Draft paper, 15 July 2024

Preparing for the next pandemic: A case for precautionary thinking and citizens' assemblies

Jonathan Birch

Centre for Philosophy of Natural and Social Science,
London School of Economics and Political Science,
Houghton Street, London, WC2A 2AE, UK.
j.birch2@lse.ac.uk
http://personal.lse.ac.uk/birchj1

Abstract

When faced with an urgent and credible threat of grave harm, we should take proportionate precautions. But what is it for a precaution to be "proportionate"? I construct a pragmatic analysis of consisting of four tests—permissibility-in-principle, adequacy, reasonable necessity and consistency—that could realistically be applied by a citizens' assembly meeting online or in person. I apply these tests retrospectively to two examples from the COVID-19 pandemic—border closures and school closures—arguing that my account captures the key questions on which it is both feasible and important to integrate expert input with democratic input. I then consider how we might try to manage the risk of future pandemics in a proportionate way.

Note: See Birch, J., *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI* (2024, Oxford University Press), chapters 7-8, for a longer discussion of proportionality.

The COVID-19 pandemic has been a global public health catastrophe. How can humanity manage the risk of pandemics better in the future? Two major challenges are combined in this question: the challenge of *responding* to new pathogens when they emerge, and the challenge of *preventing* them from emerging. My aim is to set out the core elements of an inclusive, precautionary approach to both challenges.

1. Limitations of cost-benefit analysis

In a cost-benefit analysis, the costs and benefits of a policy in different scenarios are weighted by their probabilities and valued in a common currency. This is often regarded as a gold standard for policy evaluation. Responses to COVID-19 have not generally been guided by cost-benefit analysis, leading to debate about whether such analyses should be, or should have been, carried out (Appleby 2020; Wilkes 2020). I am myself sceptical of the idea that cost-benefit analysis is the right tool for pandemic risk management, for three main reasons.

The first is the problem of *excessively wide probability ranges*. In the early stages of the pandemic, scientific advisers to the UK government avoided terms such as "forecast", "prediction", or "central projection". They instead modelled a "reasonable worst-case scenario" without assigning a probability to this or any other scenario (see, for example, MRC-CGIDA 2020). They did not think the data available supported the assignment of probabilities. By contrast, cost-benefit analysis inherently involves assigning probabilities to possible outcomes (Bradley & Bright 2020). If you cannot impose a reasonable degree of precision on the probabilities, you are unlikely to obtain unambiguous recommendations. You can generate artificial precision by encouraging scientists to go *beyond* what their evidence objectively supports and instead report their subjective degrees of belief, but scientists are right to be reluctant to do this. When evidence is scant, the scientists' subjective degrees of belief will be heavily influenced by their prior probabilities, and there is no good reason to allow the priors of scientists (which are likely to reflect their own values and biases) to set policy on momentous issues.

The second is the problem of *varying attitudes towards risk*. Suppose you have constructed a common currency for weighing outcomes, such as the "quality-adjusted life year" or QALY (Whitehead & Ali 2010). People may still be risk-averse or risk-seeking with respect to that currency. For example, a risk-averse person might prefer a guarantee of 4 QALYs to a gamble with a 50% chance of delivering 10 QALYs and a 50% chance of delivering 0 QALYS. Cost-benefit analysis normally involves an assumption of risk-neutrality: it assumes that to be risk-averse or risk-seeking is irrational. But pandemics are a context in which departures from risk neutrality are both widespread and, at face value, reasonable. Suppose you are faced with a choice between 12 weeks of self-isolation or a chance of being exposed to the SARS-CoV-2 virus, and suppose the chance is such that the expected utility for you is the same for both options. Some people may strongly prefer the risk-averse option of the 12 weeks of self-isolation; others may strongly prefer the risk-seeking option. Speaking for myself, I find both attitudes reasonable.

Departures from risk-neutrality can be incorporated into cost-benefit analysis in technically complicated ways, provided data on citizens' attitudes towards risk is available, but at present this is a significant evidence gap (Lakdawalla and Phelps 2020). Yet even if we had copious data on attitudes towards risk, a deeper problem would remain: in a cost-benefit analysis there is no way to avoid introducing controversial *normative* judgements about which attitudes towards risk are *worthy of consideration* and which are not. To dismiss all non-neutral attitudes towards risk as unworthy of consideration is clearly a normative judgement. However, one does not avoid normative judgements simply by allowing departures from risk-neutrality, for one immediately runs into the question of whether all departures are equally worthy of consideration. If some people are risk-averse and others risk-seeking, should we give departures from risk-neutrality in both directions the same weight, or should we give greater weight to the risk-averse? Should we incorporate departures *within a reasonable range*, while discounting extreme risk-aversion and extreme risk-seeking, or should we regard everyone's attitude towards risk as equally reasonable? I suggest that, in a democratic society, we should look for ways of settling a fundamental normative issue like this democratically, rather than leaving it in the hands of experts.

The third problem is that of *value conflict and incommensurability*. We assumed above that a common evaluative currency, such as the QALY, could be constructed, but pandemics cause such a diverse array of harms that any such currency will be placed under severe strain. For example, I am sceptical of the idea that there is a normatively justified common currency in which one can weigh the disvalue of educational losses to children from sustained school closures against the value of human lives that might be saved by closing schools. Reduce everything to maximizing expected QALYs and you miss morally important considerations, such as the injustice arguably involved in forcing children to make large and long-term sacrifices for the sake of older generations. A similar problem arises if we try to compare the cultural loss of the devastation to the arts sector wrought by lockdowns against lives saved. As with attitudes towards risk, these value conflicts should not be hidden or assumed away. In a democratic society, they should ideally be resolved through open, inclusive, democratic procedures, as Nordheim et al. (2021) have argued.

2. A precautionary framework for managing major risks

These problems point to the need for an alternative framework for managing major risks. An attractive initial thought is that perhaps the "precautionary principle" can provide the framework we need. It is tempting to think that we should, in some sense, "err on the side of caution" in our responses to pandemics, just as we often do in environmental policy, where the precautionary principle has been discussed a great deal.[1] However, a lot of actions may count as "erring on the side of caution", and the precautionary principle has been a great source of unclarity and confusion for policy-makers. This has led to a huge body of case law, jurisprudence and policy literature concerning how to interpret the precautionary principle in

[1] For entry points to the existing literature on the interpretation of the "precautionary principle", mainly in the context of environmental policy and technology policy, see O'Riordan & Cameron 1994; O'Riordan et al. 2001; John 2010, 2019; Munthe 2011; Steel 2015; Stirling 2016; Science for Environmental Policy 2017.

specific contexts. The context of pandemics has been somewhat neglected, at least relative to its new importance in light of COVID-19.

My aim here is to put forward one way of applying the precautionary principle to the context of setting public policy in response to major risks, such as pandemics. Many possible frameworks could be described as "precautionary" in various senses of that word, so I do not claim that what I am proposing is the *only* way to apply precautionary thinking to major risk management. But I do think it as an attractive way, in so far as it captures what I take to be the fundamental commitments of a precautionary attitude to risk.

My proposal consists of a *high-level maxim* that does not make any specific policy recommendations and a *procedural framework* for implementing that maxim in a democratic way to generate specific recommendations. Here is my proposed high-level maxim:

> When there is an **urgent** and **credible** threat of **grave harm**, **proportionate** precautions should be taken.

This aims to improve on previous general maxims, such as that in the 1998 Wingspread Statement (Montague 1998). The Wingspread slogan was that "when an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically". My proposal aims to captures the same general idea without unnecessarily limiting its scope. There is no reason to limit the scope of a precautionary framework to threats caused by "activities". A virus, for example, is not an activity, even though some human activities do increase the risk of new pathogens emerging. Nor should we limit its scope to threats to "human health or the environment" unless these are understood very broadly since, for example, threats to human well-being and threats to democracy may also warrant precautions. And scientific uncertainty need not be limited to "cause-effect relationships". Consider, for example, the R_0 parameter for a new respiratory virus: its reproductive rate in the absence of any attempts at mitigation or suppression. This is a dispositional property of a virus, not a "cause-effect relationship".

The maxim does not imply any concrete policy recommendations. No one could mistake this maxim for an algorithm or decision rule that could be applied in a technocratic manner. Clearly, it does not specify any particular policy response to any particular threat. What it aims to do is crystallize what I take to be the four fundamental commitments of a precautionary approach to risk management.

The first commitment is that some harms are sufficiently *grave* that the state is obligated to try to guard against them: to make no attempt to do so is an unacceptable violation of the duty of protection owed by the state to its citizens. Governments often draw a distinction between military risks and civil risks (such as risks to public health), and have separate frameworks for the two types of risk, the former involving much higher levels of secrecy. My focus here will on civil risks, but we need not write this distinction into the general maxim.

