

Chapter 9 : The Limits of Torture

This last chapter comprises some concluding remarks regarding the themes explored in the book; in this sense, it is closer to an afterword than a substantive chapter. As I said in the preface, the central theme of this book has been that lesser harms are always preferable to greater harms and that torture, while bad, could nevertheless be the lesser harm in exceptional cases. In Chapter 6, this ‘could’ was presented as one of philosophical thought experiment, but Chapter 7 tried to give it some traction in the real world. An important part of the project was to situate the debate about torture within the broader context of terrorism (Part I), and then to make the associated philosophical discussion (Part II) empirically engaged (Part III). I can foresee objections coming from every angle, but I have tried to forestall them in the associated discussions. The purpose of these concluding remarks, then, is not to further the defense or to recapitulate key claims, but rather to end the book by reflecting on two themes that have received insufficient attention so far. And, as the title of this chapter indicates, to acknowledge the limits of torture.

9.1 Torture Is Not a Panacea

In Chapter 6, I argued that virtually any plausible moral theory could defend the permissibility of torture in exceptional cases; the key feature of these cases was simply that torture minimized overall harms. It is important to emphasize that I really do take these cases to be exceptional or, to say it another way, very uncommon. As was made clear in Chapter 4, torture is bad, and for various reasons. Torture causes suffering, and suffering is bad. Interrogational torture also forces the tortured into a position of self-betrayal, and that is also bad. So, at the end of the book, I want to emphasize that my defense of torture acknowledges the wrongs of torture; I do not unapologetically or unreservedly

champion torture. Rather, I lament that we live in a world in which terrorism threatens lives and that torture might sometimes be our most prudent recourse against those threats.

If torture really would be so exceptional, why write an entire book about what might bear only on a handful of cases? I never have understood the point of this question insofar as, handful of cases or not, we still have to decide what to do in them. One time I went to a colloquium given by a speaker whose research area was the ethics of infectious disease; he almost seemed stupefied that applied ethicists would work on any other topic given how many annual deaths are attributable to HIV/AIDS, tuberculosis, and malaria (i.e., several million). If we think about the harms proffered by diseases or other societal ills—such as poverty or global warming—why spend so much intellectual energy worrying about how to deal with some small number of ticking time-bomb-like cases? For what it is worth, this is not just me: the literature on the moral status of interrogational torture is quite large, and growing. Why is it that we find ourselves so captivated by these questions and so untaken with others? There are many answers here but, for better or worse, torture is just such a compelling and visceral topic to many of us. While much of disease afflicts those far away, 9/11 was an attack on us and our country; it left many of us feeling personally violated and instantly catapulted terrorism into our national consciousness. Even though there is no reason to think that torture could have readily prevented 9/11, that horrible day nevertheless catalyzed the ensuing discussion about how far we may go in order to protect ourselves.

That said, it should be recognized that there are only a small number of marginal cases in which torture is a live moral topic; the rest of the cases are as morally uncontroversial as those decried in §4.2. To some extent, though, the focus on marginal cases is similar to many other debates in applied ethics. Last year, for example, there were fewer than 50 executions in the US, though capital punishment remains a lightning rod for moral debate. Human (reproductive) cloning is hotly contested—and the associated discourse enjoys a wide public forum—even though there has never been a successful

application thereof. This is not to deny that some debates in applied ethics pertain to more ubiquitous practices, such as abortion, pornography, and the treatment of non-human animals. Rather, the point is simply that torture does not stand alone in being contentious if infrequently justified. And, sadly, torture itself is hardly infrequent at all, even if the justified cases might be.

Still, people often ask why I wanted to write this book. A large part of it was simply that my views were not reflected anywhere else in the philosophical literature; there are few conservative voices in a very liberal profession. But it goes beyond that insofar as this is simply a fascinating moral question, regardless of how many ticking time-bomb cases there actually are. For most of the quotidian issues we face, the ethical answers are more or less straightforward. When we get to the margins—to those rare cases in which various values are pitted against each other in non-standard ways—is when ethics gets interesting. In other words, it is usually obvious what we should do, and we just need to have the character and resolve to do it. In the case of torturing a suspected terrorist, however, it is hardly obvious what to do, and reasonable people disagree vociferously on this issue.

