12 Chance, consent, and COVID-19

Ryan Doody

12.1 The case for (and against) mandatory lockdowns

On March 16, 2020, in response to the spread of SARS-CoV-2 (COVID-19), the premier of the Australian state of Victoria declared a state of emergency. Using special powers, he imposed a number of restrictions intended to slow viral transmission. These restrictions – colloquially referred to as a *lockdown* – required residents of impacted areas to remain in their houses and only permitted them to leave under limited circumstances to conduct essential activities. The measures also imposed a mandatory 14-day isolation period on travelers entering the state, and a moratorium on mass gatherings. They required the closure of nonessential businesses and schools. Beginning in August 2020, there was also a nightly curfew. In total, there would be 6 lockdowns in Victoria, over a period of 18 months, lasting for a cumulative total of 262 days.

Victoria was far from the only region impacted by mandatory lockdown measures during the pandemic, but given both the extent and duration of the restrictions, it serves as a particularly stark example. In particular, its example evokes a concern held by many liberally-minded witnesses to lockdown measures: namely, that they appeared to involve a somewhat tragic clash between serious public health interests and considerations of individual rights. Vocal opponents of the lockdown measures, and of the nocturnal curfew, in particular – including Victorian Liberal Party MP Tim Wilson, the plaintiff in a lawsuit contesting the measures (*Loielo v Giles*), and countless protesters – complained that the lockdowns infringed upon important *civil liberties* such as the rights to free movement and assembly. In response, the premier of Victoria, Daniel Andrews, defended the lockdown: "[I]t's not about human rights; it's about human life." The measures, Andrews argued, were necessary to suppress transmission of a lethal virus – they were necessary to protect human life.

Supposing that mandatory lockdown measures are sometimes necessary to prevent serious harm to the public, when is the state justified in using them? It's not implausible to think that, insofar as lockdowns

DOI: 10.4324/9781003310129-15

infringe important civil liberties, a high bar of effectiveness must be met – they must be necessary to prevent *a great deal* of grievous harm – to be justifiable. As a general rule, if a policy infringes on someone's rights, it shouldn't be implemented – unless the considerations in favor of doing so are of significant enough moral weight. After all, the thought goes, rights are no small thing; their infringement shouldn't be taken lightly. And an important role (indeed, on some views, the primary duty) of the state is to protect citizens' rights.

In Victoria, people like Tim Wilson, who opposed mandatory lockdown measures, argued that the very high bar for licensing rights infringements had not been met. They variously questioned how *effective* lockdown measures actually were, whether they were really *necessary* to achieve desired results, and disputed operative assumptions about what the threshold for public good should be in order to justify rights violations. On the other hand, proponents of mandatory lockdown measures, like Dan Andrews, argued that the bar *had* been met: e.g., the social damage averted and the number of lives to be saved were significant enough to override the strong presumption against violating rights – "it's not about human rights; it's about human life."

However, casting the debate in these terms, as representative as it might be of public discourse at the time, arguably concedes too much to those who, like Tim Wilson, opposed mandatory lockdowns. For, the thought goes, despite appearances, lockdown measures *do not* violate important rights. Why? Because no one has a right to impose a significant risk of grievous harm on others (this follows from the right that each of us enjoys against such impositions) – and, in the context of the COVID-19 pandemic, participating in the sorts of activities restricted by the lockdown measures would impose such a risk. Going to a bar and "getting on the beers" (to use the premier's memorable expression) would arguably impose such a risk on others. Classical liberal thought tells us that while you might enjoy a presumptive right to wave your arms around in such-and-such a fashion, if doing so on this particular occasion also amounts to doing me grievous harm, then you don't have such a right (at least, on this occasion) after all. And as for waving arms, so, we might say, for getting on the beers. Further, as mentioned earlier, the state – while perhaps lacking adequate justification to regulate issues of personal morality – plausibly *does* have an interest in regulating behavior that violates important *rights*. And so, if citizens have a right against being subjected to a significant risk of grievous harm, then the state is justified in using coercive force to prevent the imposition of such risk.

On this view, the "right against risk impositions" plays two roles in justifying mandatory lockdown measures. First, it *blocks* the objection that such measures infringe upon important rights. Second, if the state is justified in intervening to prevent rights violations, it also helps to make a *positive* case for implementing mandatory lockdown measures.

The activities that such measures prohibit are ones that we, in the context of a pandemic, do not have a right to engage in, and, insofar as those activities infringe the rights of others, the state is justified in prohibiting them.

This chapter explores whether this defense of mandatory lockdown measures can succeed. It, first, considers whether the following claim is true and, in so doing, asks how the risk to which it refers is best understood:

> *Right Against Risk*: We have a right against others not to be subjected to a significant risk of grievous harm.

Second, it considers the following objection:

> *Waived Rights*: Even if we do have a right against being subjected to a significant risk of grievous harm, those who voluntarily choose to engage in the activities prohibited by mandatory lockdown measures, if fully appraised of the risks involved, effectively waive this right.

Following this objection, if everyone who, e.g., gathers at the bar effectively *waives* their right against the risk of grievous harm that might result from contracting COVID-19, no one at the bar is in danger of infringing the (unwaived) rights of others. And thus mandatory lockdown measures cannot be justified on the grounds that they prevent rights violations – rights that have been waived cannot be violated. And if going to the bar doesn't violate anyone's rights, given the presumptive right to assemble where we please, the state's use of coercive force to prevent us from doing so appears to infringe rights after all.

The chapter argues that this objection fails. What it takes for one to waive such a right is significantly more nuanced than the objection can allow. And, in any realistic case, there will be legitimate third parties whose rights have not been waived.

12.2 A right against risk

It's not implausible to think that we each have a right against being harmed – especially if that harm is significant. Do we each also have a right against being subjected to a *risk* of harm (especially if that harm is significant)? And, if so, how is this right best understood? And would engaging in those activities prohibited by mandatory lockdown measures (e.g., gathering at the bar) constitute a violation of such a right?

These questions, to differing extents, turn on what it is to have a *right*. Opinions, of course, vary greatly. I will not argue for any particular conception of rights here. Instead, let's focus on the following three features that might plausibly characterize the possession of a right.