The second commitment is that some threats of grave harm are *urgent*, in the sense that the state is obligated to act immediately to protect its citizens against them rather than waiting for the threat to develop further. It is often very tempting, when faced with a major risk, to take no action and wait for more evidence to emerge. A precautionary approach is committed to the idea that, for a special class of urgent risks, a watching-and-waiting approach cannot be justified.

The third commitment is that, in order to justify precautions, the evidence for an urgent threat of grave harm need not be conclusive but must still be *credible*. There must be enough evidence to support taking the threat seriously. This deliberately sets an evidential bar that is lower than certainty, high probability or high confidence, while at the same time insisting that substantial evidence is needed to warrant the special treatment reserved for exceptionally urgent, exceptionally grave risks.

The fourth commitment is that the precautions taken must be *proportionate*. In very broad terms, to be elaborated upon later, this means that they should do just enough, without being excessive.

These four commitments—*gravity*, *urgency*, *credibility*, *proportionality*—have three important things in common. First, they describe imprecise, qualitative thresholds. There is no precise numerical threshold that separates a grave harm from a tolerable one, an urgent situation from a non-urgent one, a credible threat from a non-credible one, or a proportionate response from a disproportionate one. However, these are still real and important distinctions. Like the distinction between "bald" and "not bald", they are real but imprecise. Imprecision does not imply arbitrariness. We can construct procedures to rule on these questions in a principled, non-arbitrary way.

Second, the four commitments highlight questions that call for *mixed judgements*: judgements that combine both epistemic (i.e. relating to evidence and knowledge) and evaluative (i.e. ethical, social, political) considerations, with no clean way to separate out the two types of consideration (Plutynski 2017; Alexandrova 2018). The urgency of a risk may sound like a purely epistemic matter, but it depends on whether the risk of waiting for more evidence is *unacceptable* given what is at stake, and this is an evaluative question. In a similar vein, it may seem like a purely epistemic matter whether a risk is credible or not, but I take the bar for credibility to be properly sensitive to the stakes. When a risk threatens the whole of humanity, it is appropriate to take it seriously even if the evidence is quite thin, because what is at stake has so much value (a point that is often made in the context of "AI safety", e.g. by Ord 2021). The judgement about where to set the evidential bar in relation to the stakes is an evaluative one. The judgement about whether the evidence reaches that bar is epistemic. The overall assessment of credibility is a mixed epistemic-evaluative judgement.

Third, and most importantly, the four commitments point to questions regarding which *it is both feasible and important to integrate expert input with democratic input*. Due to their partly scientific character, these are judgments that should not be entirely made without

expert input. But due to their partly evaluative character, they are judgements that should not, in a democratic society, be entirely left entirely in the hands of experts. They point towards the need for a procedural framework that provides a mechanism for integrating expert and democratic input on the questions of urgency, gravity, credibility and proportionality.

What sort of procedural framework is needed? I am in favour of involving citizens' assemblies—assemblies of 100-1000 randomly selected citizens—in these decisions, as suggested briefly by Norheim et al. (2021). The technically complicated, scientifically-involved nature of a pandemic response is not a good reason to discount the value of public input. Kitcher (2002, 2011) has long argued for a role for citizens' assemblies in science policy, particularly in the context of priority setting and funding. We just need to distinguish the technical issues on which ordinary citizens are not well placed to contribute from the normative/evaluative questions on which they *are* well placed to contribute.

It is true that the challenge of responding to a pandemic calls not only for scientifically-informed policy but also for *rapid* policy decisions. Citizens' assemblies are usually called upon to discuss relatively non-urgent questions, such as new constitutions (Landemore 2020). They have also been used to consider possible responses to climate change (Climate Assembly UK 2020), which is an undoubtedly urgent problem, but the timescale is nowhere near as compressed as that of the response to a pandemic, where a delay of even a single day can have terrible consequences. Given this, perhaps we should not be surprised that such assemblies have not been a major feature of policy responses to COVID-19, although the French government established a citizens' panel to advise on vaccine confidence (see Casassus 2021). Yet I think the general idea of a citizens' assembly is one that *can* be tailored to urgent questions demanding rapid action. We just need to ensure that the questions posed to the assembly are clear, few in number, specific to the threat at hand, and answerable without high-level expertise—and we need to set the assembly a strict timetable for deliberating and voting on its answers.

In outline, a pool of around 6,000 randomly selected citizens (constructed through stratified sampling to resemble society as a whole) could be created, with frequent rotation and a term limit of about 1 year. This pool would be on standby in case of a major emergency. If a threat were detected by surveillance processes, such as surveillance for new respiratory viruses, an assembly of 150-300 would be drawn from this pool. They would meet online at short notice and initially be tasked with evaluating the urgency, gravity and credibility of the threat, with the ability to question scientific advisers and with the support of a secretariat. They would hear evidence as a single group, break into smaller grounds for deliberations, and then reassemble to vote. The question posed to them would simply be: *Is there an urgent and credible threat of grave harm?* If the assembly decides that there is, the assembly would shift to evaluating the proportionality of shortlisted policy responses. Alternatively, this second task could be taken on by a second assembly drawn from the same pool. As the crisis unfolds, they would meet regularly (e.g. weekly) to evaluate the proportionality of what the government is doing at any given time and to make recommendations.

I imagine this proposal will lead to a variety of reasonable concerns: concerns about whether this process could really be carried out fast enough to allow a timely response; concerns about the democratic credentials of such an assembly, since the members are unelected (though see Landemore 2020 on this); concerns about the assembly's relationship to the executive and legislature; concerns arising from the possibility that the assembly could be badly misaligned with, and/or completely unresponsive to, public opinion (cf. Pettit 2010; Lafont 2019); concerns about accountability for the assembly's decisions, particularly if they turn out to lead to disaster; concerns about the independence of the assembly from its scientific advisers and secretariat; and concerns about the competence of randomly selected citizens with no scientific training to make the mixed judgements required of them (though see Fishkin 2018).

These concerns are too big to address fully in one article. What I want to do here is zoom in on the specific issue of *proportionality*. What is it for a response to be proportionate to the threat, and how is this to be evaluated? What specific questions would be posed to the assembly? By giving a pragmatic analysis of proportionality, I hope to partially defuse some of the above concerns by showing that, on this specific issue, it is desirable and feasible to integrate democratic input with expert input.

3. Four pillars of proportionality

I want to propose four tests for proportionality: permissibility-in-principle, adequacy, reasonable necessity, and consistency. A proposed response to a threat that passes these 'PARC' tests can be considered proportionate.

I do not intend this as a *conceptual analysis* of 'proportionality'. I do not think ordinary language terms are in general amenable to conceptual analyses that give precise necessary and sufficient conditions. What I intend can be better described as a *pragmatic analysis*: a proposal about the questions we should be discussing, and the sequence in which we should discuss them, if we want a deliberative process to arrive at a judgement about proportionality that can command confidence.

Test 1: Permissibility-in-Principle

This is a test any proposed response to a risk must meet to deserve further consideration. It must be that, given our shared values (as represented in the panel), the response could at least in principle form part of an ethically permissible response in the right circumstances, if other conditions are met.

There is a legal component to this test: proposed policy responses to identified risks should be compatible with international human rights law. Legal experts need to be involved in the process to rule on this question, and save the panel's time by excluding options that would be straightforwardly incompatible with human rights. That does not mean, however, that proportionate measures cannot involve any element of suspending rights. Many human rights are considered by the law to be 'qualified', in the sense that it is legal for a public authority to restrict them as a proportionate means to a legitimate aim, such as a public health aim. The

COVID-19 pandemic gave us all personal experience of this, since many of our rights were temporarily suspended by lockdowns. These lockdowns were lawful suspensions of rights. If a proposal involves a suspension of qualified rights in a way that may or may not be proportionate, depending on other factors, then it should remain on the table for further discussion.

That said, international law recognizes a small number of absolute (or 'non-derogable') rights that, by international agreement, cannot be appropriately restricted in any circumstance, no matter how exceptional. The International Covenant on Civil and Political Rights recognizes several such rights: the 'right to life'; the right to freedom from torture and from cruel, inhuman, or degrading treatment or punishment; the right to freedom from medical or scientific experimentation without consent; the right to freedom from slavery and servitude; the right to freedom from imprisonment for inability to fulfil a contractual obligation; the right to freedom from the retrospective operation of criminal laws; the right to recognition as a person before the law; and the right to freedom of thought, conscience, and religion. I propose that any precaution violating any of these rights is ipso facto disproportionate and should be taken off the table.

