questions that would have to be settled if we are to know whether the view that undergirds opportunity pluralism is substantially different from views already on offer (such as Martha Nussbaum's capability view or Richard Arneson's desert-sensitive prioritarianism).

The fact that Fishkin leaves the foundations of opportunity pluralism insufficiently worked out has another implication, this one important for the constructive part of his project. Without a more fully worked-out view that underlies opportunity pluralism, we are unable to know what the latter's determinate implications are beyond the relatively uncontroversial cases. Fishkin's observations about the maldistribution of real opportunities for well-being in a context like the United States, where the lack of a university degree stunts people's ability to access a vast range of valuable goods, including, crucially, a well-remunerated job, are well-taken. But these points (whose relevance for liberal democracies other than the United States is limited), have been made before. Fishkin could have supported more radical and novel recommendations as part and parcel of his vision of opportunity pluralism. For example, he could have argued in favor of a more radical reconfiguration of work of the kind advocated by Paul Gomberg, so as to obtain a fairer distribution of the real opportunity for meaningful work. This would be congruent with his endorsement of a perfectionist view of well-being, his view of human development, and his reservations about specialization. Fishkin could additionally, or alternatively, have defended the possibility that priority of opportunity for well-being tells in favor of supporting parenting, which, he notes, is one of the few very valuable activities which are open to everyone (see Arneson,"What Do We Owe to Poor Families?"; Anne L. Alstott, *No Exit: What Parents Owe Their Children and What Society Owes Parents* [Oxford: Oxford University Press, 2004]). Fishkin does not consider at all whether opportunity pluralism supports subsidizing valuable unpaid work like parenting. As for Gomberg's proposal, Fishkin mentions it only to quickly set it aside in the name of efficiency (52). Yet it seems that curbing the market and compromising efficiency to some degree are necessary to further any of the four goals that constitute opportunity pluralism. The extent to which the latter is market-friendly is one of the issues in *Bottlenecks* which remain unresolved and cannot be resolved in the absence of a more fully worked-out principle of fair opportunity which Fishkin says should replace familiar ideals of equal opportunity.

Serena Olsaretti
ICREA-Universitat Pompeu Fabra

Frowe, Helen. *Defensive Killing.*
Oxford: Oxford University Press, 2014. Pp. xii+227. $50.00 (cloth).

Although this is not a long book, it is dense with meticulous argumentation. Its first four chapters discuss important foundational issues in the morality of individual self- and other-defense. The remaining four apply Frowe's conclusions to questions of permissible killing in war. I have found the book a source of great stimulation and inspiration. It has revealed many problems and complexities in the morality of defense of which I was unaware and has forced me to rethink var-

ious issues concerning liability, proportionality, and necessity. Rather than describe the book's many excellences, I will devote the short space of this review to a discussion of two issues on which Frowe and I disagree—proportionality in defensive harming and the liability of merely apparent threateners—in the hope of further advancing the debates.

Frowe's view of proportionality diverges in various ways from familiar accounts in the literature. She argues that the amount of harm it can be proportionate to inflict on a threatener does not always vary with the degree of harm he would otherwise cause; nor does it vary with the degree of his responsibility for the threat he poses. She also argues that proportionality does not depend on the probability that defensive action will be effective, provided it has some chance of success.

The cases in which Frowe denies that proportionality is sensitive to the amount of harm a threatener would inflict are those in which he is among other contributors to a collectively inflicted harm. She writes that "proportionality is judged not by the extent of one's contribution to a threat, but . . . [by] the magnitude of the threat to which one contributes" (78). Suppose that in a case we can call *Collaborators* (based on Parfit's "harmless torturers"), each of a thousand people maliciously wants to cause Victim to suffer but realizes that the most he can do is to cause Victim to experience a tiny amount of pain for ten hours. Although each would have inflicted his tiny pain whatever the others did, they collaborate by inflicting their pains at the same time, causing Victim to suffer great agony for ten hours. Among these people is Enemy. According to Frowe, Enemy is liable to the same harm that he would be liable to if he were inflicting all of the pain himself, provided that this would have some chance of eliminating his contribution to Victim's suffering.