DUTY: If A has a right against B's φing, then B has (at least) a *pro tanto* duty not to φ.

ENFORCEMENT: If A has a right against B, then it is *pro tanto* justified for A – and, perhaps, certain ordained third parties, like the state – to prevent B from violating this right.

COMPENSATION: If B infringes one of A's rights, then A has a right to be compensated by B in the appropriate way.

A quick word about each. According to DUTY, to have a right is, in part, for others to have a corresponding duty. This duty needn't be absolute; it can be outweighed by other considerations.[1] Also, this duty might best be thought of as a *directed duty* – so that if B were to φ, not only would B do wrong (assuming the duty hasn't been outweighed), B *wrongs* A in particular. Not all directed duties correspond to rights, however – and some philosophers appeal to something like ENFORCEMENT to account for the difference.[2] Even if that's not true in general, ENFORCEMENT is of particular interest given our purposes because the central question of this chapter concerns the sorts of restrictions that the state is justified in imposing. Like the correlative duty in DUTY, the justification one has to prevent B's infringement of A's right is *pro tanto*: it can be overridden by other considerations. Finally, COMPENSATION says that, at the very least, having a right generates a duty of appropriate compensation in the event that that right is infringed. It might be (all things considered) morally permissible for B to φ, and it might be (all things considered) morally impermissible for A (or the state) to *prevent* B from φing, but – if A has a right against B's φing – B owes A some form of compensation for the infringement.[3]

In addition to assuming that we each have a right against being harmed, I will also assume that, during the height of the COVID-19 pandemic, many everyday activities (e.g., gathering at the bar) involved imposing a *risk of harm* on others. This can be true, I contend, even if one isn't actually contagious or incubating the virus at all – all that's required is that one not be *certain* whether they're infectious. That is, we will operate with a *subjective conception* of risk.[4] For that (and other) reasons, it's very hard to say exactly *how much* risk one imposes on another by, e.g., going to the bar. Supposing (for the sake of argument) that one is experiencing no symptoms, the question turns on how likely one is to be an asymptomatic carrier anyway, how likely one is to transmit the virus if one is carrying it, and how likely those who are exposed are to suffer significant harm (e.g., death) from the exposure.

In any case, of particular interest to us are *pure risk impositions*: cases in which someone imposes a risk of an unwanted outcome, but doesn't actually cause an unwanted outcome.[5] For example, suppose that X imposes a risk on Y – by, e.g., exposing them to COVID-19 – that does not eventuate, and that Y never learns of (and, so, is not a source of

psychological distress). Cases like these are of particular interest because, if the harm *did* eventuate or if the action caused harm in some other way, Y's (uncontroversial) right against being *actually* harmed would be infringed.[6] And the question before us is whether Y also has a right against the imposition of a *risk* of harm – and, if so, how that right is best understood.

12.2.1 The risk thesis

Perhaps the most straightforward view (defended by McCarthy 1997) is the following:

> *The Risk Thesis*: We each have the right that others not impose risks of harm on us.

Because, during a pandemic, those activities prohibited by mandatory lockdown measures (e.g., gathering at the bar) impose a risk of harm on others, on this view, such activities infringe the rights of others. Consequently, we each then have a *pro tanto* duty not to engage in such activities (from DUTY), and the state is *pro tanto* justified in preventing us from doing so (from ENFORCEMENT). And so, if *The Risk Thesis* is correct, mandatory lockdown measures are *pro tanto* justified.

The Risk Thesis, however, faces a powerful objection: because nearly all actions involve imposing at least *some* risk on somebody, there's little one can do to avoid infringing someone's rights. The best one can do, it seems, is as little as possible. The only morally defensible action – whether during a pandemic or not – is to sit silently and motionlessly at home. But that, of course, is absurd. This is the *Paralysis Problem*.[7] Let's investigate it further.

	THE PARALYSIS ARGUMENT
P1	If *The Risk Thesis* is correct, then each person has a right against you imposing a risk of harm on them.
P2	Nearly all of your actions (with, perhaps, the exception of sitting silently and motionlessly at home) impose a risk of harm on someone.
P3	If φing imposes a risk of harm on someone, and they have a right against you doing so, it's morally impermissible for you to φ.
C	If *The Risk Thesis* is correct, there's nearly nothing it's morally permissible to do (with, perhaps, the exception of sitting silently and motionlessly at home).

But surely it's okay for some people to leave the house sometimes. And so, if the argument is sound, *The Risk Thesis* must go. Is the argument sound? The first premise is merely a restatement of *The Risk Thesis* itself, so there's no use denying that. The second premise, although partly an

empirical matter, is plausible enough. A risk, no matter how small, is still a risk, and there are hardly any actions certain to carry no risks whatsoever.

That leaves the third premise. If the right against risk impositions were *inalienable* and *absolute*, the third premise would be plausible: imposing a risk on someone who had an inalienable and absolute right against you doing so would be morally impermissible. But defenders of *The Risk Thesis* needn't think the right is inalienable nor absolute.[8]

Plausibly, the right against risk impositions can be *waived* (if, for example, its bearer *consents* to the risk), which potentially renders the imposition morally permissible (see Section 12.3 for more discussion). And, plausibly, the right can be *overridden* – if, for example, the amount of good that would result from the infringement outweighs its badness – thus rendering its infringement morally permissible. And, the thought goes, the smaller the risk, the easier it is to be outweighed. And so, perhaps, paralysis can be avoided after all.

But *The Risk Thesis* is not yet in the clear. Even if my *pro tanto* duty not to impose a risk on you can be outweighed, rendering my infringement of your right morally permissible, I've infringed your rights all the same. And, according to COMPENSATION, if I've infringed one of your rights, I owe you compensation. But – when the risk is vanishingly small and no actual harm eventuates – it's not plausible that I owe you compensation.[9] Furthermore, given that nearly all actions impose some risk on others, if *The Risk Thesis* were correct, we would need to engage in a quixotic attempt to identify and indemnify a potentially enormous number of people in order to leave the house. And that's its own kind of paralysis.[10]

Finally, according to ENFORCEMENT, if *The Risk Thesis* is correct, the state would be *pro tanto* justified in confining each of us to our respective rooms in order to prevent us from imposing a risk (however small) on anyone else. If we take the view seriously, mandatory lockdown measures (in their most extreme form) would be justified, not only in the face of a deadly pandemic but always. And that's too extreme. Some risks are just too small to justify that much.