The right to life is a source of complications, since national laws typically make room for exceptions, despite the supposedly 'non-derogable' nature of the right. The intentional killing of a person by the police or armed forces in the course of their duties is not always deemed a violation of the right to life. Moreover, public health emergencies, such as the COVID-19 pandemic, typically put decision-makers in a situation where, whatever they decide, some people will foreseeably die. There may be intentional allowing of a certain level of death, albeit without intentional killing. Relevantly for our purposes, the fact that restricting scientific research in certain ways may foreseeably involve foregoing medical breakthroughs does not mean that restrictions violate anyone's right to life.

I have been describing the legal element of the 'permissibility-in-principle' test, but I am not proposing that it is solely a legal test. There could be lawful options that nonetheless represent such a severe departure from our shared moral values that a citizens' panel deems them impermissible-in-principle on moral grounds. Suppose, for example, that we convene a panel to consider possible responses to a pandemic, and one of the options floated is a coercive programme of mandatory vaccination. The experts report that this is on public health grounds when proportionate. A panel might nonetheless find such an option deeply unethical, and might decide that it would be better to rule it out immediately as impermissible-in-principle rather than spend time comparing it with other options.

Test 2: Adequacy

Informally, a proportionate response to a risk must do enough. The requirement is not that the response completely removes the identified risk, or even that it renders it negligible or no longer urgent. This is often unachievable, as the COVID-19 pandemic clearly demonstrated. Many measures were taken, but none removed the risk, rendered it negligible, or removed its urgency.

We cannot specify a one-size-fits-all threshold that must be achieved (e.g. the probability of the threat materializing must be reduced to 10 per cent or less). Such a threshold would be arbitrary and insufficiently insensitive to the details of particular cases. Furthermore, we often find ourselves in situations where agreed, precise probabilities cannot be placed on scenarios or outcomes. So what is required of an adequate response? A proposal: the response should either (i) reduce the risk to an acceptable level or, if this is unachievable, (ii) deliver the best level of risk reduction that can be achieved by any permissible-in-principle option.

The question of adequacy may initially sound like a scientific question—one that can be left to the experts, without input from citizens—but it is not. This is because the judgement of an acceptable level of risk is an evaluative judgement. Our way of life depends on tolerating risks: to drive a car is to tolerate risk of injury, to interact with other people in the same room is to tolerate risk of communicable disease. But the question of what levels of risk are acceptable is a question for the people who will be taking the risks, not for an elite group of experts.

What, then, *is* the proper role for experts in debates about adequacy? Their role is that of providing estimates of the level of risk reduction different options provide. There is often enormous uncertainty about the level of risk reduction an option will deliver. It is therefore crucial that experts communicate their uncertainty about risk. Two tools are particularly useful for that task: confidence levels and probability yardsticks.

The Intergovernmental Panel on Climate Change (IPCC) has developed a 'confidence levels' framework for communicating uncertainty about risks. IPCC experts express confidence levels (very low, low, medium, high, or very high confidence) regarding claims about risk, with their confidence depending on the volume and quality of scientific evidence in support of the claim and the amount of consensus surrounding it. I suggest that this approach be generally adopted for the communication of uncertainty to citizens' panels. For example, experts should say not 'Option 1 will reduce the risk more than Option 2' but rather 'I/we have high confidence that Option 1 will reduce the risk more than Option 2'.

A probability yardstick is a standardized protocol for assigning verbal, qualitative labels to probability ranges. An influential example is the PHIA (Professional Head of Intelligence Assessment) probability yardstick, widely used in UK government circles since 2018 (Fig. 1). This yardstick maps the terms 'remote chance', 'highly unlikely', 'unlikely', 'realistic possibility', 'likely/probably', 'highly likely' and 'almost certain' to ranges of probabilities.

The PHIA yardstick is far from perfect. To require 40 per cent probability before being willing to describe an outcome as a 'realistic possibility' seems unwarranted. And to call something a 'remote chance' when it has a probability of 5 per cent seems problematic too (1 in 20 is not that remote). I also fear that the word 'likely' covers too big a range, including outcomes that are slightly more likely than not (~55 per cent), outcomes that are moderately likely (~60–70 per cent), and outcomes that have a ~75 per cent probability of occurring. Yet

this starting point illustrates the general idea. My proposal is not that we adopt this particular yardstick but rather that, in any citizens' panel, a yardstick of some kind should be used to standardize the use of probabilistic language. Experts should say, for example: 'I/we have high confidence that Option 1 will reduce the probability of causing harm, and that, although there will still be a remote chance of causing harm, this will be much reduced from the present situation, where harm is likely.'

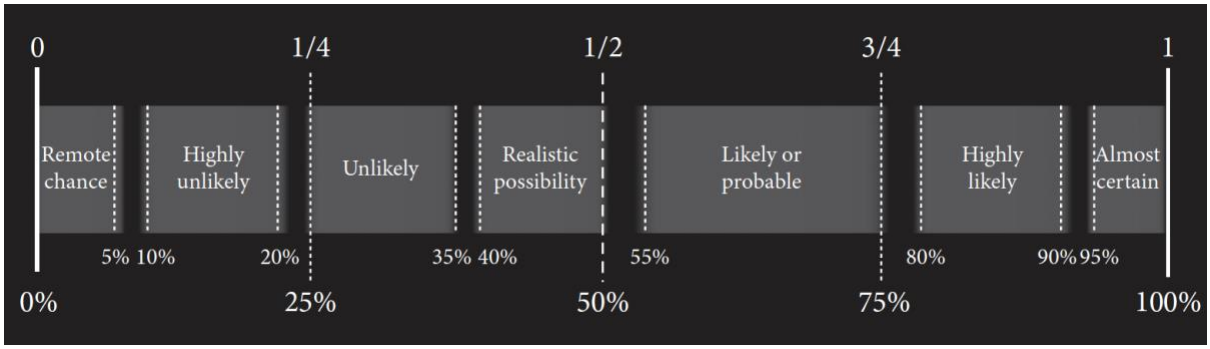

Figure 1: *The PHIA probability yardstick. Versions of this yardstick can be found frequently in UK government documents. Public sector information licensed under the Open Government Licence v3.0.*

Test 3: Reasonable Necessity

Policy responses should not impose harms or costs that go beyond what is reasonably necessary to achieve adequacy. The issue arises because there will often be a temptation to tack extra measures on to a package that is already adequate.

For example, in 2018 the EU banned three specific neonicotinoid pesticides that were credibly linked to colony collapse in bees. There was a case for going beyond a ban on these specific pesticides by banning all neonicotinoids, or even all synthetic pesticides. But tacking on these additional measures was thought to be excessive in relation to the specific threat that was identified.

Let us not get into in the details of whether this was the right decision. Perhaps the limited measures taken did not really do enough to bring the risk down to an acceptable level. Perhaps policy-makers should have gone further. If a citizens' panel had been consulted, ordinary citizens might have recommended stronger measures than those actually taken. My aim is not to take sides on this but to illustrate the general kind of debate that should be had regarding excessiveness. I am proposing a procedure, not a decision rule, so the results in particular cases cannot be pre-judged.

If any cost/risk/harm-benefit analyses of the different options have been conducted, the results, presented accessibly, can be used as expert inputs to discussions of adequacy and reasonable necessity. These methods of analysis do have a place in a good procedure. What I oppose is the use of these methods to resolve value conflict. We should never allow the

value-judgements and subjective probabilities implicit in these analyses to dictate the policy response, without proper deliberation and scrutiny.

Test 4: Consistency

Steel (2015) has rightly emphasized the importance of consistency to judgements of proportionality. Critics of precautionary thinking have argued that it often leads to inconsistency—but, to the extent this is true, it reflects poor institutional design rather than a deep flaw with the idea of precautionary thinking. Good institutional design should make sure that the consistency of the measures on the table is explicitly and carefully considered. The EU takes 'consistency' to mean 'consistency with precedent': a precaution should be consistent with other policies adopted in the past. But I disagree that this is a fundamental condition, especially in situations where the evidential picture is rapidly evolving. It introduces a conservative bias into decision-making. Sometimes the right response is to break with precedent, justifying the departure by pointing to new evidence, or to reasons that were missed by past processes but considered by this one. The consistency requirement should be that breaks with precedent can be justified, not that they cannot happen. If a deliberative process arrives at a view that breaks with precedent, they should be informed of the break, and asked to discuss whether there are good reasons for it. It is not just consistency with the past that matters.

There is also consistency among the measures being proposed right now, in the present. The COVID-19 pandemic has given us examples of packages of measures, assembled in haste, that were not particularly consistent with each other. Infamously, in summer 2020 the UK government introduced an 'Eat Out to Help Out' scheme encouraging people to return to indoor restaurants, while at the same time preparing for a second wave of infections. In September, with infections soaring as predicted, the UK government's scientific advisers gently stressed that 'a consistent package of measures should be adopted that does not appear to promote contradictory goals' (SAGE 2020b).

