THE SOLID CORE IN SOME SUSPECT APPROACHES TO TARGETED KILLINGS

Leo Katz

Targeted killings inspire great unease. Governments that engage in them defend them with arguments of self-defense and necessity that sound hollow to many. But there is more to those arguments, I think, than is generally acknowledged. There is more to them in two ways. First, they are in fact on stronger moral and legal ground than they might at first appear. And second, the dissonance between their apparent weakness and their actual strength is itself an interesting phenomenon. In this short essay I will try to substantiate both of those claims.

Let me begin with two examples that will help to make explicit—unnecessarily explicit as many see it, because the matter is so obvious to them—why targeted killings seem to many people so manifestly off limits. Their purpose, however, will not be merely to make that point, but also to carry the argument beyond it.

Example 1. Imagine a peaceable holy-man who has inspired a fanatical, violence-prone following. Members of his movement are in the process of organizing a violent attack upon us, the imperiled country. It would solve our problems if we could eliminate the holy-man with a targeted killing. But he is of course quite innocent of anything that would merit even more limited action against him. It just happens to be the case that those inspired by his cause are led to embark on violent action against us. Neither self-defense nor necessity, as conventionally interpreted, would seem to allow us to do anything to bring him down. The excuse of duress, I suppose, might fare a little better, but not much: first, because its scope might actually not reach that far, since we are not quite in the classic duress situation where someone tells us to do something lest he injure us. (Admittedly that is a very questionable restriction on the duress

defense.) But most importantly, because of the interesting fact that excuses are not the kind of defense a government can readily invoke. To be sure, there is something mysterious and perplexing in that which seems worth exploring, but I won't concern myself much with it in this essay.

Consider next

Example 2. A would-be terrorist has been tentatively thinking, and talking to others, about launching a strike against us. The attack is not imminent: he does not have a concrete plan, indeed his resolution to actually follow through is far from firm, but we have no trouble making a probability assessment and predicting with great confidence that at some point in the not too distant future, he will decide to strike, and that if we wait until he has clearly made up his mind and taken more than few steps toward his goal, it will be too late for us to stop him. (Why? Take it as a given. It shouldn't be too hard to invent suitable circumstances.) Can we kill him before things get out of hand? The general intuition is: No. Self-defense cannot be used because his attack is not sufficiently imminent. And necessity can't be used because we would essentially be killing a moral innocent—which he is until he has become enough of a threat for self-defense to apply—for the sake of saving a somewhat larger number of innocents, the kind of tradeoff we generally think the necessity defense cannot underwrite.

Not everyone of course is so very sure about that. People have questioned the moral basis of the imminence requirement. The suggestion has been made that it is just an awkward stand-in for the probability of attack. But I very much doubt that. My reasons aren't profound. I think that what those of us for whom the imminence requirement seems to have independent moral significance are guided by are our intuitions about criminal attempts more generally. Consider

what we do about the defendant who has actually formed the intent to carry out a crime but only taken a few steps towards implementing it. However sure we are that he will go further, as far as the law is concerned we are not entitled to punish him until he actually crosses the threshold between mere preparation and actual attempt—hazy though that line might be. The law here seems entirely in keeping with most people's moral intuitions.

To be sure, awkward situations have arisen where the police rightly felt compelled to intervene earlier—as in the well-known casebook case People v. Rizzo, 246 N.Y.334, 158 N.E.888 (1927), in which they foiled an impending robbery before the would-be criminals had done enough for an attempt, and they then had to let them go. Thus it is not entirely clear that just because we are not entitled to punish, we can't take preventive action. (It is hard to quarrel with what the police in *Rizzo* did. Indeed the Rizzo court, after acquitting the defendants because they could not be "guilty of an attempt to commit robbery …when they had not found or reached the presence of the person they intended to rob," then added that "The police of the city of New York did excellent work in this case by preventing the commission of a serious crime. It is a great satisfaction to realize that we have such wide-awake guardians of our peace.") Rizzo is far from unique. There are many situations where prevention and punishment get out of sync. After all, I am entitled to shoot an attacker, even though I am not allowed to inflict the death penalty as punishment for his attack. *Nevertheless,* some version of the reservations that make us insist on more than a firm resolve and a high probability that someone will act on it, when it comes to the law of attempt, is presumably at work when we insist on imminence in the self-defense context.