12.2.2 The high-risk thesis

The shortcomings of the previous view naturally suggest an alternative (defended by Song, 2019):

> *The High-Risk Thesis*: We each have the right that others not impose a *suitably high* risk of harm on us.

According to *The High-Risk Thesis*, some risk impositions violate rights while others don't. Because many everyday activities during a pandemic

impose higher risks on others than they would normally, on this view, the pandemic can turn previously permissible activities into ones that infringe rights. And, the thought goes, this affords us the resources to justify mandatory lockdown measures but avoid paralysis.

The High-Risk Thesis, however, faces several well-known problems. First, there's the question of how high is "suitably" high? It's unclear what motivates drawing the line one place rather than another – but, admittedly, we should expect some vagueness here. Even still, it's not obvious that there's a consistent way of drawing the line that will account for our intuitions about cases. Furthermore, as Altham (1983, 18–19) argues, a simple threshold will not do – the seriousness of a risk imposition is, plausibly, a function of *both* the size of the risk and the seriousness of the harm that might result. A "suitably high" risk of death is surely lower than a "suitably high" risk of mild discomfort.

Second: building on the previous problem, given that many of our actions have various potential effects, there's often no such thing as *the* harm that might result, and thus no such thing as *the* size of the risk of harm that an act imposes (Thomson 1990, 245). For example, being exposed to COVID-19 might result in a case of very mild (but unpleasant) symptoms, or of Long COVID, or of sufficient seriousness to warrant hospitalization, or of death. Suppose that, for each of these particular harms, the probability of *it* resulting is fairly low – none of them, let's say, count as "suitably high." But, suppose further, that the probability of suffering some harm or other *is* suitably high. Whether "the" risk of being exposed to COVID-19 is "suitably high" depends on how we individuate potential harms. But it's not obvious why that should matter.

Third, the view issues implausible verdicts when one can either impose a high risk on a few or spread the risk across very many more. Here's an example. It's March 2020, and you've contracted COVID-19. You have two options. You can isolate in your apartment, exposing your roommates to a reasonably high risk of contracting it also. Or you can take a series of short bus trips to a remote cabin, which involves incidental encounters with many more people. Because these encounters are so brief, you impose only a small risk on each. Let's suppose (given social distancing and mask wearing) that the risk you pose is so small that, if *The High-Risk Thesis* is correct, each of your fellow travelers has no right against you imposing it. Because the risk you pose to your roommates is (let us suppose) high enough that they do have a right against you imposing it on them, *The High-Risk Thesis* seems to imply, implausibly, that you should take the bus trip – even if you know you will encounter so many people in your travels that the overall chance of grievously harming *someone* is greater than it would be if you stayed home.[11]

In response, one could argue that the smaller risks you impose on the many travelers *add up* to something worse than staying home – ultimately rendering the bus trip impermissible. But because, for each

traveler, your trip would impose only a permissible amount of risk, if it's nevertheless impermissible to impose that risk on all of them, we have a failure of agglomeration: you ought to do this, you ought to do that, and you ought to do the other, but you shouldn't do this, that, and the other. In which case it appears there's no way for you to satisfy all that you ought to do. It's better, then, for proponents of *The High-Risk Thesis* to just accept the conclusion: you ought to impose the risk on the very many rather than impose a higher risk on the very few. On this picture, your moral duties are directed toward each *individual*, not to some aggregate or fusion or class. And, although you can be confident that *somebody* – whoever they are – is likely to be harmed by your action, because that harm is merely a *foreseen* rather than *intended* consequence, it's not clear how worrisome this objection is.[12]

Lastly, *The High-Risk Thesis*, by putting all its emphasis on the size of the risk, fails to capture an important class of intuitively impermissible risk impositions. Let 1/n fall well below the "suitably high" threshold – suppose it's, to borrow an example from Thomson (1986a, 177–179), the risk of death that turning on your gas stove would impose on your neighbor. Now imagine that B plays Russian roulette on A using a gun with one bullet and *n* chambers. B imposes the same degree of risk on A that you impose on your neighbor. Intuitively, what B does to A is impermissible, whereas what you do to your neighbor is not impermissible.[13] What's the difference?

Examples like these put pressure on *The High-Risk Thesis*. Because these examples involve the same degree of risk but inspire different moral reactions, any view that attends only to facts about the size of the risk imposition will, at best, only supply part of the story. There are three ways of responding to this pressure. First, one can resist it: continue to believe *The High-Risk Thesis*, and explain the different moral reactions in some other way. In particular, because it's implausible that you infringe on your neighbor's rights every time you light your stove (on pains of the *Paralysis Problem*), one needs to explain why we think B has wronged A despite causing her no actual harm and without infringing her rights. Second, one can reject *The High-Risk Thesis* in favor of some entirely different account of when a risk imposition infringes a right – one according to which the *size* of the risk isn't of central importance – and hope that that account fares better. Finally, one can take *The High-Risk Thesis* to provide merely a sufficient condition on what rights we have against risk impositions; imposing a "suitably high" risk of harm is part of the story, not the whole story. We'll look at each option in turn.

If B doesn't actually harm A and doesn't infringe A's right against the imposition of a "suitably high" risk of harm, why does it seem like B *wrongs* A by playing Russian roulette on her? According to Holm (2016, 921), "the moral wrongness resides in *the reasons* that the agent has for

performing the action," which, in this case, presumably involves the *intention* to impose a risk on A – after all, with Russian roulette, imposing a risk on another is basically the point! On the other hand, the risk you impose on your neighbor is *not* "the point" – the point is to make coffee; the risk is merely a foreseeable side effect. B has perverse motives, while you do not.[14] And according to Holm (2016, 921), "an agent may be blameworthy for performing an otherwise permissible risk-imposing action due to the agent's reasons for acting." So, the difference is B is *blameworthy* for imposing a 1/n risk of death, whereas you are not blameworthy for imposing a 1/n risk of death.