Most subtly of all, there is a sense of consistency that concerns *the future*: a precautionary response to one threat should not create a new threat that may be serious enough to warrant trying to undo or cancel out the original precaution. This is the type of consistency highlighted by Steel. Again, the COVID-19 pandemic provided some striking examples. The UK's advisers were acutely aware that a strategy aimed at completely eliminating a respiratory virus creates a risk of a new threat: a huge wave as soon as the measures are lifted, or as soon as a new, uncontainable variant emerges elsewhere in the world. The risk is real, but the UK's advisers initially took this to be a reason against even attempting aggressive suppression of transmission, and this was almost certainly an overreaction. They changed their advice when the likely consequences of their strategy for hospital admissions were modelled (Birch 2021a).

For a second example, consider again the case of closing schools to prevent viral transmission: a classic case of value conflict. The harms imposed on schoolchildren will naturally be raised at every stage of discussion. But it is worth noting that, even if a panel has

agreed that school closures are permissible-in-principle, adequate, and reasonably necessary, a further discussion is warranted about their consistency. Might they create a risk of harm that is so severe that it rivals the public health risk posed by the virus? If so, might it be a permissible-in-principle, adequate and reasonably necessary response to that risk to keep schools open?

That is the argument Sarah Lewis and colleagues made in a critique of school closures (Lewis et al. 2021). They argued that closing schools put children at grave risk of harm, with the harms including 'learning loss, reduced social interaction, isolation, reduced physical activity, increased mental health problems', 'potential for increased abuse, exploitation, and neglect', and 'reduced future income and life expectancy'. They concluded that 'the precautionary principle would be to keep schools open to prevent catastrophic harms to children'.

If there are terrible risks of harm associated with both doing and not doing some specific action, the question becomes: can we find some other response that is simultaneously proportionate to both risks, and that does not involve neglecting one to control the other? If we can find such a Goldilocks response, this is the path we should take. If not, then we must make a value judgement about which risk is more severe and worthy of priority. In reality, the UK government took both paths at different moments in the pandemic. In November 2020, a lockdown excluding schools was attempted, to minimize harm to children. In January 2021, with the new Alpha variant rampant, schools were again closed: the immediate risk to public health was prioritized.

4. The division of labour implicit in the PARC tests

The foregoing pragmatic analysis of proportionality is, implicitly, a proposal about how to divide the labour between experts and ordinary members of the public. For each test, there are factual sub-questions that call for expert input. Suppose we dedicate one day's deliberation to each test. This is an ideal situation, more likely to be feasible when debating how to prepare for a *future* pandemic than in the midst of one, when the process might need to be compressed into a single day. But let us think about the ideal. It would make sense for each day to start with presentations by experts that address the sub-questions pertinent to the focal test, and for the experts to remain available throughout the day to answer follow-up questions from deliberative subgroups. However, there are also clearly evaluative sub-questions that call for public input. Those would be the main focus of deliberation. Panel members would be encouraged not to defer to the experts on those sub-questions, and experts would be encouraged not to reveal their private views on them. Table 1 summarizes the division of labour, and Figure 2 summarizes the overall procedure.

On those factual sub-questions where expert input is needed, it is reasonable to expect experts to put their answers in plain language that ordinary citizens can understand (e.g. by using probability yardsticks and confidence levels). None of the questions demand inaccessible, technical answers. In answering these questions, some value-judgements will be made, unavoidably. For example, experts will have to make judgements about how to apply

coarse-grained probabilistic language (such as ‘highly likely’) to borderline cases where the yardstick does not determine what to say. But those value-judgements will not dictate the policy response and will not be beyond scrutiny. If a member of the public wants to ask ‘Wait, why did you describe that as highly likely?’ they will have the opportunity to do so. Meanwhile, on all of those evaluative sub-questions where public input is needed, ordinary members of the public are competent to answer the questions being posed. The key questions are:

- Does this option depart too far from our shared moral values to deserve further consideration?
- What level of risk reduction is acceptable, given our shared values and attitudes towards risk, and bearing in mind that our shared way of life requires us to tolerate some risks?
- Where harms and costs are imposed, can they be justified to those affected as reasonably necessary to achieve adequate risk reduction?
- Where inconsistencies arise (e.g. with other risks or with past precedent), should they be resolved by breaking with past practice, or by reconsidering our response to the current risk?

These questions do require members of the public to grasp the concept of risk, at the sort of level that might be required to complete a risk assessment form. They need to be able to understand risk as something that comes in degrees and that may be reduced by degrees. Some minimal training regarding the concept of risk (not in any way specific to the present risk) may be needed for the panel to function effectively. But no detailed scientific or mathematical knowledge is needed to debate the above questions and to apply one’s values to the issues at stake.

Indeed, there is a sense in which people who do not have specialist prior knowledge are better placed to judge whether the harms and costs imposed by a precaution can be justified to those affected, since those affected will tend to lack specialist prior knowledge too. If the justification is incomprehensible to those affected, it is not truly justifying the harms and costs to them; it is merely justifying those harms and costs to an idealized image of them.

My contention is that, if a panel of at least 150 citizens debated proportionality in relation to a specific issue following the structure provided, arriving by consensus or by majority vote at a single policy response, we could have confidence in the results. We could be confident that the major considerations would have received dedicated and extended deliberation, that the people doing the deliberating were sufficiently informed and appropriately representative of the values of the population, and that they were able to apply those values to the questions being posed.

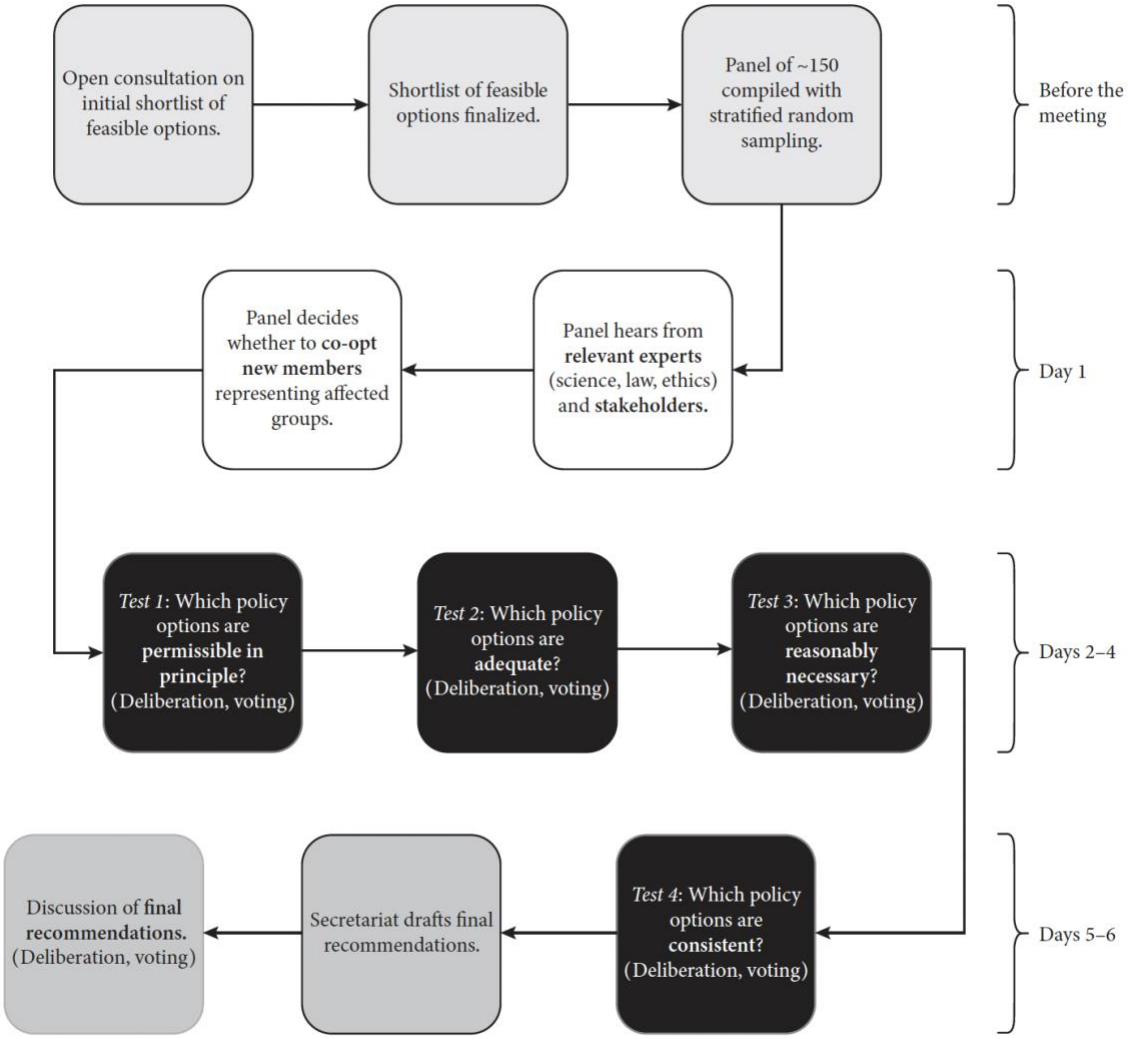

Figure 2: *A procedure for assessing the proportionality of a proposed response to identified risks. The timescale of 6 days is an ideal for normal times that might have to be compressed if implemented during (rather than prior to) a pandemic. From Birch (2024).*

