified* torture. If our allowance of justified torture leads to the spread of more (justified) torture, then this has to be a good thing. What we need to worry about is, not simply the spread of torture, but rather how the use of *justified* torture could lead to the spread of *unjustified* torture. Certainly, this is something that warrants discussion, but, from the outset, it is worth asking why anyone would practice unjustified torture. For example, imagine that we could torture ten people—nine of them innocent, but we know not which—in order to have a small chance of saving one other. Is this worth it? Filling in more details as we like, the answer can easily come out in the

negative. Are we seriously worried that torture in ticking-time-bomb-like cases will lead to torture in a case like this one? Nobody would seriously propose such torture.

Getting clear on exactly what the objection amounts to is, therefore, an important part of the dialectic. From a social-political standpoint, I propose that the prolific spread of unjustified torture is not something that we need to worry about insofar as there is no advantage to pursuing such torture and, furthermore, that there are myriad disadvantages. This has to sound naïve given the fact that torture has become normalized under some regimes; consider, for example, Argentina and Chile.<sup>25</sup> But these are *radically* different situations wherein dictators are trying to suppress their opposition. The sorts of torture that we are countenancing are, not those of political suppression, but rather those that interfere with terrorists' threats. Is the dialectic really supposed to be that the torture of a few (suspected) terrorists would lead to the widespread torture of a Pinochet-scale thirty thousand? That is radically implausible.

We have yet to get the proper point on the objection that torture will spread. Just to be charitable, here is a scenario that we might reasonably countenance. Imagine that we license torture in a ticking-time-bomb case and that the use of torture saved many thousands of lives. Another case thereafter arises that is far less clear-cut: we round up a group of people, most of whom are likely to be innocent. Only one hundred people are threatened. Intelligence reveals that the threat is hyper-imminent, such that torture may not be effective anyway (i.e., the bomb might go off before we would be able to disarm it). Should we torture? Maybe not. Would we? Maybe we are overly enthusiastic because of our recent success and want to move forward even if the moral calculus is compromised. What protections are in place to prevent the overextension of torture to inappropriate cases? Will we be able to restrain ourselves?

Contra Shue, I am more optimistic: we just owe an honest costs and benefits assessment. Why would that be so hard? Imagine that the military deployed a bomb against a terrorists' meeting, killing several terrorists and, therefore, accomplishing important strategic ends. It is possible that the next meeting takes place under the cover of a wedding and that the success of the first attack threatens judicious deliberation in the second. Does that mean that the first bomb should not have been dropped? Of course not. It means that the second bomb should not be dropped unless the collateral damage can be justified. During wartime, military officers have to make just these sorts of decisions; there is a lot of pressure on those officers to make the decisions well. Whoever authorizes torture—and at whatever level—will be under similar pressure and accountability.

Thus far, I have argued for various points. First, the spread of (justified) torture is not necessarily bad. Second, that the torture of terrorists could spread to anything close to the magnitude of torture as practiced by the worst dictators is dubious. Third, vigilance is required to prevent the spread of torture from justified to unjustified cases, though such vigilance is already required—and practiced—in various other settings wherein we might reasonably worry about slippery slopes. A fourth point, that there should be various legal safeguards in place to prevent abuses of torture, will be developed in the next chapter. The scandal at Abu Ghraib, for example, was just that, and those involved should be punished. Such punishments should be harsh, and they should be publicized. The spread of torture will be further curtailed by proper oversight and accountability; again, more on that in the next chapter.

This response paves the way for the next type of objection, which is not so much about the spread of torture across cases as about the worsening of detainee treatment within cases. In other words, we could worry that the admission of any torture threatens extension to unjustified cases, as discussed above. But another worry is that torture will be more malicious than it needs to be once torturers are desensitized toward its practice. Given studies like the Milgram experiment<sup>26</sup> or the Stanford prison experiment,<sup>27</sup> this is hardly a worry that can be easily dismissed. But the principal response is that overly zealous torture should not be tolerated and, furthermore, that it should be punished. Torturers need to exercise due restraint, and, insofar as they fail to, they should be removed from their positions and held accountable.

Important work has been done regarding the psychology of torturers,<sup>28</sup> and I can only imagine what psychological fortitude it must take to torture another human being. Furthermore, there is somewhat of a conundrum here insofar as the most effective torturers might be those who are sufficiently desensitized but, at the same time, those might be exactly the sort of people we do not want practicing torture for fear of abuse. I suspect that the best way to go would be to have torturers rotate out after tours of only a few years; this could promote a balance between desensitization and experience, both of which run at cross-purposes (i.e., we want one low and the other high, but they rise together). Again, however, I recommend that standards be put into place and that they be enforced. As will be argued in the next section, this can be done without compromising some of the institutional desiderata from §2.

### Alternatives to Torture

Everyone agrees that torture should not be practiced if we can accomplish the same outcomes—such as the retrieval of lifesaving information—in less morally costly ways. Do such alternatives exist, and, if so, what is their implication for the moral status of interrogational torture? There is the straightforward sense in which alternatives to torture exist insofar as we can simply choose not to torture, instead choosing to do something else (or nothing). It hardly follows from this, however, that there are meaningful alternatives to torture *vis-à-vis* successful intelligence gathering since nothing has yet been said about the efficacy of those alternatives. There are various ways that this dialectic could go, at least two of which are interesting: first, we could say that alternatives to torture are both more effective and less morally costly, thus making them superior; second, we could say that the alternatives are less effective and less morally costly but that these moral savings are worth recouping.<sup>29</sup> Regarding this second, the point is that we could allow the alternatives to be less effective while still preferring them insofar as they lack the associated moral hazards of torture. I have read most of the literature as endorsing this first proposal, though the second is certainly a live one as well. In what follows, I shall not maintain the distinction, but it bears acknowledgment.