But this explanation is somewhat unsatisfying. While it's true that B's character leaves something to be desired, it also seems like B *wrongs* A. And, presumably, it takes more than a bad motive to wrong someone. But, then, if B hasn't *done* anything wrong, what is B blameworthy for? At the risk of going in circles, perhaps B is blameworthy for their bad motive: their desire for the rush derived from imposing a risk on A. But now consider B*, who is motivationally just like B, but who doesn't act on their desire to play Russian roulette on A. B* fantasizes about doing so but never actually pulls the trigger. Is B* just as blameworthy? Perhaps. Has B* *wronged* A? I think no – or, at least, not to the same extent B has. Something more, or something different, must account for the difference.

12.2.3 The intention thesis

Perhaps, then, we ought to reject *The High-Risk Thesis* entirely, and instead accept something like the following view (discussed in McKerlie 1986, 243–245):

> *The Intention Thesis*: We each have the right that others not perform actions that aim to impose a risk of harm on us.

On this view, whether a risk imposition violates someone's rights depends not on the size of the risk but on whether that risk imposition is an intended result of the imposer's aim. We don't have rights against risk impositions when they are merely foreseeable side effects of the imposer's aim. We do have a right against actions that are aimed at imposing a risk of harm on us. So, on this view, A has a right against B's playing Russian roulette on her, but your neighbor doesn't have a right against you using your stove – even though both activities impose the same sized risk of death.

I'll raise three issues for *The Intention Thesis*. First, the view is arguably too lenient. Risk imposition needn't be intentional to be impermissible. Suppose B loves to fire his pistol in the air, and suppose that, when he does so, he imposes a high risk of death on A. Imposing this risk is not

B's *intention*, it's merely a foreseeable side effect of an activity he loves. There are things B could do to make it less likely that A is hit by a stray bullet (e.g., he could ensure A isn't around when he fires his gun), but he elects not to because he's utterly indifferent to what effects his behavior has on A's well-being. B's not malicious, he's grossly negligent. Or, a germane example: suppose B is contagious with COVID-19 and knowingly exposes A to the virus, not because he *intends* to impose this risk on A but because he simply doesn't care enough to do otherwise. In both examples, B acts impermissibly – he *wrongs* A. The best explanation of why (given that, let us suppose, no actual harm eventuates) is that A has a right against being treated in these ways.

The second issue is more theoretical. A compelling thought about having a right is that it grants you sovereignty over certain choices; rights provide you with a foundation on which to build your autonomy. By infringing one of your rights, I damage your autonomy by interfering in a domain of choices that, ultimately, are yours to make. But this picture is in tension with *The Intention Thesis* because an action can undermine one's autonomy without that being its *intended* purpose. As McKerlie (1986, 244–245) elaborates,

> [*The Intention Thesis*] makes the force of the rights of others depend on the structure of my plans. If I have one thing rather than another as my goal, or if something figures in my project as a by-product of my pursuit of my goal rather than as a means to the goal, then those rights have no stopping power against my action. A right becomes a sort of reflector that only has force when a hostile intention plays on it. I think instead that the rights of others have an independent force that constrains our plans and intentions. The rights start in their lives, not in our own, and they derive their force from their place in those other lives. When I discover that the route to my goal would also intrude in another life in one of the proscribed ways I must change my plans.

Why should *your* rights depend on *my* plans? Instead, shouldn't my plans, in part, depend on your rights?

Finally, let's step back and consider *The Intention Thesis* in the context of the pandemic. In particular, if that view is correct, would this bolster or subvert the case for mandatory lockdown measures? The answer is that it would subvert it. Consider an activity prohibited by the lockdown measures: e.g., going to the bar. On this view, so long as your *aim* is something *other* than to impose a risk of catching COVID-19 onto others, going to the bar doesn't infringe anyone's rights. If your motivation for going to the bar is to have a good time, not to spread COVID-19, then you are in the clear – even if your presence at the bar will foreseeably impose a risk on others. Because you (and, presumably, the other

bar-goers too) don't infringe anyone's rights by going to the bar, this forecloses a promising way in which the state might be justified in enforcing a mandatory lockdown. The lesson is that *The Intention Thesis* is implausibly permissive.

12.2.4 A hybrid view

Both *The High-Risk Thesis* and *The Intention Thesis* are too permissive but in different ways. *The High-Risk Thesis* doesn't secure us a right against being made the victim of involuntary Russian roulette so long as the risk of death is low enough. But it seems like we do have a right against others performing actions intended to impose a risk of harm on us – even if that risk is quite low. On the other hand, *The Intention Thesis* doesn't secure us rights against the risks imposed from gross negligence. Given that they each plug a hole found in the other, why not combine them?

> *The HI-brid Thesis*: We each have the right that others not perform actions that aim to impose a risk (however small) of harm on us and that others (whatever their motives) not impose a suitably high risk of harm on us.

The HI-brid Thesis is a hybrid. It grants A a right against B's playing Russian roulette on her (no matter how many chambers of the gun); it grants A a right against B imposing a high risk on her (no matter his intentions), but it doesn't grant your neighbor a right against you lighting your gas stove to make coffee. But while *The HI-brid Thesis* might get the cases right, if it does, it does so in an unsatisfying way; it's gerrymandered and seems to lack any unifying underlying theoretical motivation. Furthermore, it inherits problems from both of the views it hybridizes: e.g., it's not clear what "suitably high" means or why it should matter, it's unclear why *your* rights are sensitive to the structure of *my* plans, and so on.

These flaws aren't fatal – it's certainly possible that a compelling underlying motivation that both better unifies the thesis' two ideas and answers the objections could be provided. And while doing so is outside the scope of this chapter, I'll point to what to me seems to be a potentially promising candidate: that we each have a right to be treated with *respect*.

> *The Respect Thesis*: We each have the right that others not perform actions that fail to express proper respect for us and our projects.