An example that helps bring out the force of those imminence-based intuitions is the case of the wife who foresees that her husband is about to learn about her infidelity and who anticipates, with excellent reason, that he is then going to kill her, perhaps because he has long

told her that this is exactly what he is going to do if he should ever find out such a thing about her. Is she entitled to try to kill him, when he has not even formed the intent to kill her, because she rightly and reliably predicts that this is what he is going to do? And if she is entitled to do so, what about *his* right to defend himself against *her* attack? Surely he has such a right. Do they both, then, have a right to act in self-defense? Note that this is not a situation in which two parties act in ignorance and thus end up at understandable loggerheads. That is what makes it hard to accept that they could both be acting legitimately. And if one of them isn't, it would seem to be the woman, who is unable to argue that she was facing an *imminent* threat. (To be sure, she might still be able to claim duress.) Imminence here is certainly no mere stand-in for probability.

If despite these impediments, a government is determined to carry out a targeted killing, and to make a public defense of the practice, it is likely to pursue two strategies which I want to now examine more closely. To deal with the case of the person who is predicted to launch an attack but simply hasn't yet approached the point of imminence, a government might simply use an agent to propel the potential attacker to that point more quickly, and then kill him once he has reached it. To deal with the completely innocent holy-man, a government might just proceed to kill, and then offer a litany of necessity arguments pointing to the number of innocent civilian deaths thereby averted and perhaps even to various pernicious but indirect consequences averted as well: the great number of severe rights infringements and unjustified killings that would undoubtedly ensue if his followers were permitted to unleash the conflict they about to. Let's take a closer look at each of these approaches. As I said, they will turn out to be both more interesting and more meritorious than they might appear at first glance.

Provoking the Attack

Standard self-defense doctrine holds that you forfeit your right of self-defense if you provoke the attack. Walking into a bar, goading your worst enemy with everything you can think of until he loses his self-control and lunges at you, and then finishing him off—that's the kind of case where we generally don't want to view the person acting in self-defense as justified. Isn't this what a government using an agent provocateur is doing?

Provoking an attack to use one's right of self-defense is a particular instance of a phenomenon sometimes described as contriving to create the condition of your own defense. German criminal law calls this the problem of the *actio libera in causa*—an "action free in its origins." This refers to cases in which someone deliberately, or maybe even just recklessly, creates situations in which he then finds a need to commit an act that would ordinarily be considered justified or excused by reason of necessity, or duress, or insanity, or lack of a voluntary act, or innumerable other potential defenses, but inasmuch as he deliberately contrived to create it, or just impermissibly risked their coming about, he no longer seems entitled to invoke them. His actions being "free in their origin", though no longer free in their execution, deserve to be punished just as though he had no defense to begin with. That at least is the standard view.

But why exactly is that? One line of argument that has been offered to deal with cases of contrived defenses goes something like this. Consider a case in which I contrive to create a situation of necessity so as to be able to do damage to my closest competitor's storefront without having to fear criminal liability. To this end, I start driving down at breakneck speed the narrow alleyway where I know his store to be located. I anticipate that when I get close to the store, there will probably be a cluster of people congregating in front of it, and I will then face the choice of either running into them or veering to the side and damaging my competitor's

storefront instead. Clearly I should not manage by this subterfuge to really avoid liability. But how exactly should this kind of situation be analyzed? What the most thoughtful commentators of this situation have said—the MPC commentary actually does not qualify in this regard—is that we should think about the matter in the following way: There is no doubt that I acted rightly when at the moment of encountering the cluster of people I decided to veer into the storefront instead and should not on that account be held liable. I really do have a valid necessity defense at that moment. However, I did not have a valid claim to a necessity defense when inaugurating the chain of events that led to this denouement—in other words, at the moment at which I first began hurtling down the alleyway. All that liability requires is that the prosecution be able to identify a moment in time at which I committed an act that constituted the forbidden actus reus of damaging another person's property, and I must not have a valid defense for that act. Well, when causing my car to hurtle down the alley way, what I did could probably be described as the first step of damaging another person's property, and for that step I did not have a justification: Necessity certainly did not require me to race down the alley way. It only required me to swerve once I was already doing so. Ergo, I am liable.