Frowe is right that the harm to which a threatener contributes is relevant to his liability. If these thousand threateners all had different victims rather than the same victim, so that they together would not cause anyone to suffer more than a tiny pain, each would be liable to less harm than that to which he is liable in the example as stated. But it seems that proportionality must take account of both factors—the individual contribution and the total harm to which it contributes—rather than only one.

Compare *Collaborators* with *Contributors*, which is similar except that those who want to harm Victim are unknown to one another. Each intends to inflict a tiny pain on Victim but believes, on the basis of rumors, that many other people will each inflict a tiny pain on Victim during a specific period of ten hours. Each therefore acts independently to inflict a tiny pain during this period in the belief, which turns out to be correct, that it will be a contribution to Victim's suffering great agony. There is no collaboration but this does not make a difference to the amount of harm to which each threatener, including Enemy, is liable.

Next compare *Contributors* with a case in which Enemy discovers that for ten hours Victim will simultaneously experience a large number of tiny pains from natural causes. Enemy adds a tiny pain for this period, so that Victim suffers as much as she does in *Collaborators* and *Contributors*. This seems morally no different from what Enemy does in *Contributors*. It does not make a difference to his liability whether the other pains to which he knowingly contributes are inflicted by agents or result from natural events.

Finally, in *Kidney Stone*, Enemy learns that Victim is already suffering the pain of a kidney stone. He adds his own tiny pain for ten hours, bringing Victim during

this period to exactly the same level of suffering as in the previous three cases. Again, it does not seem to matter to Enemy's liability whether the pain he inflicts begins at the same time as, or is precisely coextensive with, the other pains to which it is added.

Return now to Frowe's view of proportionality. Suppose that it would be proportionate for Victim to kill Enemy to prevent him from acting on his own to cause Victim to suffer agony for ten hours. On that assumption, Frowe's view implies that in *Collaborators*, it would be proportionate for Victim to kill Enemy as a means of preventing him from adding his tiny pain to those inflicted by the other 999 threateners—and indeed proportionate to kill all the others as well. If, as I have suggested, there are no morally significant differences between Enemy's acts in the four cases, Frowe's view also implies that it is proportionate to kill Enemy in *Kidney Stone*. I think, however, that it is disproportionate to kill Enemy in *Collaborators*, when that would make almost no difference to Victim's suffering. But it is even more obviously disproportionate and therefore impermissible to kill Enemy in *Kidney Stone*, which involves conditions that are more familiar than those in *Collaborators*.

I have described the thousand inflictors of tiny pains, including Enemy, as malicious. Frowe's view is more plausible if we make this assumption. Yet she says she finds "the idea that proportionality is sensitive to moral responsibility quite puzzling" (though she accepts that one threatener's "greater moral responsibility means that harms to him count for less than harms to" another in determining whether an act of defense satisfies the necessity constraint) (173, 168). If she is right that proportionality is not sensitive to degrees of responsibility, it should make no difference to the harm to which the thousand inflictors are liable whether they are highly culpable or only minimally (that is, nonculpably) responsible for the threats they pose. According to her view, then, it is proportionate to kill all thousand in *Collaborators*, and to kill Enemy in *Kidney Stone*, even if the explanation of why they will otherwise inflict tiny pains is that they have been deceived into thinking that it is permissible for them to do so. This seems highly implausible. (Frowe contends that threateners who are wholly nonresponsible are not liable to defensive harm but can nevertheless have a duty to bear defensive harm. Although this may be conceptually distinct from the claim that they are liable, it is substantively the same. If a nonresponsible threatener's duty to bear defensive harm is enforceable, it would be proportionate, and permissible, to kill Enemy in *Kidney Stone*, even if he has been involuntarily drugged and is not responsible for his action.)