So what are some of the alternatives to torture? In a recent book, Matthew Alexander tells the story of the search for Abu Musab al Zarqawi, the former leader of al-Qaeda in Iraq; Alexander was one of the interrogators responsible for the intelligence work that led to al Zarqawi's death.<sup>30</sup> Throughout the interrogations, Alexander favored non-hostile techniques, and the book is meant to be a celebration of the success of those techniques insofar as al Zarqawi was eventually killed

through intelligence acquired without torture.<sup>31</sup> The techniques that Alexander advocates play off the psychological states of detainees (i.e., they are not physically threatening); these are drawn directly from the *Army Field Manual (AFM)*.<sup>32</sup> While that list in the *AFM* is somewhat longer, the techniques that Alexander emphasizes include love of family, love of comrade, fear-down, futility, and rapport. The first two invite the detainee to countenance a life without his family or friends or else suggest that loved ones are at risk and can be protected only if the detainee cooperates. Fear-down mitigates some preexisting fear of the detainee (e.g., by promising safety from it). Futility is meant to show that the detainee's cause will be frustrated regardless and that he might as well draw favor by facilitating expedient countermeasures. Finally, rapport consists in building a positive, trusting relationship with the detainee. Alexander clearly favors rapport, though he concedes that more intransigent detainees might force another approach if rapport is clearly off the table.

Certainly, Alexander thinks that the aforementioned techniques should be preferred to torture, but he allows for other psychological interventions that are hardly benign. For example, he approvingly cites the following from the *AFM*: "Limitations on the use of [violence or intimidations, including physical or mental torture or exposure to inhumane treatment] should not be confused with psychological ploys, verbal trickery, or other nonviolent or non-coercive ruses used by the interrogator in the successful interrogation of hesitant or uncooperative sources."<sup>33</sup> In other words, even though torture and inhumane treatment are explicitly proscribed, there are various forms of deception that are legitimate interrogational tactics.

Of course, nonviolent or other non-coercive means could be preferable to torture, but two important disclaimers should be registered; unpacking them will make Alexander's case far less compelling. First, whatever these alternatives amount to, they must be less morally costly than torture. Once we start looking at how some of these cases unfold, the supremacy of non-torture will hardly be obvious. Second, these alternatives need to be effective, particularly given time constraints. If they are far less effective than torture—even in a small subset of cases—then a strong argument against torture has not been made (in those cases). Imagine that someone says, for example, that some action  $\phi$  is less morally costly than  $\psi$ , but, as it turns out,  $\phi$  is far less effective than  $\psi$  in realizing some end,  $E$ . Should we  $\phi$  or  $\psi$ ? As suggested at the beginning of this section, it depends on the relative moral costs, moral benefits, and degrees of efficacy.<sup>34</sup>

The first point is that alternatives to torture are hardly benign, at least insofar as they are likely to be effective. This contention trades on two different variables, so let me illustrate by example. One alternative to torture is just to extend gracious hospitality to our detained terrorist, hoping that he will reciprocate by helping us disarm his terrorist threat. Quite benign, but not very likely to be effective. From there, rapport is less benign—for example, it is disingenuous and, therefore, not wholly morally innocuous—and probably more effective. Very effective? I certainly doubt it—especially since the interrogating personnel quite probably represent those whom the terrorist is ideologically predisposed to despise—but there is no doubt that rapport can work. But for the present argument, however, all that matters is that rapport does not *always* work, which can hardly be controversial (cf. time constraints). For emphasis, though, let me make plain that rapport should be pursued when it is reasonably likely to work or else when it is at least as reasonably likely to work as torture.

Now that we have cut the logical space down to cases in which rapport will not work, let us look at the effects of Alexander's alternatives to torture. What bore notice to me in reading his book was that these were far from docile. For example, he approvingly notes that a fellow interrogator "set the air conditioner on the lowest setting" and that it was "chilly" during an interrogation. Elsewhere he admits that a detainee "looks tired. He's got bags under his eyes. His mouth is sagging. He doesn't seem as alert as he did yesterday." He once tells a detainee: "Put your mask back on. . . . His hands tremble so badly that I have to help him get it over his head." After a tumultuous interrogation of a man who turned out to be *innocent*, Alexander admits: "I psychologically savaged him."<sup>35</sup>

While I have immense gratitude for Alexander's service, I respectfully contend that the distinctions between these treatments and torture are far more tenuous than he contends. Philosophically—if not legally—does it really make a difference that he never beat a detainee? To my mind, not really; in fact, torture might have caused less overall suffering insofar as it could have been more expedient. The alternatives to torture that he advocates apparently allow for temperature manipulation, sleep deprivation, and hooding. To my mind, these fall short of torture; nevertheless, they are clearly coercive and, therefore, morally contentious.<sup>36</sup> Beyond the stress and duress tactics, Alexander is willing to subject his detainees to extreme psychological duress by renegeing on false promises, threatening the safety of family members, and so on.

Either Alexander thinks that psychological torture is impossible, or else he thinks that these techniques fall short of it. As argued elsewhere, I am skeptical about whether psychological attacks can satisfy the specific intent requirement for torture (i.e., that the torturer specifically intends "prolonged mental harm,"<sup>37</sup> which is really orthogonal to his goals).<sup>38</sup> That said, psychological attack clearly constitutes a moral harm, and there is no reason to think that such attacks are categorically less offensive than their physical complements. In other words, even if psychological techniques are not torture, it does not follow that they are necessarily better. On the other horn of the dilemma, Alexander could acknowledge the possibility of psychological torture but deny that the crippling effects rendered by his interrogations reached the relevant threshold. To that, I can only ask: What would? Threatening death?<sup>39</sup> Threatening the indefinite detention of a detainee's family members?<sup>40</sup> Since he does both of these, one can surmise only that Alexander finds them permissible.

The argument thus far portends a red herring: maybe not only is torture impermissible, but so are many of the practices in which Alexander engages (even if they fall short of torture). By arguing that his practices are not morally innocuous, the proper conclusion might be to reject them all rather than to endorse torture. What tools are then left for interrogators? Bufacchi and Arrigo return to rapport, endorsing the "social skills method" of Hans Scharff, a master interrogator of the German Luftwaffe (i.e., air force) during World War II;<sup>41</sup> the idea here is to solicit seemingly unimportant information through cordial treatment and then, ultimately, piece together some substantial revelation. As another example, they praise Islamic clerics who have been successful in "reformulating the religious commitments of some terrorists."<sup>42</sup> Or else maybe chronically ill or badly wounded detainees would become cooperative after being offered medical service. Rejali almost seems to pay a backhanded compliment to Christian Masuy, a Belgian national who ended up working for the German Abwehr (i.e., military intelligence) during World War II.