That said, let me again note that we are arguing over a small handful of cases. I do not believe in widespread, institutionalized torture; nor, as far as I know, does anyone. A defense of torture is not meant to extend beyond highly-circumscribed cases and, in fact, part of a defense owes an accounting for how such spillage can be foreclosed (cf., §7.3). Torture should not be deployed cavalierly, and our intelligence community must be vigilant against inappropriate uses. While torture that fails to yield actionable intelligence might be justified given our epistemic status, the number of lives that hang in the balance, and so on (cf., §7.1), it nevertheless must be minimized. And, furthermore, even if torture were to be successful in saving lives, it could only be justified if it were reasonably thought to be necessary, or at least prudent.¹

¹ Imagine that either ϕ or ψ would prevent some bad outcome A; therefore neither is necessary for its prevention. Suppose that ϕ is bad, but ψ is worse. So, even though ϕ is bad and not necessary to prevent A, it is our best course of action. Alternatively, imagine that ϕ had some probability of success, P_1 , and that ψ had some other

While the preceding paragraphs have meant to emphasize the scope of the project, there is another way in which torture is not a panacea: it does nothing to address the underlying causes of terrorism that recommend it in the first place. The premise of the book has been that we find ourselves in a situation wherein torture might disarm a terrorist threat. But how did we find ourselves in this situation in the first place? Why is there terrorism at all? Much less against us? These are immensely complicated questions, and ones that have to do with far more than philosophy; geopolitics, economics, religion, and culture all play important contributory roles in explaining the root causes of terrorism. Torture is, at best, a temporary solution to a deeper problem. And, at worst, torture makes that problem even worse by exacerbating the ill will borne against us.

In approaching this discussion, it is useful to consider terms from the health care debate, particularly the distinction between curative and preventative medicine. Our health care system—reform notwithstanding—is too heavily invested in fixing people after they get sick rather than preventing them from getting sick in the first place.² Faced with a public that would rather overeat, underexercise, and spend billions on Lipitor and other cholesterol-lowering drugs, it is no surprise that American health finds itself in such dire straits. There are surely endogenous avenues that the health care industry must pursue, but many of its obstacles are exogenous. In other words, if people did not get sick—especially from easily preventable afflictions—then there would be far less to debate vis-à-vis health care. And, furthermore, the health care debate could benefit from shifting the focus from what to do with sick people rather than how to prevent people from getting sick.

probability of success, P_2 , where $P_1 > P_2$. Given our epistemic situation, we cannot say that ϕ is necessary to prevent A, since ψ might work. Even if ϕ were morally worse than ψ , ϕ could still be justified given the associated probabilities. (I.e., if ϕ were twice as bad as ψ , but four times more likely to avert the moral harms of A, then our moral calculus would support it.) In other words, ϕ needs to neither be necessary to prevent A nor the least harmful action in order to be justified; the implications for torture more generally and the above wording in particular are transparent.

² As I write, President Obama just today—March 29, 2010—signed into law a sweeping health care reform bill that does far too little to rectify this shortcoming; the principal focus of that bill is to increase access to care for the (previously) uninsured rather than to propose thoroughgoing reform as to how to prevent people from needing curative care in the first place.

The discourse about torture is isomorphic to that about health care in so many ways. Instead of fostering—or at least not dispelling—the circumstances that give rise to terrorism, we should figure out how to mitigate them. If we can make progress in this regard, then we will not find ourselves—or at least find ourselves less often—forced to countenance torture. It is far beyond my expertise to offer any substantive commentary on how this should proceed, but the suggestion is hardly novel regardless; everyone wants to lower the incidence of terrorism, presumably including the terrorists themselves.³ An exclusive focus on torture is therefore myopic insofar as, to continue the analogy with medicine, it fixates on a symptom rather than the underlying disease.

To be clear, I am not proposing an exclusive focus on torture: of course we should—and do—care about the root causes of terrorism. The discourse about torture is often insufficiently attentive to this broader context, but that shortcoming is neither here nor there when considering what to actually do in exceptional cases. In other words, supposing that we actually find ourselves the victim of a terrorist threat, it is not helpful to wonder how to prevent future threats. Rather, we must respond to the one currently at hand and, after we disarm it, we can think about how we got there in the first place and how to prevent recurrences. This book is about what to do when faced with serious threats; my hope is surely that such threats do not present themselves, but it begs off a serious intellectual task to ignore the nefarious potentiality or, more probably, eventuality.