Actions that *aim* to impose a risk on us (e.g., Russian roulette) are disrespectful – even if the size of that risk is small. When B imposes a risk on A *intentionally*, B treats A like a plaything, and that fails to express proper respect for her. Furthermore, actions that impose higher risks of

harm, arguably, express greater disrespect.[15] Treating someone with respect requires you to take their interests (including their interest in not being harmed) into account when deciding what to do. That an action *might* cause them harm provides you with a reason against doing it, a reason whose strength is proportionate to how likely that harm is. So, the more risk an action imposes on someone, the stronger the reasons against performing it. And while you may have reasons to perform it, if that risk is suitably high, performing it anyway suggests that you've improperly undervalued the interests of the person your action might harm. And that's disrespectful.

More would need to be said to fully defend the view, of course, but let's assume it is true for now. Then, people have a right against being subjected to a suitably high risk of harm. And this might serve to justify mandatory lockdown measures. If certain activities impose, on some particular people, a suitably high risk of harm, then – unless those particular people have *waived* their right – engaging in those activities violates those people's rights, and thus (from ENFORCEMENT) the state can be justified in restricting those activities by means of tools like lockdowns. It's hard to estimate exactly how much risk a person imposes on others by engaging in activities like bar-going, but it's not implausible to think that – at least, during the height of the COVID-19 pandemic – the risks are "suitably high."

But if everyone subjected to this suitably high risk has *waived* their right against having it imposed on them, imposing it on them isn't a violation – and ENFORCEMENT has no force. This brings us to our final subject of examination: whether those who knowingly and voluntarily choose to engage in activities that are known to be risky effectively waive their right against risk impositions.

12.3 Consent and COVID

Suppose that, by going to the bar, B will impose a suitably high risk of harm on A. Why? Suppose, plausibly enough, it's because A will also be at the bar – and by being in close proximity, there's a suitably high chance of B transmitting the virus to A. Assuming that A has a right against others imposing a suitably high risk of harm on her, does B violate this right of A's by going to the bar? Not if A has waived her right. And, one might argue, by wittingly and voluntarily choosing to go to the bar – an activity known to carry certain risks during a pandemic – A effectively does just that. *Mutatis mutandis* for many of the other activities prohibited by mandatory lockdown measures. And so, the argument goes, ENFORCEMENT doesn't provide adequate justification for such measures.

Let's take a closer look at this argument to see if it holds up. The argument appeals to a plausible idea governing the relationship between *consent* and the *waiving* of a right:

> *The Consent Thesis*: If A validly consents to B φing, then A effectively waives her right against B φing (supposing she has one).

If A has a right against B doing something to her, this right can be waived through an act of valid consent. What is it for an act of consent to be valid? This is a notoriously vexed issue, but there are a number of conditions that seem necessary: the person issuing the consent must be sufficiently *competent*; the consent must be given *voluntarily*; and the person issuing the consent must be *adequately informed* about what she is consenting to.[16] Each of these conditions themselves raises a host of complicated issues, but I will trust that they are clear enough for now.

We are concerned with A's right against others imposing a risk of harm on her. Substituting that into *The Consent Thesis* we get:

(1) If A validly consents to B imposing a risk of harm on her, then A effectively waives her right against B imposing a risk of harm on her.

This instance of *The Consent Thesis* raises an interesting question about *what it is* to consent to a risk. In particular, if you consent to a risk, do you then have no grounds for complaint if the risk eventuates? Thomson (1986a, 188–191) explores this issue in some detail.[17] She argues that there's no simple answer by contrasting two different cases. She observes, "[O]ne who loses in a nonfradulent lottery, which he entered without duress, has no ground for complaint when he loses" (190). (Call this *Lottery Ticket*.) But that same person surely has grounds for complaint in the following example (call it *Unpleasant Way*):

> Suppose there are two ways in which I can get home from the station at the end of the day. The first is pleasant, passes through a brightly lit middle-class shopping area, is quite safe, but is long. The second way is unpleasant, passes through an ill-lit area of warehouses, is unsafe, but is short. Nobody has ever been mugged while walking along Pleasant Way; people have from time to time been mugged on Unpleasant Way. Here I am, at the station; I am tired; I think "The hell, I'll chance it, I'll go home via Unpleasant Way." I then promptly get mugged.
>
> (Thomson 1986a, 189–190)

It's not clear what accounts for the difference.[18] This, in my opinion, remains "a nice problem" (190). But it needn't be our problem. Whether or not A would have "grounds for complaint" were she to contract COVID-19 from B at the bar, our problem concerns whether A's actions suffice to waive her right against B imposing that *risk* on her.

(2) By wittingly and voluntarily going to the bar, A validly consents to B imposing a risk of harm on her.

If A wittingly and voluntarily goes to the bar during the pandemic, she thereby accepts the risks that come with it. That she "accepts the risks" is just another way of saying that she consents to them. This wouldn't be true, of course, if A were unaware of the dangers involved or if she was coerced or unduly induced into going. But she's not and she wasn't: she chooses to go to the bar wittingly and voluntarily.

But (2) can be resisted. While it may be that by wittingly and voluntarily going to the bar, A accepts the relevant risks, it doesn't follow that she validly consents to *B in particular* – or anyone else, for that matter – imposing such a risk on her. Consent is inter*personal*: it's something that one person grants to another who they stand in some relation to. One can "accept the risks" of some course of action – like leaving their umbrella at home when it might rain – without there being anyone to whom they give their consent. And, one might think, this can be so even when the risks in question stem not from the weather but from others.

In fact, Thomson (1986a, 190) makes a similar point concerning the difference between *Lottery Ticket* and *Unpleasant Way*: "there is no person or persons such that I consented to his or their imposing a risk of being mugged on me." In buying a lottery ticket, however, there is a person or persons (e.g., whoever is administering the drawing) such that one consents to the risk of losing to. But, upon reflection, it's not clear this difference makes a difference. Consider the following:

> *Poker Table*: You sit down at the poker table at your local casino and settle in for a night of high-stakes play. You've never met anyone else at the table, and players come and go throughout the night. After hours of play, luck is not on your side: you lose everything to the late-arriving man in sunglasses on your left.