Masuy's preferred method of torture was to hold the victim's head under water in the bathtub, but with a curious twist. Following the torture,

Masuy's "patient would be conscientiously dried in a bathrobe, warmed up, rubbed with eau de cologne, and consoled with Cognac." Masuy would then praise the victim for his courage, but remind him that it was useless, urge him to come clean, and if not, repeat the torture.<sup>43</sup>

Is Masuy's rapport-building method better than torture that lacks such amenities? Maybe, but this is hardly an *alternative* to torture—nor does Rejali suggest it is—but rather just a variant thereof.

What about the suggestions of Bufacchi and Arrigo? I cannot imagine that "reformulating the religious commitments of some terrorists" is even reasonably expedient, and it must be granted that such an approach would fail on at least some, if not many, terrorists. Medical treatment? Sure, this has some potential, but I suspect the more promising approach is to deny medical treatment—especially for painful afflictions—under promises of its dispensation after intelligence disclosures. Regardless, such considerations can be only a part of our intelligence apparatus insofar as some of our detainees will not be in immediate medical need. That still leaves torture on the table for the rest. So now we are down to Scharff's social skills as the principal alternative to torture in at least some cases. Any semblance of time constraints would severely compromise an interrogator's ability to string together some relevant intelligence finding from sundry information, as would a detainee's unwillingness to communicate at all. It does not hurt my position at all to allow that there are some cases in which these techniques—or the allied ones more directly involving rapport—are superior to torture. Let psychological profiling play an important determination in which way we go as well as assessments by the intelligence community regarding the imminence of attack. Regardless, torture will still be the only viable recourse when we are short on time and/or dealing with a recalcitrant detainee.

A final point brings us back to Alexander: this is an experienced military interrogator who has spent extensive time with important intelligence assets. In his book, he lauds rapport building but acknowledges its limits. Medical generosity never came into play with any of the interrogations detailed in his book, and, certainly, the reformulation of religious commitments was never considered within his operational time frame. Ultimately, I do not see a tremendous moral distinction between the practices advocated by Alexander and physical torture. The proposals made by Bufacchi and Arrigo have lower moral costs than the practices employed by Alexander, yet he does not explicitly endorse them. Why? The obvious answer would be that they are not efficacious given time constraints. But herein lies the rub because neither are Alexander's. As he says of a particular psychological approach:

This will take days, maybe weeks, to develop. And in the meantime, how many innocent Iraqis will die? How many American soldiers will be ambushed and killed? How many suicide bombers will turn marketplaces into bloody bedlam?<sup>44</sup>

We do not always have the days or weeks that Alexander requires. Yes, his techniques netted al Zarqawi, and this should be a regaled accomplishment. How long

did it take? His assignment began sometime in March 2006, and al Zarqawi was killed June 7, 2006.<sup>45</sup> That is over two months and maybe even over three. How many more lives could have been saved were we to have accepted torture as a legitimate tool in our intelligence arsenal?

In closing this section, let me reiterate that torture is not a panacea and that there will surely be cases when we should do things other than torture.<sup>46</sup> In fact, those cases will constitute the overwhelming preponderance of our intelligence activity. Nevertheless, there are limited cases in which torture offers advantages to its alternatives, particularly given time constraints. I suspect that torture has greater efficacy as well; I will remain circumspect on this issue for now but will return to it in the next section. Even if torture effects greater moral harms than its alternatives, it can still be justified if its probability of success is sufficiently better than those alternatives. When time constraints eliminate those alternatives out of hand, even a low success rate for torture could be justified given the threat of a serious harm. Let us now turn to whether such cases ever exist, starting with whether the existence of such cases even matters.

### The Folly of Cases

Critics of torture often propose a dialectic in which they call for a single real-world example of a ticking-time-bomb case and then deny—either rhetorically or substantively—that such a case exists.<sup>47</sup> Given the moral hazards of torture, this methodology is supposed to shift the burden of proof to the advocate of torture, which is not completely unreasonable. Certainly, it would be odd to think that this goes the other way: the critic of torture cannot be asked to prove that no ticking-time-bomb case has ever existed. That said, the cases card is often overplayed, and torture could be justified even in the absence of unequivocal, historically justified cases. There are at least two compelling reasons for this.

First, the contemporary advent of terrorism is much more dangerous than its predecessors. Many of the more noteworthy terroristic events caused the death of, say, one to four hundred people, yet 9/11 caused the death of almost three thousand.<sup>48</sup> Aside from 9/11-like risks, new forms of bioterrorism can also threaten thousands of lives (e.g., contamination of water supplies); the simple point is that terrorism can and has become more lethal than it was in the past. As terrorism threatens more lives, the moral imperative to stop it becomes greater. So, even if we cannot agree on cases in which torture was justified against lesser threats, it does not follow that we would not agree on cases in which torture will be justified against greater threats. The fact that we do not have these latter cases may simply be that the terroristic landscape has continued to evolve and that such cases have not yet presented themselves. In the absence of such cases, there is no harm in thinking of them hypothetically: if terrorism threatened some great moral harm (e.g., one on the order of thousands of lives lost), would torture be permissible? So, even though the relevant cases might not (yet) be on offer, we should not therefore assume that torture will never be justified.

Second, the critic is under a misapprehension if he thinks that successful cases of torture will be publicized. If ticking-time-bomb cases cannot be produced, then there are two conclusions we could draw. The first is that no ticking-time-bomb case has ever existed in the real world. The second is that either no such case has existed or else such a case did exist and we do not know about it. As a simple point of logic, the

second conclusion is more likely to be true than is the first; a disjunction is more often true than either disjunct.<sup>49</sup> And, more substantively, there is good reason to believe that successful cases of torture (i.e., ones that lead to the retrieval of actionable intelligence) will not be widely advertised since there are obvious political costs to this sort of publicity.

By way of example, consider that stress and duress tactics have been widely deployed by the United States in its War on Terror and also that waterboarding has been used in more limited circumstances. Either these practices have yielded lifesaving intelligence, or they have not. Which scenario is more likely? To deny that these techniques have accomplished anything takes a very dim view of our government and military: the allegation would be that our intelligence community has, over a nontrivial period of time, continued practices that were not producing results. What possible reason could there be for this? By analogy, imagine that I went around hitting people with tire irons in the hopes of acquiring the contents of their wallets. Furthermore, imagine that this never worked; after numerous attempts, I had made no economic progress. Finally, imagine that those attempts took place over a protracted time period, over various demographics, with some experimental variations on a core theme, and so on. Given that they were not leading to the desired outcome, would I continue them? Only if I lacked any sort of instrumental rationality or else were masochistic.