³ This might sound like a strange claim, but remember that the point of terrorism is to use extreme fear in order to effectuate some sort of ideological change. In other words, there is nothing intrinsically motivating about the fear itself; that fear is meant to be used instrumentally. Therefore, there is no reason that terrorists should want to perpetuate terrorism if they can achieve their ideological goals by alternative means; this is especially true given the costs and risks of terrorism. If all someone cared about were causing terror—i.e., rather than the ideological goals—then we would not even call him a terrorist, but rather a sadist or a psychopath. To be sure, there might be incompatible ideologies that different terrorists wanted to promote, thus, from their perspective at least, making terrorism a practical necessity. But, from either of their perspectives, there would be no reason to prefer more terrorism to less. See Chapter 1 for more discussion on the conceptual foundations of terrorism.

9.2 How Far Should We Go?

A second theme worth considering in these closing remarks is that the previously-offered defense of torture could authorize some completely heinous acts, including highly brutalized torture and/or the torture of innocents. To wit, consider the claim that torture is justified if and only if it portends a lesser harm than we might reasonably expect to otherwise absorb. Such a claim puts no “upward bound” on how bad the torture could be; insofar as the terrorist threat gets worse, we are allowed to consider more egregious torture in order to prevent that threat. This does not deny that all torture is bad—of course it is—but only allows that it can either be bad or worse. If Khalid Sheikh Mohammed was waterboarded 183 times, that is a lot of suffering; it would have been less morally bad, for example, were he to have been waterboarded some fewer number of times insofar as fewer instances would have led to less suffering.⁴

In §5.1, we saw Jeremy Bentham's utilitarian basis for torture, and one of his precepts was that we should torture no more than was necessary to elicit our goal. While there are obvious epistemological and other empirical challenges to this moral edict (cf., §7.1), the moral principle is straightforward. Torturers should be held accountable for the torture that they inflict, and overly-zealous torture—i.e., torture that goes beyond what anyone might reasonably expect to be warranted—should be punished. That said, I still think that extreme torture can be justified, but only when such torture could reasonably be expected to be the lesser harm. For example, if intelligence reveals a terrorist threat against many lives, then greater latitude should be extended vis-à-vis torture than when fewer lives were threatened. Even were that torture to ultimately be unsuccessful, it would at least have been deployed against a greater threat; the calculus of expected outcomes could therefore support it.

That said, I still think that there will ultimately be limits on the sorts of torture that reasonable people could ultimately contemplate. Imagine, for example, that intelligence officers were to reach the

⁴ Claims like this do not presuppose a thoroughgoing science of torture—such as that objected to by Darius Rejali in §7.1—but rather are just meant to be self-evident.

conclusion that some detainee had no intelligence value, or that they at least had no reason to think otherwise. They could either stop the battery of interrogations, or else press forward with more aggressive techniques. While it *could* be the case that brutalizing the detainee prevents the next 9/11, there is simply no basis for believing that it would. In other words, even if the magnitude of the terrorist threat they sought to prevent was immensely grave, the likelihood that this particular detainee could diffuse the threat could be so remote as to not license his torture.

Of course we never know exactly what the likelihoods are; if we found ourselves in perfect epistemic situations, we could talk about certainties rather than likelihoods. Rather, we do the best we can with whatever information we have, and we hold those accountable who make decisions that are unsupported by that information. If some intelligence officer were to brutally torture some detainee, we could ask him why he did it. Maybe he has a good reason, in which case the torture could (but would not necessarily) be justified.⁵ Or else he does not, in which case he should be punished.

But still, the critic might press, given the potentiality of some catastrophic threat, the torture calculus could come out positive even if the likelihood of torture being efficacious were extremely low. Agreed. That said, there should still need to be some non-negligible reason to think that torture would work at all. If we were to round up some random person on the street and torturing him might prevent a grave threat, should we countenance his torture? In such a case, we have no reason to think that the likelihood of torture is anything other than (marginally above) zero, so, no, we should not torture; I cannot fathom any reasonable person saying otherwise. Furthermore, if we are just talking about some non-specific terrorist threat (e.g., there may be such a threat, but we have no positive reason to think that there is), then whatever rational basis there could have even been for torture (e.g., against a specific threat) is further reduced.

⁵ In other words, having a good reason to torture does not mean that there are not better, countervailing reasons not to torture. The existence of such reasons does not always presuppose that ignorance thereof is culpable, which is to say that a torturer might torture unjustifiably while, at the same time, not be blameworthy.