Table 1: *A summary of the PARC tests and the division of labour they induce between experts and members of the public. From Birch (2024).*

Condition	Description	Aspects calling for expert input	Aspects calling for public input
Permissibility in principle	The response is compatible with our shared values. Given those values, the response could in principle form part of a proportionate response, if other conditions are met.	If the response involves suspending rights, which rights will be suspended and in what ways? What other harms and costs are likely to result from the proposed response?	In cases where a response is legally permissible -in-principle, does it nonetheless depart too far from our shared moral values to merit further consideration?
Adequacy	The response either (i) reduces the risk to an acceptable level or, if this is unachievable, (ii) delivers the best level of risk reduction that can be achieved by any permissible-in-principle option.	What levels of risk reduction are likely to be delivered by different options? Confidence levels and probability yardsticks should be used. If cost-benefit or harm-benefit analyses are available, the results can be presented.	What level of risk reduction is *acceptable*, given our shared values and attitudes towards risk, and bearing in mind that our shared way of life requires us to tolerate some risks?
Reasonable necessity	The response does not impose any harms or costs (including suspensions of rights) beyond those reasonably necessary to achieve adequacy, and minimizes any harms or costs that are reasonably necessary.	If there are multiple permissible ways to achieve adequacy, what are their comparative harms and costs? What steps are proposed to minimize these harms and costs? If any cost-benefit or harm-benefit analyses are available, the results can be presented to the panel.	Where harms and costs are imposed, can they be justified to those affected as *reasonably necessary* to achieve adequate risk reduction?
Consistency	The response can be reconciled with our attitude towards other risks, including any new risks created by the very response under consideration. Sometimes reconciliation can take the form of justifying a departure from past practice.	What are other comparable risks, and how have they been managed? Will a new risks be generated by the proposed responses? If so, what are they?	Where inconsistencies arise (e.g. with other risks or with past precedent), should they be resolved by breaking with past practice, or by reconsidering our response to the current risk?

5. Proportionality in the COVID-19 pandemic: Border closures

Would the PARC tests have helped a citizens' panel think through the issues posed by the COVID-19 pandemic? Let us consider two examples: border closures and school closures.

Border closures have been a major element of the global response to COVID-19. 348 countries introduced either partial or complete border closures in the week of 11 March, and 189 countries have, for at least a short period, introduced a near-complete border closure, where all non-citizens are denied entry (Shiraef 2021). Several countries have also compelled some or all travellers to spend two weeks in a quarantine facility, usually a co-opted hotel (Walsh 2021). When a country both completely closes its borders to all non-citizens *and* compels all returning citizens to quarantine in an approved facility, let us call this a "strict border closure".

Australia, New Zealand and Taiwan have been notable for their early and sustained use of strict border closures. All three closed their borders to foreign nationals on 19 or 20 March 2020, bringing in hotel quarantine for all returning citizens. Was this proportionate? One point of contrast is the UK, which initially took a very different approach, before bringing in a limited form of hotel quarantine in 2021 (Walsh 2021). In February 2020, the UK's Scientific Advisory Group for Emergencies (SAGE) had advised that it would be "draconian" to close the UK's borders entirely (SAGE 2020a). "Draconian" does not imply "unjustified", but it is not an evaluatively neutral term either. It suggests a (questionable) judgement of excessiveness in relation to the risk.

Permissible-in-principle? Strict border closures can preserve absolute rights, as long as the quarantine conditions are not cruel, inhuman or degrading. The right to travel is not absolute. Even the right to return to one's own country, though generally recognised as an important right, is not an absolute right under international law. The UN Human Rights Council has commented that "there are few, if any, circumstances in which deprivation of the right to enter one's own country could be reasonable" (UN HRC 1999), leaving open the possibility of pandemics creating such circumstances (Martha & Bailey 2020). So, although the Australian government in particular has been criticised for partially suspending this right in practice by setting a cap on the number of international arrivals at any given time, this is not automatically disproportionate (Evershed 2021). The question then becomes whether it is excessive, or whether it is reasonably necessary to achieve adequate risk reduction.

Adequate? At the time strict border closures were first introduced, there was a clear risk that borders would have to be reopened before any vaccine or effective treatment had become available, and, in such a case, they would merely have postponed a disastrous epidemic rather than preventing one. However, all this shows is that border closures do not completely remove the risk of a disastrous epidemic. This makes it proportionate to combine border closures with other preparatory measures that also satisfy the other conditions (such as building up healthcare system capacity, testing and tracing capacity, and vaccine distribution capacity), so that the country is as well prepared as possible for the time border closures end.

It remains plausible that a package of measures containing border closures will perform at least as well as any other package in reducing that risk.

Reasonably necessary? The crucial question here is whether strict border closures do more than is reasonably necessary to satisfy the other conditions. This seems to have been the view of SAGE in describing this approach as "draconian". However, border closures do appear to have been effective in the early stages of the pandemic as a means of preventing the import of SARS-CoV-2 (Grépin et al. 2021). It is less clear at present that persisting with border closures is effective (e.g. for preventing import of new variants) once there is already sustained community transmission of the virus.

The practice of setting caps on the number of international arrivals, as Australia did, in order to keep arrivals within the capacity of designated quarantine hotels, raises a concern about excessiveness. Could an acceptable level of risk reduction have been achieved while removing the cap and allowing some returning citizens to quarantine at home? This is an issue that only a democratic process could settle in an appropriate way. As with so many of the issues concerning the response to COVID, the answer depends partly on what it is appropriate to assume about human behaviour: to what extent can people be trusted to self-isolate properly at home, without enforcement and monitoring? Those on both sides of the debate about the proportionality of strict border closures should be able to agree that a public-involving process could and should have been involved in that judgement call.

Consistent? Do sustained border closures create a new threat that itself calls for a precautionary response? There is a risk that they will merely delay the epidemic until borders are reopened but, as discussed earlier, this is best seen as a challenge to the adequacy of border closures rather than their consistency, because a delayed epidemic is not a "new threat" in the relevant sense. In some cases, border closures can create a risk of food shortages in countries that are far from self-sufficient regarding food (Charumbira 2021; Thomas 2020), and this should be taken seriously when evaluating consistency.

6. Proportionality in the COVID-19 pandemic: School closures

On 4 March 2020, the Italian government ordered the closure of all schools and universities. By 16 March, most European countries had closed their schools. The UK did so on 20 March. Schools were closed for around three months (see Aspinall 2020 for a timeline). The aim was to suppress the reproductive rate (R) of the virus to below 1, so as to prevent the healthcare system becoming overwhelmed by exponential growth in the number of cases. Was this proportionate? Not all governments agreed that it was. A notable exception was Sweden, which partially closed its schools but kept them open for children up to the age of 15 (Vlachos et al. 2021).

Permissible-in-principle? The right to education is an important human right, but not an absolute one. It is generally recognized that it could be permissibly suspended in exceptional circumstances, such as a public health emergency.

Adequate? No government has suggested that school closures are adequate by themselves, but many have regarded them as a reasonably necessary component of an adequate package of measures for bringing *R* below 1. On 17 March 2020, SAGE's subgroup on epidemiological modelling, SPI-M-O, produced a "consensus statement" that contains a crisp statement of the justification for school closures:

> It is almost certain that school closures will not make the epidemic worse, and that they would reduce both the epidemic peak and expected number of cases. Our best assessment is that they would reduce the reproduction number by between 10% and 20%. We do not know how likely it is that this will change the reproduction number from being above 1 to below 1 (SPI-M-O 2020)

The combination of all *other* measures put together was judged "likely or highly likely" to lead to the critical care capacity of the health service being exceeded. School closures could tip the balance (though no probability is attached to the claim that they will), and they would not make things worse. The claim is implicitly that a package of measures including school closures would deliver a level of risk reduction at least as great as any other package. The advisers were in a position to be confident of this comparative judgement, despite being unable to estimate the absolute probability of the strategy succeeding.