Possibly an even more compelling way to make this point is the following example: A person throws himself out the window, so as to land on top of his intended victim, whom in this fashion he kills. He then claims the involuntary act defense, because at the moment at which his body made contact with his victim's he was not committing a voluntary act. We would deny him the defense because we would view the involuntariness of his landing on his victim's head as irrelevant: He engaged in a voluntary act when he threw himself out of the window and that voluntary act qualifies as the actus reus required by a homicide—"causing another person's death."

Interestingly enough, however, this argument against the legitimacy of contrived defenses does not work very well when applied to contrived self-defense. The problem has to do with proximate causation. We are going to feel uneasy saying that about the defendant who goaded his victim into an ill-considered attack, to which the defendant then responds with lethal force, that he (the defendant) inaugurated a chain of events, at the end of which he may have had a legitimate right of self-defense, but at the beginning of which he did not. In other words, proceeding as we did with necessity and involuntary acts, does not seem to work: That is because the chain of events he inaugurated did not proximately trigger the victim's death. The victim's response was an intervening event.

We might feel uneasy, but it certainly wouldn't be absurd. Whether the provoked victim's actions qualify as an intervening act that breaks the chain of proximate causation between the defendant's provocation and the victim's death is an uncertain matter under the usual understanding of the intervening cause doctrine. The archtypical intervening act is one in which the intervening actor intends to bring about the consequence he brings about. Since the victim here quite obviously was not trying to bring about his own death, he is not an archtypical intervening actor. But the intention to kill the provoker quite possibly is close enough to the paradigmatic case for the doctrine to still apply.

But one might try an entirely different tack altogether to argue for the impermissibility of the provoking defendant's making a legitimate self-defense claim. Self-defense is legitimate if it is truly necessary for the defendant to save his life. But if he provoked the victim, was it truly necessary? Could he not simply have saved his life by not provoking him, rather than by provoking and killing him?

The difficulty with this argument becomes apparent if we try to make the time period involved in the process of goading a very extended one. Suppose a woman embarks on a relationship with a man who has some track record in starting to behave abusively toward his partners many months, or even years, into the relationship. She is well aware of this. Do we want to say that she loses her right of self-defense on the ground that she could have avoided the whole problem by not entering into the relationship with him in the first place? What this brings out is an important ambiguity in the notion of necessity as it enters into the definition of self-defense. We don't really mean that you are only entitled to self-defense if that is logically necessary in some absolute sense. Necessity is relative to background circumstances. And we are surely willing to restrict the background circumstances that we allow as being relevant, most especially by only focusing on recently developed background circumstances. Thus if the agent provocateur began his operations sufficiently in advance of the final denouement, it will be very hard to argue that his use of lethal force was therefore not necessary in that it could have been avoided by not getting into a relationship with the potential terrorist in the first place.

A third way to argue for depriving the provoking actor of his right of self-defense is to say that what we have here is a form of entrapment. In a literal sense, we certainly do. But in a legal and moral sense, the situation is far from clear. What we are talking about here certainly is not the entrapment defense in the strictly legal sense—where it is simply an argument someone is allowed to raise against being convicted of a crime into which the government entrapped him. This is not a case of a government seeking to punish someone for something they entrapped him into doing, for the simple reason that they killed him before the issue of punishment even arose. Moreover, if they had tried to arrest and prosecute him, he would have had trouble making a doctrinally sound entrapment defense because it is usually restricted to non-violent crimes, and

putting even that technical issue aside, he might have had trouble proving the lack of predisposition that the entrapment defense requires. But more important than any of these slightly formalistic points, is the fact that the entrapment defense is hardly an uncontroversial moral principle. In fact its moral footing continues to be much debated, and is if anything more widely questioned than the restriction of the right of self-defense with regard to provoking agents. In a similar vein, one might try to make some sort of estoppel argument against the provoking agent's trying to claim self-defense, but estoppel arguments are on an even less sound footing than entrapment, from which in this case they seem barely distinguishable.