Although I see no way that her view can avoid the implication in *Kidney Stone*, she suggests a way in which it need not imply that it is proportionate to kill all thousand in *Collaborators*. "There are," she writes, "general moral reasons not to cause massive amounts of harm that apply even when those harms are directed at people who are liable to them" (which, on her view and mine, entails that they are proportionate) (209). She does not, however, say what these general moral reasons are, or explain how they can make it impermissible to inflict harms to which the victims are liable, particularly when there are no other means of preventing a threatened harm. I suspect her claim derives from her view that liability is not a justification but only the removal of one moral barrier to harming (106). I think, by contrast, that liability is a positive justification because a person can be liable only

when harm is unavoidable, so that if the liable person is not harmed, someone who is not liable will be harmed instead and that will be less just. Hence, if a large number of people are actually morally liable to be killed, I do not think it can be impermissible to kill them just because of the magnitude of the overall harm to them.

To avoid these implications, we should accept that how much defensive harm a threatener is liable to suffer is sensitive both to the amount of harm he alone would otherwise inflict and to the degree to which he is responsible for that threatened harm. This is compatible with the harm it is proportionate to inflict on a threatener also being sensitive to the total harm to which he would contribute.

Another way in which Frowe's view of proportionality differs from more orthodox views is that she denies that whether an act of defensive harming is proportionate can depend on the probability that it will succeed. She claims that proportionality is a relation between the harm inflicted defensively and the harm the defensive act *would* prevent *were it successful*. Her arguments to show that proportionality is not sensitive to probability are that if it were, (1) one victim could be permitted to act in self-defense while another would not, even though their situations were identical except for a difference in probability, and (2) it could be disproportionate to inflict a lesser harm with a lower probability of success but proportionate to inflict a greater harm with a higher probability. The deeper source of her view, however, is that she doubts that defensive harming can violate the rights of threateners just because the probability of success is low. If the probability of success is an element of proportionality and if, as Frowe and I agree, a threatener cannot be liable to disproportionate harm, then a low probability of success can exempt a threatener from liability to defensive harming. But Frowe resists the idea that people can escape liability simply by being efficient threateners.

I think that a threatener *can* be wronged by defensive harming that has a low probability of success, even when the harm would be proportionate if it had a higher probability of success. The examples to which Frowe appeals are all of culpable threateners but, as we have seen, it ought not to matter, on her view, if the threateners were instead only minimally responsible. Yet this seems to make a difference. Suppose Threatener is only minimally responsible for threatening Victim with a broken arm. If it were certain that Victim could prevent the breaking of her arm by breaking Threatener's arm, it would be proportionate for her to do that. But if breaking Threatener's arm would have less than a one percent probability of preventing him from breaking Victim's arm, it seems that it would be disproportionate for Victim to try to defend herself this way (or rather, as I will explain below, there are some instances in which this is univocally true and others in which there is a sense in which it is true). The probability is just too high that, if she does break his arm, she will be pointlessly harming a morally innocent person.

Frowe responds that if Victim breaks Threatener's arm and, improbably, succeeds in preventing him from breaking her arm, it will be clear that her act was not disproportionate. This is an instance of a pervasive problem in ethics—namely, the divergence between expected outcomes and actual outcomes. One response here is to understand proportionality the way Parfit understands permissibility, by recognizing that it has both evidence-relative and fact-relative dimensions. An act of defense that has a low probability of success but nevertheless succeeds may be disproportionate in the evidence-relative sense but proportion-

ate in the fact-relative sense. Each of these facts may be relevant to different issues. For example, whether the act is proportionate in the evidence-relative sense is what must be action-guiding and is also relevant to whether Victim is blameworthy.

The most important question is which dimension of proportionality is internal to liability. There are four possibilities.