With the virus again resurgent, the UK closed schools for a second sustained period in early 2021. By this time, more evidence both for and against school closures had become available. On 28 January 2021, SAGE cautioned that:

> The opening of primary and secondary schools is likely to increase effective R by a factor of 1.1 to 1.5 (10% to 50%) (medium confidence). Options with fewer children in attendance (such as selected year groups or cohorts) are likely to fall towards the lower end of this range. While prevalence is falling and vaccinations are continuing, a later opening of schools would result in less community transmission and fewer hospitalisations. Attending school is important for the wellbeing of children and SAGE advises that reopening schools should be a priority when infection rates allow it. (SAGE 2021).

Reasonably necessary? This condition, combined with the significant harms imposed by school closures, points to the need to prefer partial school closures over full closures, and to close as little of the school as possible in order to achieve critical epidemiological goals (in this case, $R < 1$). SAGE's advice recognizes this. At the time, however, *R* was estimated by SAGE to be in the range 0.7 and 1.1 with schools currently closed. This implied that, in their opinion, scenarios in which $R > 1$ or ≈ 1 remained credible, suggesting no scope at that time for reopening schools even partially without compromising adequacy.

I find it plausible that a public-involving process (such as a citizens' assembly) would have accepted these rationales for the adequacy and non-excessiveness of a lengthy one-off school closure in 2020 and another in early 2021. I find it less likely that it would have accepted the

idea of indefinite, repeated school closures, as floated in Ferguson et al. (2020). What troubles me is no such democratic exercise took place. Arguments for the proportionality of these measures are implicit in the SAGE minutes and related documents, but the public had no input on this momentous question in the UK or, to my knowledge, any other country. Here too, those on both sides of the debate about the proportionality of school closures should be able to agree that a public-involving process could and should have been involved in the decision.

Consistent? Lewis et al. (2021) object to the proportionality of school closures on the grounds that they create a risk of grave harm to children, including "learning loss, reduced social interaction, isolation, reduced physical activity, increased mental health problems", "potential for increased abuse, exploitation, and neglect", and "reduced future income and life expectancy". They conclude that "the precautionary principle would be to keep schools open to prevent catastrophic harms to children" (Godfrey-Smith 2021 has defended a similar view).

This objection should be taken seriously. The "suppression" strategy laid out on 16 March in an influential paper by the Imperial COVID-19 Response Team (Ferguson et al. 2020) involved *indefinite* school closures for two-thirds of the time until a vaccine was developed. A graph in the paper illustrates the strategy by depicting *seven* periods of prolonged school closure, but more than seven could have been needed, had vaccines taken longer to develop. This strategy, as depicted, was never pursued. If it had been, the objection from consistency would have been very strong indeed. A defender of school closures could reply that the long-term effects on learning, income and life expectancy of a one-off school closure, or two of them, can be effectively prevented by a good enough recovery package. Some harms, though, are immediate, especially mental health problems and the potential for increased abuse, exploitation and neglect, so the objection is not so easily defused.

To be clear, in considering these two examples, my aim has not been to declare some actions proportionate and others disproportionate. To do so would be in tension with my view that proportionality is a matter for public-involving procedures, such as citizens' assemblies, to decide. I have no reason to think that such an assembly would share my own views. My aim has been to show that the PARC tests capture the right questions to ask: the mixed judgements that should properly be at the heart of decision-making when credible, grave and urgent risks are identified.

7. Managing the risk of future pandemics

Once a novel pathogen is spreading freely from one human to another, a proportionate response to the threat is very likely to involve substantial harms, costs, and suspensions of qualified human rights, as COVID-19 has shown. By contrast, proportionate management of the risk of a future pandemic *before* the pathogen emerges may be achievable at far lower cost. I want to close with some reflections on this risk.

Since the early 2000s, countries around the world have identified pandemic influenza as one of the gravest threats they faced. H5N1 has been a source of particular concern. H5N1 is a highly pathogenic subtype of the influenza A virus. When H5N1 has infected a human (and this has happened a few hundred times), it is estimated to have had an infection fatality rate of somewhere around 60% (Akhtar 2021). The world is spared, for now, because H5N1 cannot transmit between humans, just as, prior to late 2019, SARS-CoV-2 could not spread between humans.

We can hope that a virus as pathogenic as H5N1 would be easier to contain than SARS-CoV-2, because there would be fewer asymptomatic cases, but there is no guarantee of this. We know from the 1918-1919 influenza pandemic (driven by another subtype of influenza A, H1N1) that influenza viruses can combine high fatality rates (the case fatality rate of the 1918 influenza virus was over 2.5%, with an unusually young age profile; those aged 25-35 were at very high risk) with very high transmissibility. A widely cited review concluded in 2006 that:

> Even with modern antiviral and antibacterial drugs, vaccines, and prevention knowledge, the return of a pandemic virus equivalent in pathogenicity to the virus of 1918 would likely kill >100 million people worldwide. A pandemic virus with the (alleged) pathogenic potential of some recent H5N1 outbreaks could cause substantially more deaths. (Taubenberger & Morens 2006)

Is this a risk we can mitigate now? It is, because many of the main drivers of the emergence of new zoonotic pathogens are anthropogenic (UNEP 2020; see also **Figure 3**). The risk arises from interaction between humans and animals, and between wild and domesticated animal populations, and we have substantial influence over both types of interaction. The wildlife trade is a crucial risk factor, as is the destruction of natural habitats that brings humans into closer contact with wildlife. Another risk factor is surging global demand for meat, combined with intensive methods of meat production (UNEP 2020). The dense packing of genetically similar animals into small spaces creates conditions that can increase the likelihood of a virus evolving into a more pathogenic form, as appears to have happened with H5N1 (Akhtar 2012). When potential new hosts are everywhere in the immediate environment, a virus is no longer under selection to spare its host's life while it awaits a new transmission opportunity, and the conventional wisdom that viruses evolve to become less pathogenic need not apply (Sober 2020). A third type of risk factor is laboratory research in which the transmissibility of pathogens is deliberately enhanced (Ord 2021).

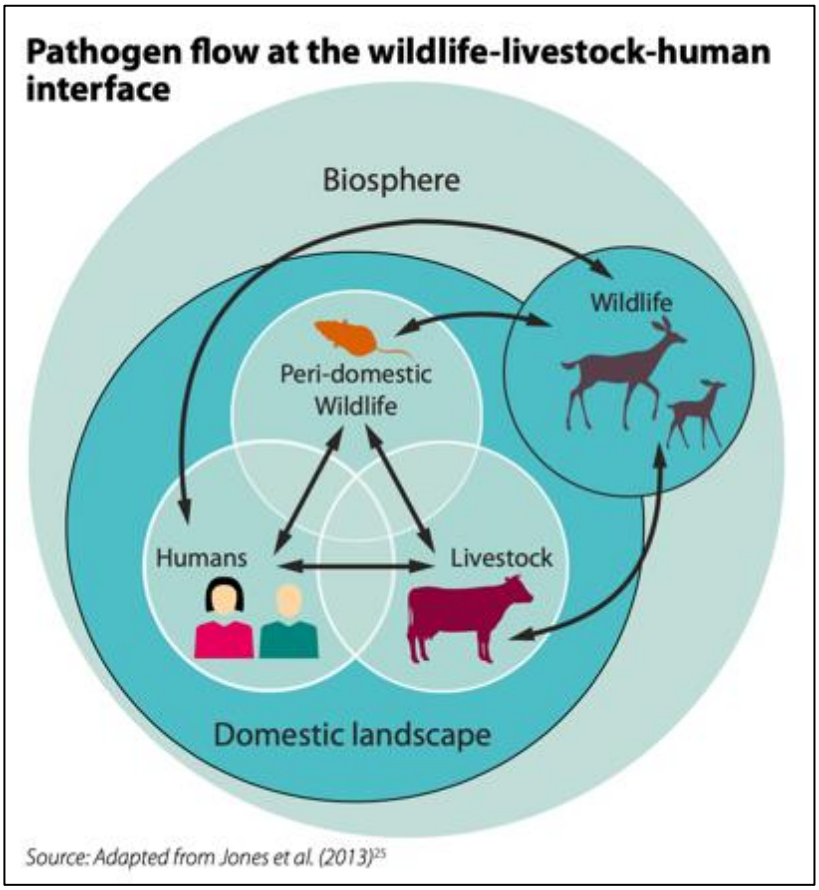

Figure 3: Figure reproduced from UNEP (2020), adapted from Jones et al. (2013). *The risk of zoonotic transmission can be managed by reducing the amount of interaction between each group shown, or by improving the biosecurity of the interactions.*

There is an urgent need for meaningful democratic reflection and decision-making, either in the form of citizens' assemblies or in some other form (such as referendums), on the question of what steps would be proportionate to this risk. We can see that, if the risk is allowed to develop to the stage of human-human transmission, proportionality will justify imposing substantial harms on billions of people. If we can control the risk by imposing proportionate measures now, with much less harm, there is a clear case for doing it.