You wittingly and voluntarily agree to play poker, accepting the risk that you might lose it all. But just like A's trip to the bar and Thomson's trip down Unpleasant Way, there is no person or persons in particular to whom you've consented to potentially lose it all. When you joined the table, you didn't know who all else might join as the night wore on. But *Poker Table* seems more like *Lottery Ticket*: you've consented and you have no grounds to complain when you lose.

Here's a suggestion. When you join the poker table, in addition to accepting the risk of losing, you also *impose* a similar risk on your fellow players. Similarly, when you buy the lottery ticket, in addition to accepting the risk of losing, you also accept a chance of winning – which imposes a risk on whoever must pay out if you do. Nothing similar is true in *Unpleasant Way* however: by taking that route, you accept the risk of being mugged, but you do not impose a comparable risk on those who impose that risk on you.[19] The difference appears to be one of *reciprocity*: in accepting an imposition of risk from a person or persons, you in turn

impose a risk on them.[20] Why might this matter? Here's a thought: wittingly and voluntarily imposing a (suitably high) risk on someone functions to *personalize* the relationship – so that accepting the risks imposed on you suffices for consent. By imposing a risk on B, A opens herself to consenting to the risk that B imposes on her.[21]

(2′) By wittingly and voluntarily engaging in a reciprocal risk imposition with B, A validly consents to B imposing a risk of harm on her.

In our example, by wittingly and voluntarily going to the bar, A thereby wittingly and voluntarily engages in a reciprocal risk imposition with B (as well as whoever else might be at the bar that day). And so, from (1) and (2′), we can conclude:

> By wittingly and voluntarily going to the bar, A effectively waives her right against B imposing a risk of harm on her.

This, of course, also holds for B, and for whoever else goes to the bar as well. So, each person who goes to the bar effectively waives their right against each of the others imposing a risk on them – so long as their going is done wittingly and voluntarily.

(3) Many of those activities prohibited by mandatory lockdown measures are such that those who would engage in them would do so wittingly and voluntarily.

And so, for a wide class of activities (like going to the bar), so long as the public is made well aware of the risks involved, those who choose to engage in those activities effectively waive their rights against having the risks of COVID-19 imposed on them. And so, because no one who voluntarily engages in those activities is in danger of having their rights violated, the state's justification for instituting mandatory lockdown measures cannot be grounded in preventing rights violations.

12.3.1 Responding to the consent argument

This is a compelling argument. There is something to the thought that, if I am doing basically the same thing to you that you are doing to me, I lose my standing to complain about what you're doing to me. There is, at the least, something *hypocritical* about condemning someone for the very same thing you're doing to them. Of course, it doesn't follow from the fact that no one can non-hypocritically complain about what you've done, that you've done nothing wrong.[22] But even if it did, I think the argument is ultimately unsuccessful anyway – for somewhat prosaic reasons.

Let's grant (1) and (2′) for the sake of argument. I will raise a problem for (3). It's simply not true that those who would engage in the activities prohibited by mandatory lockdown measures would do so knowingly and voluntarily. Perhaps some would, but it's unrealistic to imagine that all would. Here are two examples.

First, consider someone – call her "Missy" – who has formed many radically false beliefs about the pandemic (perhaps owing to the large amount of misinformation especially prevalent at the beginning), which leads her to systematically underappreciate the risks that she imposes on others and that they present to her. If Missy decides to spend the day at the bar, does she engage in that activity *wittingly*? We can assume, of course, that she knows what a bar is, and so there's a sense in which she knows what she is doing by going to the bar. But given her many and varied false beliefs about the dangers of COVID-19, she fails to appreciate that to spend the day at the bar is to do something that exposes her and others to an elevated level of risk. There's a sense, then, in which she doesn't know what she's doing; she doesn't go to the bar wittingly. And so it's not obvious that her right against being subjected to a suitably high risk isn't violated.

Second, consider someone – call her "Essie" – who is an essential worker, working in food service. Essie spends the day at the bar, let's suppose, because she gets paid to work there. And while it's true that Essie isn't forced to work at the bar, she definitely wouldn't spend her time there if not for the promise of remuneration and the threat of unemployment. Does Essie go to the bar voluntarily? In a sense, yes. She appreciates the risks involved, and given the size of her paycheck and the other options available to her, she figures it's worth it. But her presence at the bar seems importantly less voluntarily than that of her clientele. And so it's not obvious that her right against being subjected to a suitably high risk isn't violated.

One final point. The argument, if successful, establishes that no one who voluntarily engages in activities that would be prohibited under the lockdown is in danger of having their rights violated. We've just questioned whether that argument is successful. But there is a further worry. Even if it is, it doesn't follow that the state isn't justified in instituting mandatory lockdown measures on the grounds of preventing the rights violations of *others* – in particular, the rights of health-care workers and all of us who have a stake in a well-functioning health-care system. Here's the thought. Even if you, as well as everyone else at the bar, have elected to waive your right against the risks of COVID-19, you have all collectively made it more likely that some number of you will require additional health care, including hospitalization. And while that might be a risk that you and the other bar-goers are happy to bear, it's not a risk that you bear alone – we all have an interest in a well-functioning health-care system that isn't overrun and depleted by victims of COVID-19. This might provide the state with ample justification to institute restrictions when doing so is necessary to sustain the health-care system.

12.4 Conclusion

This chapter did several things. First, it explored whether we each – in addition to having a right against being harmed – have a right against being subjected to a *risk* of harm. We surveyed several of the extant positions and cautiously sketched a novel one. Second, it considered an argument against mandatory lockdown measures on the grounds that those whose rights the measures are meant to protect have, in virtue of engaging in the activities in question, effectively waived those rights. We then looked at several objections to this argument and closed by noting that even if the argument is successful, pandemic bar-goers shouldn't rush to get on the beers. The integrity of the health-care system, and the rights of its workers and clientele, may yet supply the state with a plausible rights-based ground to justify mandatory lockdowns.

Notes

1 Gewirth (1981, 16) argues that there is at least one absolute right: the "right not to be made the intended victims of a homicidal project." If he's correct, then B will have more than a *pro tanto* duty – hence, the clause "at least" in DUTY.