Surely, some of its more vociferous critics would proffer precisely these criticisms against the Bush administration, but such criticisms are too often disingenuous and rhetorical. Vice President Cheney, for example, has been adamant that “enhanced interrogation techniques were absolutely essential in saving thousands of American lives and preventing further attacks against the United States.”<sup>50</sup> Critics take this to be hubris, but I am not cynical enough to issue wholesale indictments against our political or military leadership: if enhanced interrogation techniques were continued over some number of years, it would defy credibility that such techniques never produced any positive outcomes.

Consider, for example, the immoderate waterboarding of Khalid Shiekh Mohammed—who was allegedly waterboarded 183 times<sup>51</sup>—which is meant to reveal the excesses of an overzealous Bush administration: if waterboarding were to work at all, it surely would not require this many iterations.<sup>52</sup> Maybe, but maybe not; there is no a priori reason that this should be true. And, if the waterboarding did lead to the stoppage of the “Second Wave” and, in particular, an attack on Los Angeles’s Library Tower, then it would have been justified.<sup>53</sup> The details of this case are uncertain enough that we should not place too much stock in it,<sup>54</sup> but that is precisely the point: we do not always know—if we ever do—the details of the cases because those details are hidden behind public relations considerations, security clearances, and so on. If the details ever do come out, the sources (e.g., disgruntled and *ex*-insiders) do not always inspire confidence.

The argument thus far has simply been that we should not be too greedy in our demand for actual and clear-cut ticking-time-bomb cases. First, some skepticism is warranted insofar as to whether the present and future will resemble the past: the changing face of terrorism makes ticking-time-bomb-like scenarios more likely. It therefore follows that the lack of historical cases has limited implications for their ensuing prospects. Second, even if there has been historical precedent, there are all sorts of reasons to doubt that we would know about it. To wit, there are political and security factors that straightforwardly make it unlikely that actual cases

would enter public consciousness. While some of us believe the Bush administration when it says that enhanced interrogation has saved lives, many will obviously be less charitable. Surely, however, it takes immense cynicism to allege that those enhanced techniques would be continued over several years if they were patently worthless.

I therefore submit to my interlocutors that it is a faulty dialectic in which we ask for real-world ticking-time-bomb cases and then try to draw any substantive conclusions when none is on offer. That said, the proponent of torture in exceptional cases does owe more than a circumstantial and impressionistic argument. While the considerations offered above go some way toward mitigating the burden-of-proof game, they hardly obviate it. Therefore, in the remainder of this section, let us consider some plausible ticking-time-bomb cases. The actual details of these cases are less important than the critics would have us believe for the simple reason that what matters is what *could* happen, not what *has* happened. Regardless, it will be illustrative to consider some actual cases that at least bear relevant isomorphisms to the philosophers' construct. Toward that end, I will offer three cases, one celebrated, one uncelebrated, and one fictional; the reasons for these choices, as well as the ill-formed juxtaposition between the actual and the fictional, will become clear as we proceed.

### Ticking-Time-Bomb Cases Redux

For our first case, I start with one that Alan Dershowitz has championed: the torture of Abdul Hakim Murad by the Filipino police in 1995. Murad was arrested in a security sweep ahead of Pope John Paul II's visit to the Philippines and was subsequently tortured for sixty-seven days. The torture was effectuated through beating—which broke most of his ribs—forcing water into his mouth, forcing him to sit on ice, and burning his genitalia with cigarettes.<sup>55</sup> During his ordeal, Murad disclosed information that helped foil “plots to assassinate the pope and to crash eleven commercial airliners carrying approximately four thousand passengers into the Pacific Ocean, as well as a plan to fly a private Cessna into CIA headquarters.”<sup>56</sup> Following his ultimate confessions, Murad was handed over to the U.S. authorities, along with the information that the Filipino police had gleaned.

While this looks like a model ticking-time-bomb case, there are some features that bear further discussion. Rejali, for example, has been critical of it for at least two reasons. First, the sweep also resulted in the acquisition of an encrypted computer and disks; once these files were decrypted, they revealed the same information—less the Cessna plot—that was disclosed under torture. Second, the torture lasted sixty-seven days and was, therefore, not used against an imminent threat (i.e., lest that threat would have been realized in the interim). Furthermore, Murad began to talk only once a new team of interrogators introduced themselves as Mossad agents and said that they were taking him to Israel.<sup>57</sup> Rejali thinks that this fact therefore means it was Murad's “imagination and personality [i.e., susceptibility to the faux Israeli threat], not actual torture, that got him talking.”<sup>58</sup> Both these are fair objections, so let me take them in turn.

Regarding the computer and disks, they were not decrypted in time such as to make the torture redundant. Therefore, it is misleading to portray the torture as unnecessary given that the information gained through torture was also on the computer and disks since that information was unknown and inaccessible to the authorities.<sup>59</sup> Rejali approvingly evokes a discussion initiated by Jay Winik in which

Winik wonders what would have happened if Murad had been handed over to the Americans earlier; Rejali thinks it is straightforward that the decryption would have been faster, presumably in such a time frame that torture would not have been needed. The idea is that, even if the Filipinos could not run the decryption, the Americans could have and that the computer and disks should, therefore, have been immediately handed over to the Americans. But, supposing that the Americans could have readily decrypted these materials or else could have done so much faster than the Filipinos—neither of which seems obvious to me—there was no reasonable basis for the Filipinos to know what would have been on them in the first place and, therefore, no motivation to seek American help.

A quick refrain might be “if you don’t know what you have, hand it to the Americans”—or some other technologically sophisticated country—but there are all sorts of reasons to find this problematic. These range from considerations of national sovereignty, to national security (namely, that of the Philippines), to resources (e.g., the United States cannot provide integral support to everyone else’s counterterrorism), to simple logistics. To put it another way, the costs—whether economic or moral—of torturing Murad seem to me far lower than whatever alternative Rejali means to suggest (i.e., he never makes a concrete proposal). Regardless, the Filipino authorities *did* contact the Americans:

Philippine investigators called in their American counterparts for help. This was standard operating procedure. According to US and Philippine officials interviewed... both the CIA Manila station chief and the resident FBI legal attaché were notified. A team of intelligence agents flew in from Washington.<sup>60</sup>

Furthermore, the only mention of decryption is that the “data were encrypted and in Arabic, but Philippine technicians eventually deciphered the code and translated the texts.”<sup>61</sup> Nothing in the news article says that the Filipino authorities were unable to decrypt the information in a reasonable time frame or that the Americans would have been able to do any better. For all these reasons, the existence of the computer and disks does not undermine the case for torture.