Given my emphasis throughout this paper on democratic processes, I do not take my own view on these questions to be decisive. Nonetheless, in the interests of provoking discussion, I will state my own view anyway.[2] I think measures taken up to now to mitigate the risk of the emergence of new pathogens have been far from proportionate, and that proportionality in fact justifies much more ambitious action. More ambitious steps might include a total worldwide ban on the wildlife trade and the farming of wild species for any purpose (with reasonable compensation for those who have previously relied on it), very strict biosecurity protocols for all animal farming, removal of government subsidies for meat production and a shift towards subsidising plant-based alternatives, other actions aimed at reducing demand for

[2] See also Birch 2021b; Whitmore 2021; Cao 2021 and Akhtar 2021 for more discussion-starters on these issues.

meat (such as campaigns to make one day of the week a meat-free day), and a moratorium on "gain of function" laboratory research while tighter international biosecurity norms are developed.

There would be significant costs to all of these steps. Compensation will be appropriate for those affected but not culpable for creating the risk. Yet these costs are likely to be far lower than the costs of a future pandemic of H5N1 or another highly pathogenic influenza virus. At present, the world is akin to a drunk driver who, after one bad accident, is back on the road again, still drunk.

Acknowledgements
I thank Liam Kofi Bright for comments on a draft. I thank virtual audiences at the University of Salzburg, "COVID Philosophy Week" (University of Johannesburg), and "Thinking the Pandemic" (Federal University of Minas Gerais) for their reactions, and I thank Charlotte Werndl, Ben Smart, Alex Broadbent, and Felipe Carvalho for creating those opportunities.

References

Akhtar, A. (2012) *Animals and public health: Why treating animals better is critical to human welfare*. Basingstoke: Palgrave Macmillan.

Akhtar, A. (2021, 5 March) COVID-19 is the consequence of our cruelty to animals. *LSE Covid-19 Blog*. https://blogs.lse.ac.uk/covid19/2021/03/05/lse-festival-2021-covid-19-is-the-consequence-of-our-cruelty-to-animals/

Alexandrova, A. (2018). Can the science of well-being be objective? *British Journal for the Philosophy of Science*, 69, 421-445.

Alvelda, P., Ferguson, T., Mallery, J. C. (2020, 18 November). To save the economy, save people first. *Institute for New Economic Thinking.* https://www.ineteconomics.org/perspectives/blog/to-save-the-economy-save-people-first

Appleby, J. (2020). Tackling Covid-19: are the costs worth the benefits? *British Medical Journal* 369, m1496. http://doi.org/10.1136/bmj.m1496

Aspinall, E. (2020). COVID-19 timeline. *British Foreign Policy Group*. https://bfpg.co.uk/2020/04/covid-19-timeline/

Blaine, J. (2021). Exploring the psychosocial consequences of mandatory quarantine during the COVID-19 pandemic in Hong Kong. *Psychology and Behavioral Sciences*, 10, 96-103. http://doi.org/10.11648/j.pbs.20211002.13

Birch, J. (2021a). Science and policy in extremis: the UK's initial response to COVID-19. *European Journal for Philosophy of Science*.

Birch, J. (2021b, 9 March). Humans, animals and pandemics: What needs to change? *LSE Philosophy Blog*. https://www.lse.ac.uk/philosophy/blog/2021/03/09/animals-humans-and-pandemics-what-needs-to-change/

Birch, J. (2024). *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI*. Oxford University Press.

Bradley, R., & Bright, L. K. (2020). Managing our uncertainty amidst the crisis. *The Philosophers' Magazine*, 90, 32-35.

Broadbent, A., & Smart, B. (2020, 23 March). Why a one-size-fits-all approach to the COVID-19 pandemic could have lethal consequences. *The Conversation*. https://theconversation.com/why-a-one-size-fits-all-approach-to-covid-19-could-have-lethal-consequences-134252

Cao, D. (2021, 5 March) To avoid more pandemics, we need to stop eating wild and factory-farmed animals. *LSE Covid-19 Blog*. https://blogs.lse.ac.uk/covid19/2021/03/05/lse-festival-2021-to-avoid-more-pandemics-we-need-to-stop-eating-wild-and-factory-farmed-animals/

Carrington, D. (2013, 29 April) Bee-harming pesticides banned in Europe. *The Guardian*. http://www.theguardian.com/environment/2013/apr/29/bee-harming-pesticides-banned-europe

Casassus, B. (2021, 18 February). Vaccine-wary France turns to citizens' panel to boost trust in COVID-19 shots. *Science*. https://www.sciencemag.org/news/2021/02/vaccine-wary-france-turns-citizens-panel-boost-trust-covid-19-shots

Casey, B. H. (2020, 18 December). Covid-19: Is there a trade-off between economic damage and loss of life? *LSE European Politics and Policy (EUROPP)*. https://blogs.lse.ac.uk/europpblog/2020/12/18/covid-19-is-there-a-trade-off-between-economic-damage-and-loss-of-life/

Charumbira, S. (2021). Landlocked Lesotho faces food crisis amid Covid border closures. *The Guardian*. https://www.theguardian.com/global-development/2021/jan/19/landlocked-lesotho-faces-food-crisis-amid-covid-border-closures

Climate Assembly UK (2020). Report: The path to net zero. https://www.climateassembly.uk/report/

Evershed, N. (2021, 7 July). Data reveals Australia's new international arrivals cap is harshest yet. *The Guardian*. https://www.theguardian.com/news/datablog/2021/jul/07/australias-travel-restrictions-how-the-cap-on-international-arrivals-has-changed

European Commission (2000). *Communication from the Commission on the precautionary principle*. https://op.europa.eu/en/publication-detail/-/publication/21676661-a79f-4153-b984-aeb28f07c80a/language-en

Ferguson, N. M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Dorigatti, I., Fu, H., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Okell, L. C., van Elsland, S., Thompson, H., Verity, R., Volz, E., Wang, H., Wang, Y., Walker, P. G. T., Walters, C., Winskill, P., Whittaker, C., Donnelly, C. A., Riley, S., & Ghani, A. C. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Retrieved from: https://www.gov.uk/government/publications/impact-of-non-pharmaceutical-interventions-npis-to-reduce-covid-19-mortality-and-healthcare-demand-16-march-2020

Ferré, M. E. (2021, 11 July). In response to rising COVID-19 cases, Taiwan continues banning foreign nationals. *The Mobile Workforce*. https://www.mobilework.law/2021/07/in-response-to-rising-covid-19-cases-taiwan-continues-banning-foreign-nationals/

Fishkin, J. S. (2018). *Democracy when the people are thinking: Revitalizing our politics through public deliberation*. Oxford: Oxford University Press.

Godfrey-Smith, P. (2021). Covid heterodoxy in three layers. https://petergodfreysmith.com/wp-content/uploads/2021/08/Covid-Heterodoxy-PGS-v4D.pdf

Grépin, K. A., Ho T.-L., Liu, Z., Marion, S., Piper, J., Worsnop C. Z., Lee, K. (2021). Evidence of the effectiveness of travel related measures during the early phase of the COVID-19 pandemic: A rapid systematic review. *BMJ Global Health*, 6, e004537. http://10.1136/bmjgh-2020-004537

Habershon, S. (2021, 27 March). New Zealand's Covid quarantine fee change puts politics over a citizen's right to return. *The Guardian*. https://www.theguardian.com/world/commentisfree/2021/mar/27/new-zealands-covid-quarantine-fee-change-places-politics-over-a-citizens-right-to-return-sarah-habershon

Jemrosik, E., & Selgelid, M. J. (2020). COVID-19 human challenge studies: Ethical issues. *Lancet Infectious Diseases*, 20, E198-E203. https://doi.org/10.1016/S1473-3099(20)30438-2

John, S. (2010). In defence of bad science and irrational policies: An alternative account of the precautionary principle. *Ethical Theory and Moral Practice*, 13, 3–18.

John, S. (2019). The politics of certainty: The precautionary principle, inductive risk and procedural fairness. *Ethics, Policy & Environment*, 22, 21-33 http://doi.org/10.1080/21550085.2019.1581418

Jones, B.A., Grace, D., Kock, R., Alonso, S., Rushton, J., & Said, M.Y. (2013). Zoonosis emergence linked to agricultural intensification and environmental change. *Proceedings of the National Academy of Sciences of the United States of America*, 110, 8399–8404. https://doi.org/10.1073/pnas.1208059110

Kitcher, P. (2002). *Science, Truth and Democracy*. New York: Oxford University Press.

Kitcher, P. (2011*). Science in a Democratic Society*. Buffalo, NY: Prometheus.