2 See, for example, Wenar (2013, 209, 214), and "in my view enforceability (…) distinguishes directed duties we call rights from those we do not" (Cruft 2013, 209). Thomson (1986b, 161) also regards something like ENFORCEMENT as "a plausible idea."

3 Consider the discussion in Thomson (1980) of the beleaguered backpacker, who, imperiled by a blizzard, breaks into someone's boarded-up cabin to escape the elements. Because the backpacker would otherwise die, it's morally permissible for them to do so; it would be morally impermissible for us to prevent them from doing so; but they nevertheless infringe the cabin owner's property rights – and, correspondingly, owe them compensation.

4 Among subjective conceptions of risk, we can distinguish belief-relative from evidence-relative conceptions. The former concern the actual beliefs of the agent, the latter concerns the beliefs the agent *should* have given their evidence. Nothing in the ensuing discussion will turn on this distinction. These two notions can, in turn, be distinguished from a fact-relative conception of risk, which understands risk in terms of facts that are independent of agent's beliefs and evidence (e.g., objective chances) – but, for reasons outside the scope of this chapter, I confess to finding this notion obscure.

5 The distinction between pure and impure risk impositions is introduced and discussed in Thomson (1986a).

6 According to some, an imposition of a risk of harm – whether or not it eventuates – is itself a harm (Finkelstein 2003; Oberdiek 2017; Placani 2017). On these views, there is a sense in which pure risk impositions don't exist. For reasons outside the scope of this chapter, I don't find these views particularly plausible (for some indication of why, see Maheshwari 2021; Rowe 2021), and so I shall set them aside.

7 The name for the objection comes from Hayenhjelm and Wolff (2012, e37), who regard it as "the central philosophical problem" concerning the morality of risk-impositions. It's been discussed by many, including Fried (1970, 192–193), Kagan (1989, 87–88), McCarthy (1997), Nozick (1974, 73–78), Thomson

(1986b), and Holm (2016). Thomson (1990, 244–245) raises a similar, but distinct issue (which Holm, 2016, refers to as the "Proliferation Problem").

8 McCarthy (1997, 215) is one such defender. He thinks that an imposition of risk can be rendered morally permissible (despite infringing a right) if the imposition is consented to or "if the good that would come of bringing about [the imposition] would sufficiently outweigh the burden to the bearer of the right." He calls the former point the Consent Idea, and calls the latter the Trade-off Idea. See Holm (2016) for more discussion.

9 Thomson (1986b, 165) makes a similar argument: because I do not owe you compensation for imposing a risk on you, we have "a ground for thinking that I do not infringe a right of yours when I do so."

10 McCarthy (1997) argues that this problem can be overcome because he thinks, in cases like this, the only compensation owed is the promise to make amends in the event that the risk of harm eventuates. Although, of course, we should make amends if the harm eventuates; it's unclear why this is all that should be required if it's true that we have a right against *pure* risk impositions (see Holm 2016, for further discussion).

11 McKerlie (1986) and Railton (1985) are, to my knowledge, the first to raise this issue. McCarthy (1997, 213–214) discusses the problem, arguing that it provides a point in favor of *The Risk Thesis*. Holm (2016, 922–923) argues that the objection is inconclusive because proponents of *The High-Risk Thesis* needn't find this implication of their view objectionable. Song (2019, 776) argues that proponents can appeal to moral considerations other than rights violations (e.g., expected harm) to rescue the verdict that it's worse to impose risk on the many rather than the few.

12 Aboodi et al. (2008, 266–268) make the same point in response to a different but closely related objection to the "threshold version" of a deontological moral theory under uncertainty. Although not presented as such, *The High-Risk Thesis* is essentially just that.

13 This puzzle, as well as the diagnosis below, appears in Thomson (1986b, 167) and (McKerlie 1986, 241). (Nozick (1974, 74, 81–82) raises the puzzle, too, but floats a different solution.) It's raised as an objection to *The High-Risk Thesis*, in particular, in McCarthy (1997, 213), who takes it to be a reason to reject the view; and in Holm (2016, 920–922) and Song (2019, 775–776), who both think the objection can be answered.

14 Appealing to the distinction between *intended* effects and *foreseeable* side-effects to account for the difference is popular (Holm 2016; McCarthy 1997; McKerlie 1986; Song 2019; Thomson 1986b), but not universal. Nozick (1974, 74), in discussing a similar example concerning driving and Russian roulette, offers two other explanations: first, that Russian roulette, unlike driving, has insufficient social value; and second, that it's "not a normal part of almost everyone's life," (Nozick 1974, 82). Nozick's first explanation (about social value) fails when applied to our example: so long as B enjoys the activity (finds it energizing, etc.), it's not obvious that it has less social value than your home-brewed coffee habit. His second explanation fares better: making coffee is a normal part of almost everyone's life. But as Nozick himself points out, this approach places a lot of weight on the scheme used to classify actions.

15 Lazar (2019, 25) develops a similar idea, arguing that "killing someone more riskily shows greater disrespect for him by more grievously undervaluing his standing and interests".

16 See Eyal (2019) for a general survey of issues concerning *informed consent*. See Dougherty (2021) for a discussion of how ignorance can affect the content of what one counts as having consented to.

17 Davis (2022, 475) explores this issue, too – specifically in the context of COVID-19: "your having accepted the risks that I might unintentionally transmit a deadly virus to you means that, if indeed I do unintentionally transmit the virus to you and you later die from it, I do not thereby wrong you. You knew the risks and you accepted them." He ultimately rejects this argument.

18 One possibility is that you "know the probabilities" in *Lottery Ticket* but not in *Unpleasant Way*. But that won't do. Consider, instead of playing the lottery, playing a slot machine. You don't "know the probabilities" of winning, but if you walk away from the machine a loser, you don't have grounds for complaint.

19 Suppose, instead, that you take Unpleasant Way, not because it's a shorter route home, but because you're interested in finding someone to mug. In this version, you *do* impose a risk of mugging on the others – and so the risks involved are reciprocal. But now this case seems more like *Poker Table*. If one of the other muggers gets the drop on you, that's just fair play.

20 Many have highlighted the potential importance of reciprocal risk-impositions. See, for example, Fried (1970); Song (2019). For discussions of it in the context of the COVID-19 pandemic, see Davis (2022); Lang (2020); Lazar and Barry (2020).