Tweak the case somewhat such that we apprehend someone who is a former neighbor of a known terrorist; our captive steadfastly denies knowledge of any terrorist plot. In reality, this sort of case probably comes up pretty often, so what should we do? Of course the details are going to matter but, if the case really were this simple, there would be little reason to think that torture would be effective: just because our detainee lived next to a terrorist, we should not expect that the former would have details of the latter's machinations, both because the relationship is attenuated and because the operational details of the plot would not be casually shared. And, while some sort of stern interrogation might be justified, there is simply no cause for optimism that a course of protracted torture would be useful. Maybe it would have been, or maybe some released terrorist would have capitulated under one more application of waterboarding; we never know. The best that we can do is think critically about the expected outcomes, even as messy and ill-informed as those can get on the ground. The more confident we are that our detainee is involved in some nefarious attack or the more confident we are that said attack is against a greater number of people, the more latitude we have in trying to preempt it. Ultimately, though, the point is that we simply cannot torture without some reasoned basis. We assuredly disagree as to what constitutes those bases, but arbitrary, capricious, or gratuitous torture plays no role in my account.

But what about the (reasoned) torture of the innocent? In §5.3, we considered ticking time-bomb case variants, one of which involved the torture of a terrorist's innocent daughter such that the terrorist would relinquish information about his pending attack. Not surprisingly, support for torture in these cases was lower than when the terrorist himself was to be tortured. The critic could press that nothing in my account precludes the torture of the innocent daughter, and this could presumably be held up as a defect of the account. So, first, let me acknowledge that my account allows for the torture of innocents when such torture could reasonably be thought the least morally offensive way to prevent

a greater harm. And then, second, let me explore what such a commitment might amount to in practice.

Starting with the torture of an innocent daughter, this is at least doubly improbable as a standard ticking time-bomb case. To wit, such a case asks us to imagine not only that we find ourselves in a standard ticking time-bomb case, but also that we have apprehended a terrorist's innocent daughter and, furthermore, that, while impervious to torture himself, he would crack were his daughter to be assaulted. If one thinks that ticking time bomb cases are rare, then one could reasonably expect that the cases involving the daughter are almost impossible. In other words, the conjoined probabilities of two exceedingly unlikely scenarios is vanishingly small, far smaller than the already unlikely probability of either scenario obtaining by itself. Terrorists simply do not travel around with their young daughters and, even were there some independent way to apprehend the daughter, we might expect a dedicated terrorist to remain true to his cause despite the torture of his daughter. Now it just seems that I am evading the question, all the worse since that was the same allegation I hurled at opponents of ticking time-bomb methodology back in §5.2. As I already said above, I think that such torture could be justified, but I just deny that it ever would be.

The question of torturing innocents also arose in §7.1, where I presented the following case:

Red Sweatshirt. Our intelligence reveals that the 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 only be disarmed if both men are tortured, one of whom is innocent.

Empirically, this sort of case is far more likely than the daughter case, so I see it as more of a worry. That said, my strategy is still the same: to deny that it is very likely to happen or else to say that, if it does, then the torture is justified. As I suggested in that earlier discussion, one way to look at the torture of innocents is as a sort of collateral damage, just as we think of the killing of non-combatants during war. Maybe torture is worse than killing or maybe not; regardless, the point is that some sort of moral costs

are allowable in the pursuit of a greater moral good. This does not deny that we must be vigilant against minimizing those costs, that some costs are not worth bearing, and so on. My view on torture is that it constitutes a moral wrong, but that wrong is commensurable with all sorts of other moral values. To put it another way torture still has to compete against other values in our moral calculus; if we should not torture, it is because our moral calculus recommends against it and not because there is anything about torture that sets it aside from that calculus altogether. Kantians undoubtedly see it differently, but their vision is not mine.

As a final point of discussion in this vein, consider what we might call "preventative torture". Roughly, I have in mind a distinction—sometimes promulgated under the aegis of the "Bush Doctrine"—between the use of force against actual aggressors and the use of force against (potential) future aggressors. While the just war tradition acknowledges the right of self-defense, that right is usually articulated in terms of defense against actual, as opposed to potential, aggression. So if some state is attacking ours, we can defend ourselves. If that other state is gathering on the border and attack appears imminent, then we can attack before we are attacked. This second scheme often goes under "preemptive war", but can still be well-grounded in classical just war theory, particularly if we understand the amassing of troops as an act of aggression, whether or not those troops have yet to attack.