(1) When a defensive act is proportionate in both senses, in part because its probability of success is high and because it succeeds, the threatener is clearly liable (if other conditions of liability are met). Frowe agrees.

(2) When an act is disproportionate in both senses in part because the probability of success is low and the act fails, I think Threatener was not liable and has been wronged. Frowe disagrees. I think the act is fact-relatively disproportionate because there is no relevant good effect to outweigh the bad. Frowe says that the act was neither proportionate nor disproportionate because it was not a means of defense at all (152)—though it is unclear how this is compatible with her claim that proportionality is a relation between the harm an act inflicts and the harm it would prevent were it to succeed. (On this latter view, defense can be fact-relatively proportionate even when it fails. This is the basis of Frowe's denial that Threatener has been wronged. On this understanding of her view, we can in principle agree about when *successful* defense is fact-relatively proportionate.)

(3) Defense may be evidence-relatively disproportionate in part because the probability of success is low but fact-relatively proportionate because it succeeds. In such a case, it is hard to believe that the threatener is wronged, even though I think the victim acted impermissibly in the evidence-relative sense and thus is blameworthy. I therefore think that fact-relative proportionality is sufficient for liability. But that does not imply that it is internal to liability. The claim that proportionality is internal to liability means that it is necessary for liability, so that no one can be liable to harm that is disproportionate. If Threatener can be liable to harm that is evidence-relatively proportionate but fact-relatively disproportionate, fact-relative proportionality is not internal to liability.

(4) I think evidence-relative proportionality can be sufficient for liability. Suppose an act of defense is evidence-relatively proportionate in part because it has a high probability of success but fact-relatively disproportionate because it fails. Frowe says that a threatener can be liable to defensive action that fails. (180) That must be right. Suppose that Victim engages in defensive action that it is reasonable to believe will succeed, and would be proportionate if it were to succeed, but that nevertheless fails. Threatener cannot reasonably claim that he has been wronged. A person can be liable to defensive harm when (1) harm is unavoidable, (2) it is possible to distribute the harm in different ways, and (3) it would be more just to harm him than to allow someone else to be harmed. If an attempt to harm him fails, that does not entail that he was not liable. That liability is instrumental does not imply that an act with a liability justification necessarily succeeds but only that there is a liability justification for the act's instrumental aim. If the act fails, harm has not been distributed in accordance with liability, but Threatener was nonetheless liable to the attempt.

Evidence-relative proportionality is therefore sufficient for liability and fact-relative proportionality is thus not necessary for liability. But cases in the third category above show that fact-relative proportionality is also sufficient, so that evidence-relative proportionality is not necessary either. What is internal to liability, therefore, is the disjunct of evidence-relative and fact-relative proportionality—that is, what is necessary for Threatener to be liable to defensive action is that the action be proportionate in either the evidence-relative or the fact-relative sense.

There are six probabilities that can be relevant to proportionality. These probabilities are: (1) that a threatened harm will occur in the absence of defense, (2) that a threatened harm will be of a certain magnitude, (3) that a defensive act will succeed, (4) that a defensive harm will be of a certain magnitude, (5) that defensive action will harm innocent bystanders, and (6) that harm to bystanders will be of a certain magnitude. Probabilities 3 and 4 are relevant to narrow proportionality, 5 and 6 to wide proportionality, and 1 and 2 to both. All raise problems that I cannot discuss here.

I have claimed that when a defensive act is disproportionate in both senses because the probability of success was low and the act failed, Threatener was not liable to the harm inflicted, even though it had a chance of success. This is intuitively most plausible when Threatener is minimally responsible (or nonresponsible). But it raises the question of whether Threatener has a right of defense against the harm to which he is not liable.

Consider first a different problem raised when people appear to pose a threat but do not. What I say about this problem provides an answer to the question about self-defense by threateners who are not liable because defense against them would be disproportionate.