To the general skepticism about allowing the provoking agent a self-defense claim it is perhaps best to respond, though, with something more general that shows why what the provoking agent does probably has to be tolerated.

Elsewhere I have argued—in a joint paper with the economist Alvaro Sandroni—that all rule-based decision-making produces certain kinds of cycles and that these cycles necessarily invite maneuvers of the kind that the provoking agent is engaged in. Let me here give a very brief sketch of that argument.

First, a bit of background. Sandroni, who constructed the model on which our argument builds, had been interested in the kind of cyclical choices that psychologists often hold up as the hallmark of human irrationality. Here are two telling examples.

(1)When someone is given the choice between watching movie X and watching movie Y, he chooses to watch movie Y. When given the choice between watching movie Y by himself or with a handicapped person, he chooses to watch it with a handicapped person. But when given the choice between watching movie Y with a handicapped person or movie X, without a

handicapped person, he chooses movie X. In other words, his choices are cyclical. Why is that? What he is doing seems bizarre and irrational, at least from a purely formal point of view. Informally, it is not so hard to make out what is going on here. The person would be embarrassed to admit not wanting to watch a movie in the presence of a handicapped person, which is why he chooses to watch it *with* him rather than *without* him. But given the opportunity to escape doing so by watching an entirely different movie, which he has the option of watching all by himself, he has no trouble doing so, because he can say that he happened to prefer that movie, though of course we know the opposite to be true, since we know that given the opportunity he would choose Y over X.

(2)When someone has to choose between staying in the office to work, and going home to relax instead, he would choose home. When given the choice between going to visit a sick friend in the hospital or staying at home to relax, he would choose to visit his friend in the hospital. But when given the choice between visiting the sick friend in the hospital or working at the office, he chooses work at the office. Once again we have a cycle, which at least on a formal level seems strange and perplexing. Of course informally, we find nothing strange and perplexing about it. The person is too embarrassed to admit that he would rather stay away from the hospital, and since he can only do so if he has work as an excuse, he chooses to stay at the office.

Sandroni pointed out that a natural way to describe, or model, what is going on here is to think of the chooser as maximizing his preferences subject to a constraint that requires him to respect, first, the usual feasibility constraints (also known as the budget constraint), i.e. only choose something that is actually available to him, and second, and this is where Sandroni's interesting innovation comes in, the requirement that one's choice be "rationalizable," his term

for a choice that can be justified in terms of applicable legal or moral rules. That extra constraint somewhat unexpectedly gives rise to cycling. The cycling in turn gives rise to the possibility of certain strategic actions that allow one to attain a seemingly forbidden objective by being a little manipulative: One is supposed to visit his friend, but one can avoid doing so by making sure not to be home but at work. One is supposed to not discriminate against a handicapped person, but one can nonetheless do what comes to the same thing, by choosing a movie different from the one the handicapped person is going to be watching, and so on.

As we point out in our joint paper, these situations have ready analogues in the law, most especially in the law of criminal law defenses. Consider the law of duress. If I am threatened with torture, unless I commit a rather serious crime, including something as serious killing several people, I might well qualify for the duress defense if I submit: The law cannot expect me to be a hero, the explanation would presumably run. By contrast, if I were threatened with the destruction of a manuscript I have devoted my life's work to, unless I commit the identical crime—killing several people—I naturally will not qualify for the duress defense if I submit. Now suppose that I have in the past endured significant torture just to save the manuscript. Then it seems we just landed in another cycle, this time a legal one. As between suffering the destruction of the manuscript or being tortured, I would rather be tortured. As between torture and committing several killings, I would rather carry out the killings. Yet as between committing several killings and suffering the destruction of my manuscript, I would rather allow my manuscript to be destroyed. This cycle too gives rise to possibilities of circumvention, like the psychological cycles considered earlier.