The second central issue in this case is the sixty-seven days of torture. If these written accounts are really meant to be believed, Murad gave up nothing until the very last day of interrogation, the day on which he was lied to about a pending transfer to the Israelis. For the sake of argument, let us grant that this is true. Does it undermine the case for torture? First, why would it? The reason is supposed to be that ticking-time-bomb cases require imminent threats; if the associated threat were more than two months distant, then there would be no cause for torture since other intelligence avenues could be pursued.

The facts of this case are revealing as to why this sort of response fails. We can rest assured that, while Murad was being tortured, other intelligence avenues were being pursued: it is certainly not the case that the entire intelligence community did nothing else in the interim other than torture Murad in the hope that such torture would pan out. Investigations continued, leads were followed, and so on. And, as it turns out, none of this revealed the nature of the threats: only his confession did. The electronic materials would have given similar information, but they were still encrypted, and there was no reasonable basis for suspecting that they contained relevant material.

This leads to a second point, which is more general: what matters is, not imminence, but necessity.<sup>62</sup> Ticking-time-bomb cases are often formulated in terms of imminence, but that is misleading insofar as imminence and necessity come together in those cases. In other words, a threat may be imminent, but it can also be true that no method but torture would lead to its abrogation; torture is necessary to prevent the threat. When these features are run together, we lose sight of which feature matters. Regardless, it is easy enough to pull them apart:

*Imminence.* A bomb will go off in the near future and kill thousands; the torture of a detainee will ensure the deactivation of the bomb and the preservation of all the lives. Aside from torture, there is some other act,  $\varphi$ , that will also ensure the preservation of all the lives. Unlike torture,  $\varphi$  has no moral hazards.

*Necessity.* A bomb will go off in two years and will kill thousands; the torture of a detainee will ensure the deactivation of the bomb and the preservation of all the lives. The detainee has acted alone and is the only one who knows the bomb's location. The bomb is undetectable and will not be found without his testimony. Unfortunately, the detainee is dying of an aggressive form of cancer; while he is still strong enough to undergo torture, he will assuredly die in the near future, thus severing any hope of deactivating the bomb.

I submit that torture is not justified in the imminence example, that the necessity example is the same in every morally relevant way as standard ticking-time-bomb cases, and that, if we think that torture is justified in those cases, then we should think that it is also justified in the necessity example. Imminence is often suggestive of necessity, and non-imminence also tends to run contrary to necessity; nevertheless, these are at best probabilistic relationships that bear exceptions. What we should look for when countenancing torture is necessity—not imminence—even if it is true that necessity will often require imminence.

Returning to Murad's case, the prolonged length of his torture is morally adventurous. The question is not, how long it took to get the information from him, but whether torture was necessary to get it. And, furthermore, the fact that he ultimately relented under threat of interrogation from Israeli interrogators does not undermine the case for the antecedent torture: that torture may well have undermined his resolve and led to capitulation. Would he have caved had that threat been wielded earlier? Maybe, maybe not. One has to suspect that it *was wielded* in sixty-seven days as that is a long enough period of time that a lot must have been said; the ultimate appearance of officers dressed as Mossad agents just reified the threat. Even granting that there might have been some more expedient way to engender the confession, that shows, not that torture was not justified, but only that some other regimen could have been advisable. That said, the intelligence community should be afforded reasonable—and probably even extensive—latitude in this regard given the stakes, the vagaries of real-world interrogations, epistemic limitations, and so on.

This first case I chose for its celebrity as well as the opportunity to draw the aforementioned distinction between imminence and necessity. The second case—to be presented shortly—was chosen precisely because of its lack of celebrity: it appears nowhere in the torture literature, and I have never heard anyone discuss it. In fact,

I just happened to come across it in the *New York Times* one morning, without any research or agenda. The point is that there are probably various cases like this one and that many of us are completely unaware of their existence; it is important for the defender of torture not to singularly cling to cases like Murad's. There are few enough cases of legitimate torture so as not to call into doubt the presupposition that such cases need to be fairly rare and be the exception rather than the rule, but we can reasonably expect there to be at least a few cases per year. Let us now turn to the details of this second case.

Since the United States deposed the old Iraqi regime, American and Iraqi security forces have worked together, both to increase the overall security and to train the Iraqi forces ahead of the prospective departure of American forces. As was reported in April 2007—with the date of the incident unspecified but presumably somewhat proximate—Iraqi soldiers were on patrol in Ghazaliya, a neighborhood in the northwest of Baghdad. These soldiers picked up Mustafa Subhi Jassam after Jassam was seen loitering around the patrol route twice in the same day; two other suspected insurgents were picked up separately. The Iraqis beat Jassam in front of the other two detainees, and the stripes on Jassam's back indicated that he was probably beaten with electrical cables. Jassam was then handed over to the Americans and took them to one of the houses in which al-Qaeda agents made bombs. In this house, those agents had been making fuses for improvised explosive devices (IEDs), and various other materials were found: large quantities of soap (for making explosives), coils of blasting wire, two large antiaircraft guns, three propane tanks, and an oxygen tank that had been modified to house a bomb. Large piles of homemade explosives were drying on the roof. According to the commanding officer, Captain Darren Fowler, the detainees also gave the names of al-Qaeda agents (including those of high rank), identified safe houses and weapons caches, disclosed holding cells for kidnap victims, and provided other tactical information.<sup>63</sup>

This case has all the features of a ticking-time-bomb case, or at least as many as any empirical finding would reasonably be able to identify. Did the Iraqis know that Jassam was a terrorist? They did not, but they apprehended him because of suspicious activity in a conspicuous place. If we were even minimally charitable, we would grant that the Iraqis had some cause to start the beating, such as guilty or erratic behavior after Jassam's apprehension. Might Jassam have been innocent? Sure, but the number of lives hanging in the balance invites at least some risk. And, regardless, Jassam was not innocent insofar as he was complicit in terroristic activity and had firsthand knowledge of plans that threatened many.