Lafont, C. (2019). *Democracy without shortcuts: A participatory conception of deliberative democracy*. Oxford: Oxford University Press.

Lakdawalla, D. N., & Phelps, C. E. (2020). Health technology assessment with risk aversion in health. *Journal of Health Economics*, 72, 102346. https://doi.org/10.1016/j.jhealeco.2020.102346

Landemore, H. (2020). *Open democracy: Reinventing popular rule for the twenty-first century*. Princeton, NJ: Princeton University Press.

Lee, P-C., Chen, S-C., Chiu, T. Y., Chen, C. M., Chi, C. (2020, 21 July). What we can learn from Taiwan's response to the covid-19 epidemic. *The BMJ Opinion*. https://blogs.bmj.com/bmj/2020/07/21/what-we-can-learn-from-taiwans-response-to-the-covid-19-epidemic/

Lewis, S. J., Munro, A. P. S., Smith, G. D., Pollock, A. M. (2021). Closing schools is not evidence based and harms children. *British Medical Journal*, 372, n521 http://doi.org/10.1136/bmj.n521

Martha, R., & Bailey, S. (2020, 23 June). The right to enter his or her own country. *EJIL:Talk!* https://www.ejiltalk.org/the-right-to-enter-his-or-her-own-country/

Montague, P. (1998, 19 February). The precautionary principle. *Rachel's Environment and Health Weekly*, 586. https://ratical.org/co-globalize/REHW586.html

MRC Centre for Global Infectious Disease Analysis [MRC-CGIDA]. (2020). *Timing and local triggering of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demands.* https://www.gov.uk/government/publications/timing-local-triggering-of-non-pharmaceutical-interventions-npis-to-reduce-covid-19-mortality-and-healthcare-demands-5-march-2020

Norheim, O. F., Abi-Rached, J. M., Bright, L. K., Bærøe, K., Ferraz, O. L. M., Gloppen, S. & Voorhoeve, A. (2020). Difficult trade-offs in response to COVID-19: the case for open and inclusive decision making. *Nature Medicine*, 27, 10–13. https://doi.org/10.1038/s41591-020-01204-6

O'Riordan, T., Cameron, J. (Eds.) (1994). *Interpreting the precautionary principle*. London: Routledge.

O'Riordan, T., Cameron, J. & Jordan, A. (Eds.) (2001). *Reinterpreting the precautionary principle*. Berlin: CMP Publishing.

Ord, T. (2021). *The precipice: Existential risk and the future of humanity*. London: Bloomsbury.

Pettit, P. (2010). *On the people's terms: A republican theory and model of democracy*. Cambridge: Cambridge University Press.

Plutynski, A. (2017). Safe, or sorry? Cancer screening and inductive risk. In Elliott, K. & Richards, T. (Eds.) (2017). *Exploring inductive risk: Case studies in values in science*. Oxford: Oxford University Press., pp. 149-169.

Munthe, C. (2011). *The price of precaution and the ethics of risk*. Dordrecht: Springer.

Science for Environment Policy (2017) *The Precautionary Principle: decision making under uncertainty*. Future Brief 18. Produced for the European Commission DG Environment by the Science Communication Unit, UWE, Bristol. http://ec.europa.eu/science-environment-policy

Scientific Advisory Group for Emergencies [SAGE] (2020a). *SAGE 3 minutes: Coronavirus (COVID-19) response, 3 February 2020*. https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-3-february-2020

Scientific Advisory Group for Emergencies [SAGE] (2020b). *Summary of the effectiveness and harms of different non-pharmaceutical interventions, 21 September 2020.* https://www.gov.uk/government/publications/summary-of-the-effectiveness-and-harms-of-different-non-pharmaceutical-interventions-16-september-2020

Scientific Advisory Group for Emergencies [SAGE] (2021). *SAGE 78 minutes: Coronavirus (COVID-19) response, 28 January 2021*. https://www.gov.uk/government/publications/sage-78-minutes-coronavirus-covid-19-response-28-january-2021

Shiraef, M. A. (2021, 18 March). Closed borders, travel bans and halted immigration: 5 ways COVID-19 changed how – and where – people move around the world. *The Conversation*. https://theconversation.com/closed-borders-travel-bans-and-halted-immigration-5-ways-covid-19-changed-how-and-where-people-move-around-the-world-157040

Siddique, H. (2021, 2 June). Covid catch-up plan for England pupils 'pitiful compared with other countries'. *The Guardian*. https://www.theguardian.com/education/2021/jun/02/union-criticises-pitiful-covid-catch-up-plan-england-pupils

Smart, B., Combrink, H., Broadbent, A., & Streicher, P. (2021). *Report: Direct and Indirect Health Effects of Lockdown in South Africa*. Center for Global Development. http://www.jstor.org/stable/resrep30890

Sober, E. (2020). Infectious diseases and the evolution of virulence. *Auxiliary Hypotheses Blog*. http://www.thebsps.org/auxhyp/evolution-of-virulence-sober/

SPI-M-O (2020). *SPI-M-O: Consensus view on the impact of mass school closures, 17 March 2020*. https://www.gov.uk/government/publications/spi-m-o-consensus-view-on-the-impact-of-mass-school-closures-17-march-2020

Steel, D. (2015) *Philosophy and the precautionary principle: Science, evidence, and environmental policy*. Cambridge: Cambridge University Press.

Stirling, A. (2017) Precaution in the governance of technology. In Brownsword, R., Scotford, E., & Yeung, K. (Eds.), *The Oxford Handbook of Law, Regulation and Technology*.

Sunstein, C. R. (2001). *Risk and reason: Safety, law, and the environment*. Cambridge: Cambridge University Press.

Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*. Cambridge: Cambridge University Press.

Taubenberger, J. K., & Morens, D. M. (2006). 1918 Influenza: the mother of all pandemics. *Emerging Infectious Diseases*, 12, 15-22. http://doi.org/10.3201/eid1201.050979

Tetlow, G., Pope, T., Dalton, G. (2020). *Coronavirus and unemployment: a five-nation comparison*. Institute for Government. https://www.instituteforgovernment.org.uk/publications/coronavirus-unemployment-five-nation-comparison

Thiessen, T. (2020, 13 July). Australia hotel quarantine: You will pay $3000-$5000 returning home, here are state details. https://www.forbes.com/sites/tamarathiessen/2020/07/13/australia-nsw-hotel-quarantine-travelers-pay-3000/?sh=1954d17f4283

Thomas, D. (2020). Supermarkets try to calm food shortage fears amid border chaos. *BBC News*. https://www.bbc.co.uk/news/business-55393076

UK Department for Education (2021, June). *Education Recovery Support for early years settings, schools and providers of 16-19 education*. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/993053/Education_recovery_support_June-2021.pdf

UK Government (2021a). *Coronavirus (COVID-19) in the UK: Deaths*. https://coronavirus.data.gov.uk/details/deaths

UN Environment Programme (2020). *Preventing the next pandemic: Zoonotic diseases and how to break the chain of transmission.*

UN General Assembly (1966, 16 December). *International covenant on civil and political rights*. https://www.ohchr.org/en/professionalinterest/pages/ccpr.aspx

UN Human Rights Committee [UN HRC] (1999, 18 October). CCPR/C/21/Rev.1/Add.9. https://tbinternet.ohchr.org/_layouts/15/treatybodyexternal/Download.aspx?symbolno=CCPR%2fC%2f21%2fRev.1%2fAdd.9&Lang=en

Vlachos, J., Hertegård, E., Svaleryd H. B. (2021). The effects of school closures on SARS-CoV-2 among parents and teachers. *Proceedings of the National Academy of Sciences of the United States of America*, 118, e2020834118 http://doi.org/10.1073/pnas.2020834118

Whitehead, S. J. & Ali, S. (2010). Health outcomes in economic evaluation: the QALY and utilities. *British Medical Bulletin*, 96, 5–21. https://doi.org/10.1093/bmb/ldq033

Wilkes, G. (2020, 27 November). The doubtful case for an impossible Covid-19 cost-benefit analysis. *Institute for Government Blog*. https://www.instituteforgovernment.org.uk/blog/doubtful-case-impossible-covid-cost-benefit-analysis

Walsh, D. (2021, 30 March). Quarantine hotels: Which countries are using them to stop arrivals spreading COVID? *Euronews*. https://www.euronews.com/travel/2021/02/24/quarantine-hotels-which-countries-are-using-them-to-stop-arrivals-spreading-covid

Whitfort, A. (2021, 5 March) China should ban the farming of wild animals now. *LSE Covid-19 Blog*. https://blogs.lse.ac.uk/covid19/2021/03/05/lse-festival-2021-china-should-ban-the-farming-of-wild-animals-now/

Worldometer (2021). *COVID-19 coronavirus pandemic*. https://www.worldometers.info/coronavirus/