21 This – admittedly underdeveloped – suggestion finds some indirect support from thew ways in which reciprocal risk impositions *don't* seem to matter. Imagine, for example, that I (unbeknown to you) subject you to an involuntary game of Russian roulette. I impose a risk on you. But also suppose that I know the gun is just as likely to backfire, killing me, as it is to kill you. From my point of view, the risks of the activity are mutual. But this goes no way at all toward justifying my actions. Imagine instead that, while I'm playing Russian roulette on you, you are (unbeknownst to me) playing Russian roulette on me – perhaps with the aid of stealth drones. The risks are reciprocal (albeit unbeknownst to us), but again this fails to justify our actions. Lastly, imagine that I know the person I'm imposing a risk on is, in turn, imposing a risk on me – and that I know that they know this too. This case is different from the previous two, which suggests that it's not the reciprocity of the risks *itself* that matters, but what we know about the relationship that obtains between us.

22 Davis (2022, 481) (drawing on Cohen 2006) makes the same point: "the fact that we behaved the same way does not make the wrongdoer's actions any less wrong, but it does make us poorly positioned to criticize them." On the other hand, Lang (2020) appears to disagree (at least about this case): "If I am wronging you and am also being wronged by you, and you are wronging me and are also being wronged by me, then it will turn out…that neither of us is wronging the other." I'm inclined to agree with Davis here, but the objection in the text lands either way.

References

Aboodi, Ron, Adi Borer, and David Enoch. 2008. Deontology, individualism, and uncertainty: A reply to Jackson and Smith. *Journal of Philosophy*, 105(5): 259–272.

Altham, J.E.J. 1983. Ethics of risk. *Proceedings of the Aristotelian Society*, 84: 15–29. DOI: 10.1093/aristotelian/84.1.15

Cohen, G.A. 2006. Casting the first stone: Who can, and who can't, condemn the terrorists? *Royal Institute of Philosophy Supplement*, 58: 113–136. DOI: 10.1017/S1358246100009334

Cruft, Rowan. 2013. Why is it disrespectful to violate rights? *Proceedings of the Aristotelian Society*, 113: 201–224. DOI: 10.1111/j.1467-9264.2013.00352.x

Davis, Jeremy. 2022. The ethics of killing in a pandemic: Unintentional virus transmission, reciprocal risk imposition, and standards of blame. *Journal of Applied Philosophy*, 39(3): 471–486. DOI: 10.1111/japp.12574

Dougherty, Tom. 2021. *The Scope of Consent* (Oxford: Oxford University Press).

Eyal, Nir. 2019. Informed Consent. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.). URL: https://plato.stanford.edu/archives/spr2019/entries/informed-consent/

Finkelstein, Claire Oakes. 2003. Is risk a harm? *University of Pennsylvania Law Review*, 151(3): 963–1001

Fried, Charles. 1970. *Anatomy of Value* (Cambridge, MA: Harvard University Press).

Gewirth, Alan. 1981. Are there any absolute rights? *The Philosophical Quarterly*, 31(122): 1–16. DOI: 10.2307/2218674

Hayenhjelm, Madeleine and Jonathan Wolff. 2012. The moral problem of risk impositions: A survey of the literature. *The European Journal of Philosophy*, 20(S1): 26–51. DOI: 10.1111/j.1468-0378.2011.00482.x

Holm, Sune. 2016. A Right against Risk-Imposition and the Problem of Paralysis. *Ethical Theory and Moral Practice*, 19: 917–930. DOI: 10.1007/s10677-016-9697-6

Kagan, Shelly. 1989. *The Limits of Morality* (Oxford: Oxford University Press).

Lang, Gerald. 2020. Costs and risk impositions in a pandemic. *Public Ethics*. URL: https://www.publicethics.org/post/costs-and-risk-imposition-in-a-pandemic

Lazar, Seth. 2019. Risky killing: How risks worsen violations of objective rights. *Journal of Moral Philosophy*, 16(1): 1–26.

Lazar, Seth and Christian Barry. 2020. Justifying lockdown. *Ethics and International Affairs*. URL: https://www.ethicsandinternationalaffairs.org/2020/justifying-lockdown/

Maheshwari, Kritika. 2021. On the harm of imposing risk of harm. *Ethical Theory and Moral Practice*, 24: 965–980. DOI: 10.1007/s10677-021-10227-y

McCarthy, David. 1997. Rights, explanation, and risk. *Ethics*, 107(2): 205–225.

McKerlie, Dennis. 1986. Rights and risk. *Canadian Journal of Philosophy*, 16(2): 239–251.

Nozick, Robert. 1974. *Anarchy, State, and Utopia* (New York, NY: Basic Books).

Oberdiek, J. 2017. *Imposing Risks: A Normative Framework* (Oxford: Oxford University Press).

Placani, A. 2017. When the risk of harm harms. *Law and Philosophy*, 36(1): 77–100.

Railton, Peter. 1985. Locke, stock, and peril: Natural property rights, pollution, and risk. In *To Breathe Freely*, edited by Mary Gibson (Rowman & Littlefield Publishers), pp. 89–123.

Rowe, Thomas. 2021. Can a risk of harm itself be a harm? *Analysis*, 81(4): 694–701. DOI: 10.1093/analys/anab033

Song, Fei. 2019. Rights against high-level risk impositions. *Ethical Theory and Moral Practice*, 22: 763–778. DOI: 10.1007/s10677-019-09994-6

Thomson, Judith Jarvis. 1980. Rights and compensation. *Nous*, 14(1): 3–15.

———. 1986a. Imposing Risks. In *Rights, Restitution, and Risk*, edited by William Parent, chap. 11 (Cambridge, MA: Harvard University Press), pp. 173–191.

———. 1986b. Some questions about government regulation of behavior. In *Rights, Restitution, and Risk*, edited by William Parent (Cambridge, MA: Harvard University Press), pp. 154–172.

———. 1990. *Realm of Rights* (Harvard University Press).

Wenar, Leif. 2013. The nature of claim-rights. *Ethics*, 123: 202–229.