Was torture necessary to force Jassam to talk? Let us assume that the Iraqis asked Jassam for whatever he knew before they grabbed the electrical cables; this is only the most minimal extension of charity. Or, to put it another way, Jassam could have confessed when the prospective beating was imminent but before it had commenced, yet he chose not to. Would there have been another way to find all the information that he surrendered? As suggested earlier, the intelligence community does not sit around waiting for torture to produce information: intelligence officers are constantly—if not always for the overwhelming majority of them—pursuing other recourses. As it turns out, none of those had yet worked. Would something else have panned out before these weapons were deployed? While we will never know for sure, there are reasons to suggest a negative answer. First, it is virtually impossible that the Iraqis were building weapons they never planned to use. Second, if they had no plans to use them in the (at least somewhat) near future,

why build them prematurely? Having weapons on hand increases the chances that they would be confiscated, that the insurgents would be captured, and so on. It is, therefore, only logical that these weapons would have been used at least reasonably soon after Jassam led the Americans to them. We cannot rule out the possibility that there would have been another way to eliminate the threat, but we also have no reason to think that there was.

This case identifies the straightforward torture of a terrorist such that life-threatening weapons were captured; other potentially lifesaving personnel and tactical evidence were also gained. Furthermore, there was good reason to think that Jassam was up to something: his repeated presence along the patrol route suggests that he was doing reconnaissance either for future attacks or else to learn the vulnerabilities of that route (e.g., in order to move weapons). The readiness of the weapons suggests imminence, and the lack of other intelligence measures to have discovered this very active facility at least makes necessity (for near-term recovery) a plausible claim. In point of fact, the torture did work, and it hardly required any sort of elaborate institutional configuration, just a few Iraqis and some electrical cables. And there are probably a handful or two of cases like this that have developed since the American incursion in Iraq; I strongly doubt that this one is completely unique in any morally relevant ways. Regardless, I certainly think that it represents a case of justified torture.

As a final case, consider the following:

Height of the [Australian] summer, Mercury at the century-mark; the noonday sun softened the [asphalt] beneath the [tires] of her little Hyundai sedan to the consistency of putty. Her three year old son, quiet at last, snuffled in his sleep on the back seat. He had a summer cold and wailed like a banshee in the supermarket, forcing her to cut short her shopping. Her car needed [gas]. Her tot was asleep on the back seat. She poured twenty [liters] into the tank; thumbing notes from her purse, harried and distracted, her keys dangled from the ignition.

[While] she was in the service station a man drove off in her car. Police wound back the [service station's close-captioned television], saw a heavy set Pacific Islander with a blonde-streaked Afro entering her car. "Don't panic," a Constable advised the mother, "as soon as he sees your little boy in the back he will abandon the car." He did; police arrived at the railway station before the car thief did and arrested him after a struggle when he vaulted over the station barrier.

In the [patrol car] on the way to the police station: "Where did you leave the [car]?" Denial instead of dissimulation: "It wasn't me." It was—property stolen from the car was found in his pockets. In the detectives' office: "[It's] been twenty minutes since you took the car—little tin box like that car—It will heat up like an oven under this sun. Another twenty minutes and the child's dead or brain damaged. Where did you dump the car?" Again: "It wasn't me."

Appeals to decency, to reason, to self-interest: "It's not too late; tell us where you left the car and you will only be charged with Take-and-Use. That's just a six month extension of your recognizance." Threats: "If the child dies I will charge you with Manslaughter!" Sneering, defiant and belligerent; he made no secret of his contempt for the police.



Part-way through his umpteenth, “It wasn’t me,” a questioner clipped him across the ear as if he were a child, an insult calculated to bring the Islander to his feet to fight, there a body-punch elicited a roar of pain, but he fought back until he lapsed into semi-consciousness under a rain of blows. He [had gotten into fights from time to time], but now, kneeling on hands and knees in his own urine, in pain he had never known, he finally [realized] the beating would go on until he told the police where he had abandoned the child and the car.

... When found, the stolen child was dehydrated, too weak to cry; there were ice packs and dehydration in the casualty ward but no long-time prognosis on brain damage.<sup>64</sup>

There are various reasons that I have chosen to use this case, so some commentary is warranted. First, this is obviously a case, not of terrorism, but rather of kidnapping. As argued elsewhere, torture need not always be practiced against terrorists and that kidnapping cases might serve as other appropriate contexts.<sup>65</sup> What ultimately matters is simply that the benefits of torture outweigh the costs of abstinence, and cases like the above make such a supposition plausible. Unlike with earlier sections of this chapter, I shall not go through myriad long-term effects of the criminal justice system as might be effected *vis-à-vis* torture, but suffice it to say that analogous arguments could be developed.<sup>66</sup> Note that, in this case, we know that the detainee is guilty: he is caught on tape and has stolen property from the car in his possession. Note also that this case postulates saving, not the lives of thousands, but the life of a single small child. While cases become more dramatic—and, to my mind, easier to reconcile—as greater numbers are at risk, I do not think that such numbers are always required. And, as with the previous case, note that no grand institutional structure is required for this torture, at least not any more than is already in place (e.g., a police force). Finally, we really do have straightforward necessity in this case: the outdoor heat and the unknown—and, in the relevant time scale, unknowable—location of the car force a decision between the child’s death and the beating of the detainee.

As it turns out, this case is fictional; it was written by a former police officer, perhaps as a teaching tool for cadets. So why use it? Nothing in the case strains credulity: there are no wild stipulations, and the case could easily enough be seen on the evening news, if rarely. Ticking-time-bomb cases are criticized as being overextensions of our empirical circumstances,<sup>67</sup> but nothing in this case sounds even remotely foreign. To be sure, I am not proposing that cases like this are often realized; the point is simply that they are absolutely empirically possible. As suggested in §4, actual cases are obviously going to go underreported, so the critic should not overplay their apparent dearth. But, whether we have a catalog of such cases or not, the point of using a fictional case is to drive home the point that such cases are less fantastic than commonly alleged. Regardless of whether the cases are actual or not, we can still ask what should be done in them, and it is reasonable to assume that something like the above will happen, somewhere or sometime, in human experience. So the point of this case—aside from illustrating a nonterrorism context of torture—is to make plain that the actual and the fictional need not be that far apart.