For instance, imagine that I am able to protect myself against the people who are threatening the destruction of my manuscript (unless I do what they want) by borrowing a lot of

money from a loan shark which I then use to hire people to safeguard my manuscript. When it comes to paying off the loan shark, I don't have the funds, as a result of which he threatens to torture me unless I commit several killings in his behalf. Now I am likely to qualify for the duress defense. If one tried to create doctrines that made such exploitation of legal cycles impossible, they would turn out to be tantamount to the most extreme and unyielding kind of utilitarianism—the sort that would inflict punishment every time I choose to do anything at any moment in time that does not maximize overall social utility, more or less understood in traditional Benthamite ways. (If there are other ways to block the exploitation of legal cycles, they are going to be at least equally unappealing.)

There are entirely uncontroversial examples of this kind of cycle in the law of self-defense. Suppose a robber says to me "Your money or your life." I stand fast and refuse to give him my money, whereupon he attacks me, whereupon I kill him, which would of course be perfectly legitimate on my part. Hidden within this simple interaction is the exploitation of a cycle much like the one involved in duress. Given the choice between killing a thief and losing my money, I am required by law to lose my money. Given the choice between losing my money and losing my life, I am free to lose my life. Given the choice between losing my life and taking the attacker's life instead, I am free to kill my attacker. The result is that if I want to get around the prohibition against choosing my money over my attacker's life, I can do so by first choosing to put my life at risk for the sake of my money, and then saving my life by killing the attacker.

Something analogous is really being done when I contrive to create the circumstances of self-defense by putting an agent provocateur in the victim's vicinity. Although the details can surely be quarreled with, and to achieve complete immunity certain strictures of indirectness

might need to be observed, if the government wants to be completely in the clear morally and legally, but now we are just haggling about the price, as George Bernard Shaw would say.

Minimizing Rights Violations

A standard argument for a targeted killing is the following. If we do not eliminate the potential attacker, or for that matter the holy man, on the grounds that we would be violating their rights if we did, we are facilitating a much greater number of right violations down the road. To begin with, there are the rights violations carried out by the would-be attacker and by the followers of the holy-man, but those are not the only ones to worry about. There are also the rights violations that are likely to be carried out by our side in the course of any conflict. Not that we think we are justified in those—if we were they wouldn't be rights violations—it is just that we can statistically predict that there will be many of those, whether we like it or not. That's just how it is. If one cares about rights, surely one should prefer a single rights violation now in lieu of many more later on.

This kind of argument is what Robert Nozick called a utilitarianism of rights, and he viewed is as one which an advocate of rights, or at least a deontologist, would generally reject. If one is a deontologist, one ought to respect rights, not seek to minimize rights violations. That means, somewhat paradoxically, doing things that involve avoiding violating a right here and now, even if in the long run the person, or governmental entity, respecting those rights, would be causing fewer rights violations, indeed would himself be committing fewer rights violations, if he did not respect the a certain right here and now, such as the right not to be the victim of an unjustified targeted killing.

I am willing to go along with that. What I am not willing to go along with is a conclusion usually drawn from it: namely that it would be better if one did not violate a right than if one did. "Would be better" turns out to have several possible senses, and while it might well be better not to carry out the targeted killing *in the sense that it is what one ought to do*, it turns out, on reflection, not to be better in some other, equally important sense—namely in the sense that one might well be guiltier of fewer rights violations if one were willing to engage in such a violation here and now. In other words, if someone were to review the moral ledger of someone who decided to respect a certain right here and now at the expense of violating many more such rights later on, and compared that to the moral ledger of someone who ignored that self-same right here and now, so as to forestall having to commit many more rights violations later on, this last person might well end up with the better moral ledger.