Suppose Enemy puts one bullet in a two-chambered gun with a lock that allows the trigger to be pulled only once. He spins the chamber and culpably points the gun at Victim with the intention of pulling the trigger. Neither knows which chamber the bullet is in, though it is in fact not in the chamber that the hammer will strike if the trigger is pulled. Victim can try to kill Enemy, but Enemy has another weapon he will use to kill her, though only if she attempts to defend herself.

For Victim, harm is avoidable, though she cannot know this. If she attempts to defend herself, harm will then be unavoidable for Enemy, for either she will kill him or he will kill her in self-defense. Because liability is instrumental, there can be no liability in situations in which harm is entirely avoidable. But whether harm is unavoidable is relative to agents. When Victim chooses to try to kill Enemy, harm is avoidable for her but not for him.

Frowe would say that Enemy is liable to be killed. But she accepts that liability is instrumental and that, because he poses no threat, he cannot be liable to *defensive* harm. Yet she says he is liable to "harms that Victim inflicts in the course of trying to defend [herself]. The relevant moral responsibility for forfeiting rights against this sort of harm is responsibility for the fact that Victim believes that if [she] does not kill Enemy, Enemy will kill [her]" (85–86). Yet if liability is instrumental, and thus subject to an effectiveness condition, Enemy cannot be liable to the harms that Victim would inflict; for those harms cannot achieve anything—

and Frowe rightly distinguishes between harm that cannot be effective and harm that could be effective but fails.

What motivates Frowe to claim that Enemy is liable to be killed is her belief that if he is not, then when Victim tries to kill him, she wrongs him and he has a right of self-defense against her—and that seems false. I think, however, that Victim does wrong Enemy, as Enemy does not forfeit his right not to be killed when he in fact poses no threat. That is, he is not liable to be killed because killing him would not be instrumental to the achievement of any justified aim (putting aside considerations of desert, deterrence, the affirmation of Victim's moral status, and so on).

But even though Enemy is not liable to be killed, he has no right of defense. This is because, for him, harm is unavoidable and he bears greater responsibility for that fact than Victim does. Assuming that the possible harms are comparable in magnitude, it is a matter of justice that he rather than she should suffer the harm he is more responsible for having made unavoidable. He is therefore liable to *allow* himself to be killed. (This point may be clearer if we imagine that Enemy cannot defend himself but that a third party with full knowledge can. Although Enemy is not liable to be *killed*, he is liable to be *allowed* to be killed. He has forfeited his right to be saved when the alternative is killing Victim.)

Return now to the question whether a wrongful threatener against whom the only possible defense would be disproportionate in both senses is permitted to harm his victim in self-defense. Assume that Threatener is not liable to disproportionate defensive action and that Victim acts wrongly in attempting it. It may nevertheless be true that Threatener is not permitted to defend himself. For he is in a situation in which harm is unavoidable and he may bear greater responsibility for that than Victim does. If so, while he is not liable to be harmed by Victim, he may be liable to allow himself to be harmed—unless, perhaps, he could defend himself by causing her only minor harm and would not harm her further.

This review does not come close to doing justice to the richness of Frowe's intricately argued, insightful, and challenging book. I have benefited enormously from thinking as carefully as I am able about the positions she defends and the arguments she gives for them. I am confident that the same will be true of others who read this splendid book.

Jeff McMahan
University of Oxford

Harel, Alon. *Why Law Matters*.
Oxford: Oxford University Press, 2014. Pp. 256. $65.00 (cloth).

Alon Harel's *Why Law Matters* is an unusual book in many respects. To begin with, few publications are reviewed both in *Ethics* and in the *Washington Post*. This fact suggests—rightly, in my view—that the book will be of interest both to the author's fellow academics and to sophisticated general readers who pay attention to the more theoretical aspects of political and legal affairs. Moreover, Harel's book is also unusual among recent works in political philosophy and the philosophy of law in that it defends a noninstrumentalist view of what justifies political and legal