As I said at the end of the last section, I think that these three examples provide an illustrative basis for what actual cases do or could look like. Critics might allege that more cases are needed, but, to my mind, even one would be enough.

Regardless—and to repeat an ongoing point—more cases have to exist than we know about. For present purposes, these will suffice, especially given space constraints. Dershowitz talks about others, as, less sympathetically, does Rejali.<sup>68</sup> The objective has been, not to be comprehensive, but rather to give some sort of concrete context to the previously abstract treatment of ticking-time-bomb cases; these examples and the associated discussion make good on that.

This essay has meant to take seriously real-world objections to torture. Defenders of torture have had curiously little to say about these challenges, so the responses deserve adequate space. Surely, more empirical criticisms will or have come in, yet I cannot stem them all at present; this essay represents part of a broader (and more philosophical) project against which adjudications must be made for space. My principal hope is to rebut claims of empirical obliviousness, whether through direct empirical engagement or else through further elucidation of the theoretical project. On the former, many empirical claims have been made in this chapter, and they are either true or false; discussion of them is certainly invited. On the latter, I have continued to articulate the position espoused most directly in the previous chapter: the central point is just that the vagaries of actual experience (e.g., the loss of epistemic certainties) complicate rather than derail ticking-time-bomb thinking.

## Notes

1. See Fritz Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (Chicago: University of Chicago Press, 2012).
2. Darius Rejali, *Torture and Democracy* (Princeton, NJ: Princeton University Press, 2007), 446–447.
3. For more discussion, see Ronald Melzack and Patrick Wall, *The Challenge of Pain* (Harmondsworth, England: Penguin, 1982). See also Peter Deeley, *Beyond Breaking Point* (London: Arthur Baker, 1971).
4. Rejali, *Torture and Democracy* (see note 2 above), 449.
5. *Ibid.*, 455–456.
6. See also Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), chap. 9.
7. Rejali, *Torture and Democracy* (see note 2 above), 455–456.
8. *Ibid.*, 464. For more discussion, see Aldert Vrij, *Detecting Lies and Deceit* (Chichester, NY: John Wiley & Sons, 2000). [See also Samantha Mann, Aldert Vrij, and Ray Bull, “Detecting True Lies,” *Journal of Applied Psychology* 89, no. 1 (2004): 137–149.
9. It is worth recognizing that this point is made by supposing a (false) dichotomy between torturing and doing nothing; the structure of the dialectic is set up only to show the limited force of the misinformation objection. More will be in §4 about alternatives to torture.
10. Rejali, *Torture and Democracy* (see note 2 above), 458.
11. *Ibid.*, 474.
12. Jean Maria Arrigo, “A Utilitarian Argument against Torture Interrogation of Terrorists,” *Science and Engineering Ethics* 10, no. 3 (2004): 1–30, 22 (emphasis added), quoted in Henry Shue, “Torture in Dreamland: Disposing of the Ticking Bomb,” *Case Western Reserve Journal of International Law* 37, nos. 2–3 (2006): 231–239, 236.
13. Vittorio Bufacchi and Jean Maria Arrigo, “Torture, Terrorism and the State: A Refutation of the Ticking Time-Bomb Argument,” *Journal of Applied Philosophy* 23, no. 3 (2006): 355–373, 362–366.
14. See Fritz Allhoff, “Physician Involvement in Hostile Interrogations,” *Cambridge Quarterly of Healthcare Ethics* 15 (2006): 392–402. See also Fritz Allhoff, ed., *Physicians at War: The Dual-Loyalties Challenge* (Dordrecht: Springer, 2008).
15. See, e.g., Rejali, *Torture and Democracy* (see note 2 above), 245–247.

16. Bufacchi and Arrigo, "Torture, Terrorism and the State" (see note 13 above), 364.
17. Arrigo, "A Utilitarian Argument against Torture Interrogation of Terrorists" (see note 12 above), 236.
18. Jessica Wolfendale, "Training Torturers: A Critique of the 'Ticking-Bomb' Argument," *Social Theory and Practice* 32, no. 3 (2006): 269–287, 270.
19. Shue, "Torture in Dreamland" (see note 13 above), 236.
20. Ibid., see also Wolfendale, "Training Torturers" (see note 18 above), 273, 274.
21. A related issue is whether we should think of torture in individual *cases* or else whether we should think about our social *policies*, which will either support or abhor torture. This idea is somewhat redolent of rule utilitarianism, but I would just say that our social policies should be to torture only in justified cases, and the associated moral rule should be to torture when doing so maximizes the moral benefits. It is a red herring to say that we need to pick between "always torture" and "never torture" as surely we do not. Policies just represent agglomerated treatments of individual cases, and there is nothing misleading about focusing on cases or else on proposing that our policies should be nuanced rather than blunt. The arguments regarding the institutionalization of torture in §3 are also relevant in terms of thinking about the policy issues, though see also Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), chap. 8.
22. Andrew Fiala, "A Critique of Exceptions: Torture, Terrorism, and the Lesser Evil," *International Journal of Applied Philosophy* 20, no. 1 (2006): 127–142.
23. Henry Shue, "Torture," *Philosophy & Public Affairs* 7, no. 2 (1978): 124–143, 139.
24. Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), §5.2.
25. For more discussion, see Aibid., §2.1.
26. Stanley Milgram, "Behavioral Study of Obedience," *Journal of Abnormal and Social Psychology* 67 (1963): 371–378. See also Stanley Milgram, *Obedience to Authority: An Experimental View* (New York: Harper & Row, 1974). A short and accessible excerpt from the book is available as Stanley Milgram, "The Perils of Obedience," *Harper's Magazine* (December 1973): 62–77.
27. Craig Haney, W. Curtis Banks, and Phillip G. Zimbardo, "Interpersonal Dynamics in a Simulated Prison," *International Journal of Criminology and Penology* 1 (1971): 69–97. See also Craig Haney, W. Curtis Banks, and Phillip G. Zimbardo, "Study of Prisoners and Guards in a Simulated Prison," *Naval Research Reviews* 9 (1973): 1–17.
28. See, e.g., Wolfendale, "Training Torturers" (see note 18 above).
29. Logically, there are two other possibilities as well. The first is that the alternatives are less effective and more morally costly, which obviously makes those alternatives ill advised. The second is that the alternatives are more effective and more costly, though I do not think that anyone holds this: the principal argument against torture is its moral cost, and it is hard to envision other intelligence practices that would carry an even higher burden. The alternatives that are proposed in the literature certainly do not bear out this possibility, so I take it that the two options mentioned in text are the ones that should be considered.
30. Matthew Alexander and John R. Bruning, *How to Break a Terrorist* (New York: Free Press, 2008).
31. As a secondary theme to the book, Alexander paints a clash of cultures among the interrogators, contending that there is a deep divide between those advocating hostile and those advocating non-hostile means of interrogation.
32. Department of the Army, *FM 2-22.3 (FM 34-52), Human Intelligence Collector Operations* (Washington, DC: Author, 2006), <http://fas.org/irp/doddir/army/fm2-22-3.pdf>, §§ 8-6–8-20.
33. Ibid., § 5-26, quoted in Alexander and Bruning, *How to Break a Terrorist* (see note 30 above), 69. Alexander and Bruning actually have the wrong reference, citing to the *Army Manual [sic] Field Manual*, 34–52, i.e., Department of the Army, *Intelligence Interrogation. Human Intelligence Collector Operations* superseded *Intelligence Interrogation* in 2006. Given the publication date of Alexander's book, this oversight is curious.
34. Mathematically, let  $PS_x$  stand for the probability that some technique will succeed in eliciting the necessary information,  $MC_x$  stand for the moral costs of that technique—whether successful or not—and  $MB$  stand for the technique's moral benefits (if successful). The question then is whether  $PS_\varphi \times MB - MC_\varphi > PS_\psi \times MB - MC_\psi$ . If so, we should  $\varphi$ ; otherwise, we should  $\psi$ .