As the temporal links become more attenuated, however, we shift the nomenclature from preemptive war toward preventative war: instead of preempting some particular attack, preventative action aims to prevent that attack from ever being either imminent or actual. If we bomb terrorist training grounds in Afghanistan, for example, that might reasonably be understood as preventative insofar as, let us suppose, attacks emanating thereof would neither be imminent nor actual. What about preventative torture? Just to have a concrete case for discussion, consider:

Media Officer. The intelligence community is on the trail of a terrorist organization's chief media officer; this officer is solely responsible for the dissemination of propaganda used to recruit

future terrorists. He is very effective at his job and produces anti-American flyers and DVDs that have been very successful for recruitment. If he is not apprehended in the near future, the media officer's imminent marketing campaign will ensure the recruitment of many new terrorists, as well as the associated calamities they would eventually perpetuate; if the media officer were captured, distribution of the marketing campaign will be forestalled and no new terrorists would be recruited. While the intelligence community has no other direct leads on this media officer, it has in custody several of his known, yet recalcitrant, associates. Were torture of the media officer's associates judged to be the only way to learn his location—and therefore secure his detention—before the marketing campaign was disseminated, would it be justified?⁶

Were this case to gain any empirical purchase, I would be deeply conflicted. The problem with preventative action is that it might ultimately prove unnecessary; given the time lag between the preventative action itself and whatever action it aims to prevent, any number of things could change. Maybe the political landscape changes in the interim such that the would-be terrorists are no longer compelled toward terrorism. Or maybe other terrorists are able to effect the sort of policy changes that would forestall other terrorist acts. Regardless, the simple point is that a number of things could happen, any of which would render the preventative attack an unnecessary cost.

On the other hand, preventative action might minimize overall harm, such as when we destroy one weed before it takes over the rest of our garden; were we to wait, the intervention would be far more costly. I suspect that many would take exception to the details of the above case, questioning whether a media officer really has an effect on the recruiting of terrorists, or whether he might be easily replaced by another (or even more effective) media officer such that the torture became superfluous. As it turns out, these possibilities are ruled out by the facts of the case, but we could always remain suspicious about how empirically plausible such cases are.

One thing that matters is how far removed the preventative action is from the action that it prevents; the greater the discrepancy, the more could change in the meantime. If we are countenancing the bombing of a training camp for very young children, then there is cause for skepticism about the necessity of the attack. Let us suppose that those children would not participate in terrorist attacks for

⁶ This case is loosely inspired by discussion in Matthew Alexander and John Bruning, *How to Break a Terrorist* (New York: Free Press, 2008).

some number of years, so political or diplomatic inroads would be plausible, or even simple intervention from regular law enforcement. If the training is instead geared toward more proximate attacks, then the case for prevention is stronger.

Were there to be a case exactly like the one posited above, I think that torture could be justified. That said, I think that such a case is exceedingly unlikely, precisely because we will rarely have good reason to believe that the media officer is that effective, that irreplaceable, that there would not be other political or diplomatic ways to address the training camps, or that normal military action would not be appropriate for dismantling them. So, as with the daughter case, my theoretical commitments are straightforward; such commitments do not have any obvious application to real world decisions. Rather, I suspect that torture would only ever be justified in cases reasonably close to ticking time-bomb cases and that the torture of innocents or preventative torture are not likely to recommend themselves.

9.3 Final Remarks

Before concluding, let me again emphasize that I think torture comprises various moral harms. One of myriad questions put to us by terrorism, though, is whether torture could be less morally harmful than the alternative. Terrorist threats portend grave moral harms, and we are surely able to defend ourselves against such harms. There will not be a great multitude of cases in which torture is the most prudent reply to terrorist threats, but nor are such cases ruled out, whether philosophically or empirically. To be sure, we must be extremely careful in choosing whether to utilize torture, paying close attention to our epistemic situations and our ability to constrain the use of torture to appropriate cases. These challenges are serious, but not insurmountable. Hopefully we may eventually find ourselves in a world where we need not countenance torture as a response to terrorism, especially were that world bereft of terrorism altogether; unfortunately, such a world is not yet ours.

As I said in the preface, my own views on torture and terrorism are indelibly connected to the horrendous terrorist attacks of 9/11; this book appropriately concludes on the eleventh page of the ninth chapter.⁷ While these attacks were among the most extreme imaginable, far lesser ones should still be vociferously opposed. The lives of innocents—particularly lives under nefarious assault—deserve our most ambitious protection. Let us not forswear any countermeasures in that defense.

⁷ [Author's note: Page positioning owes to estimate of 465 words/page from publisher; ensure correctness—or remove—during page proofs.]