To see this more clearly, let's consider for a moment the venerable *Regina v. Dudley and Stephens*, the lifeboat case in which Thomas Dudley decides to kill the youngest of the sailors, Richard Parker, so that he and the other two sailors, Stephens and Brooks, would have a shot at survival by feeding on his remains. Let us suppose, not implausibly, that Dudley had acted with criminal recklessness in taking a comparatively slight boat like the Mignonette on an ocean voyage to Australia—so criminal in fact that if Parker, Stephens and Brooks had all died in the resulting shipwreck, he would be guilty of their murder. Finally, let us suppose that we agree with the court's judgment that Dudley's killing of Parker was wrong, that it was too much like the classic utilitarian horror stories in which killings that achieve net savings of lives are nevertheless impermissible: cases like the organ-harvesting surgeon who saves five people one of whom needs a heart, two of whom need kidneys, and two of whom need lungs, by killing a

fifth; or the prison warden who surrenders a prisoner to a lynch mob so as to avert a riot in which many more innocent lives would be lost.

Let us now compare what would have happened if Dudley had done the right thing, and abstained from killing Parker, and what would have happened if he had done what he actually did, killed him. In the first case, he would have been guilty of three murders—that of Stephens, Brooks and Parker. In the second case, he would only be guilty of what he actually was found guilty of, the murder of Parker. In short, Dudley's legal position is improved by doing the wrong thing.

What I am imagining here is not as unusual as it might sound. After all, all we really need to construct this kind of scenario is to take any one of the classical utilitarian dilemmas—like the organ harvesting case or the lynch mob case—and to imagine someone who has wrongfully inflicted some kind of serious harm that he could mitigate if he engaged in one of these unpalatable utilitarian tradeoffs. Imagine for instance a ruler who first whips up the frenzy of the crowd against an innocent prisoner, but then wants to backtrack: he will minimize his blameworthiness by agreeing to the execution of the innocent prisoner, so as to prevent the lynch mob from doing even worse things, for which he could otherwise be blamed, given his role in whipping them up.

Minimizing Innocent Deaths

The most natural and unsophisticated sounding argument a state will make for undertaking a targeted killing is of course the simple-minded one that it saves lives—or rather, that it achieves a net saving of lives. Let's focus here in particular on its attempt to defend the killing of someone like the holy-man in the second of my examples. To recall: The situation I

have been envisioning regarding the holy-man is one in which his followers are about to launch an attack against us, we in turn would be able to defend against the attack by suitably preemptive action, and in the course of such preemptive action would end up killing many innocents who happen to be bystanders of the conflict. The government will argue that by killing him, who is really only "technically" innocent, we have saved many others from the deadly consequences of the attack which his malign influence would unleash. Such an argument is usually dismissed because it sounds at first indistinguishable from the kind of standard utilitarian claim, which many would dismiss by drawing the analogy to organ-harvesting doctors and lynch-mob appeasing prison wardens. But in fact our situation is different in some crucial respects and the net saving of lives argument here on much stronger, even if not entirely unassailable ground.

To see why, let's begin by adapting the infamous trolley example to our kind of situation. Imagine that we could take preemptive action against would-be attackers by building a long ramp, an incline, down which we could chase a trolley which would run over, and kill, the innocents that happen to surround our would-be attackers and in due course would kill our would-be attackers as well. This is something we are obviously entitled to do.

Now add the familiar second track to this incline, such that if we were turn the trolley, it would run into our holy-man instead, killing him, sparing the innocents, and stopping in their tracks the attackers who depended on the sustenance they drew from his message. Would it not be perfectly acceptable to turn the trolley here, just as it is in the standard trolley case?

To be sure, we cannot ordinarily use this kind of a trolley arrangement to liquidate just anyone by setting things up in such a way that we take aim at one group of innocents and then avoid hitting them by turning the trolley onto a second track. But that is because we are

ordinarily not entitled to take aim at a group of innocents in the first place. Here things are different, because we are entitled to take aim at them, in the course of exercising our right of self-defense. Turning the trolley now seems an entirely different matter.