In other words, assume that torture has the same costs whether it is successful or not, which is at least a reasonable approximation. It might be false insofar as failed torture could raise more umbrage, distrust, and so on, but set that aside for simplicity. What we then care about is the probability that torture will work, times the moral benefits that would accrue if it does, minus its moral costs. That quantity then gets compared to the (lesser) costs of the alternative to torture, as subtracted from the expected moral benefit of that alternative. The unweighted moral benefits are the same on each side of the inequality since both torture and its alternatives aim at the same end. This is not to say that they have the same consequences, but I find it more intuitive to think that those different consequences play out under moral costs rather than under moral benefits (i.e., the alternatives to torture avoid certain moral costs rather than probabilistically aspiring toward different moral benefits).

35. Alexander and Bruning, *How to Break a Terrorist* (see note 30 above), 31, 111, 158, 165.

36. For more discussion, see Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), §4.1.

37. United Nations Treaty Collection, IV.9.

38. Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), §4.1.

39. Alexander and Bruning, *How to Break a Terrorist* (see note 30 above), 103.

40. *Ibid.*, 157–158, 160–163.

41. Bufacchi and Arrigo, “Torture, Terrorism and the State” (see note 13 above), 367.

42. *Ibid.*, 368.

43. Rejali, *Torture and Democracy* (see note 2 above), 111, quoting Mellor, *La torture*, 218.

44. Alexander and Bruning, *How to Break a Terrorist* (see note 30 above), 236. The reference is to a colleague who is attempting to leverage hope in exchange for intelligence.

45. *Ibid.*, 5, 277–279.

46. For more discussion, see Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), §9.1.

47. See also Shue, “Torture” (note 23 above).

48. For more discussion, see Allhoff, *Terrorism, Ticking Time-Bombs, and Torture* (see note 1 above), chap. 2., n. 40.

49. At least if the disjuncts are distinct, and neither is necessarily false.

50. Bill Sammon, “Cheney: Enhanced Interrogations ‘Essential’ in Saving American Lives,” *Foxnews.com*, August 30, 2009, <http://www.foxnews.com/politics/2009/08/30/cheney-enhanced-interrogations-essential-saving-american-lives/>. There is no shortage of these sorts of quotes from Cheney. President George W. Bush has been somewhat more circumspect but has, ultimately, said similar things. See, e.g., Peter Hamby, “Bush Defends Interrogation Program in Michigan Speech,” *CNN.com*, May 29, 2009, <http://www.cnn.com/2009/POLITICS/05/29/george.bush.speech/index.html>.

51. “September 11 Mastermind Khalid Sheikh Mohammed ‘Waterboarded 183 Times,’” *Times Online*, April 20, 2009, [http://www.timesonline.co.uk/tol/news/world/us\\_and\\_americas/article6130165.ece](http://www.timesonline.co.uk/tol/news/world/us_and_americas/article6130165.ece). There is at least some dispute as to how to count instances of waterboarding. See, e.g., Joseph Abrams, “Despite Reports, Khalid Sheikh Mohammed Was Not Waterboarded 183 Times,” *Foxnews.com*, April 28, 2009, <http://www.foxnews.com/politics/2009/04/28/despite-reports-khalid-sheikh-mohammed-waterboarded-times/>.

52. Another al-Qaeda suspect, Abu Zubaydah, was allegedly waterboarded eighty-three times, yet a former CIA officer has maintained that this waterboarding did not engender any more information than Zubaydah was willing to give from the outset. See Abrams, “Despite Reports” (see note 51 above). See also Rejali, *Torture and Democracy* (see note 2 above), 504–507.

53. Marc A. Thiessen, “The CIA’s Questioning Worked,” *The Washington Post*, April 21, 2009, <http://www.washingtonpost.com/wp-dyn/content/article/2009/04/20/AR2009042002818.html>.

54. Timothy Noah, “Bored: Al-Qaida’s Plot to Bomb the Library Tower Was Not Worth Torturing Anyone Over,” *Slate*, April 21, 2009, <http://www.slate.com/id/2216601/>.

55. Alan M. Dershowitz, *Why Terrorism Works: Understanding the Threat, Responding to the Challenge*, (New Haven, CT: Yale University Press, 2002), 137. The more comprehensive account is found in Matthew Brzezinski, “Bust and Boom,” *Washington Post Magazine*, December 30, 2001, W09, <http://www.washingtonpost.com/ac2/wp-dyn?pagename=article&node=&contentId=A14725-2001Dec21>. See also Peter Maass, “If a Terror Suspect Won’t Talk, Should