Most of the details of the trolley arrangement are surely inessential for this argument to work. If we simply directed a strike against the holy-man, that seems just as permissible as engaging in the trolley maneuver—once one has registered that it is not essentially different from the trolley case. And so it would seem that at least in a context in which we are entitled to take innocent lives in the course of defending ourselves—and probably for other reasons as well--we will often be entitled to cash in that chit for other innocent lives, in this case for a targeted killing.

Some people are going to find my analysis here too coarse-grained. There are many situations in moral theory in which chits like this just don't work—you really can't cash them in, however Pareto-optimal it would be if you could. A good illustration is what one might call the problem of delayed self-defense: Take an ordinary case of self-defense, in which rather than killing my attacker outright I allow myself to be injured by him, hoping to survive his attack without having to kill anyone. Having been injured badly, however, it turns out I can only survive with a heart transplant, for which he is the only potential donor. Can I avail myself of his heart? Can I claim that this is a kind for extended self-defense? Some people might say yes (Michael Moore has so argued) most however would say No. Not killing him in the first place did not give me a chit I can cash in later on, however mutually beneficial it would be for him and me alike if I could.

But it hardly follows that such chits never work, and at the moment I see no reason why they shouldn't work in the case of the holy-man.

Root Causes

I said at the outset that in this essay I would try to advance two claims, first, that there is more to the standard governmental strategies for coping with objections to targeted killings than might at first appear. This I consider myself to have done. Second, that there is something interesting to be said about the dissonance between the actual strength and the apparent weakness of those strategies. This I will now try to do in a very, very sketchy way.

What exactly is it that gives us such pause about the three strategies I have discussed for getting targeted killings done? Even if you are persuaded that by pursuing those strategies one successfully manages to get such killing done without running afoul of either law or morality, you are bound to feel puzzled about the fact that they work. There seems something strange and troubling about them. What makes them strange and troubling in my opinion is that they reveal features of legal rules and legal reasoning that we would regard as irrational in the course of ordinary decision-making. The strategy of provoking an attack so as to then be able to counterattack depends on the ubiquity of cyclical, or intransitive, choices in a legal system, something which many would regard as the hallmark of an irrational person, if an individual decision-maker were to exhibit it. The strategy of minimizing rights violations goes against a desirable feature of rational decision-making that decision theorists like to call non-monotonicity. We expect sound moral and legal choices to have the basic property that if we choose to do the right rather than the wrong thing, we end up in a legally and morally better position than someone who chooses to the wrong rather than the right thing. (We expect moral

choices to be monotonically related to overall moral status.) Finally, the strategy of minimizing deaths seems to violate what is sometimes called the independence of irrelevant alternatives. As between killing the holy-man and not killing him, the latter is the only acceptable option, but somehow introducing a further option, that of killing other innocents, then changes the relative ordering, i.e. the relative permissibility, of those two.

There is nothing accidental about the fact that legal reasoning violates these canons of rational decision-making. It has long been known that collective decision-making routinely violates these canons. Those are the familiar paradoxes of social choices—the impossibility theorems of Arrow, Sen and others. It has also come to be accepted that individual decision-making, if it seeks to reflect a multiplicity of judgmental criteria, closely resembles collective decision-making and produces analogous paradoxes. But that is exactly what legal decision-making is: individual decision-making based on a multiplicity criteria. It thus cannot avoid those paradoxes either. Our paradox-fraught assessment of targeted killings is just a special case of all that.

What I have so sweepingly asserted in this last paragraph is really a mouthful. I certainly don't consider it be obviously right. Indeed I have made the defense of this thesis the subject of an entire book called *Why the Law Is So Perverse*. This larger claim about the root causes of the dissonance between the actual solidity and the apparent feebleness of the arguments for targeted killings is thus one the reader has to either take on faith—or read the book.