Durham Research Online

Deposited in DRO:

02 July 2013

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Beyleveld, Deryck and Brownsword, Roger (2012) 'Emerging technologies, extreme uncertainty, and the principle of rational precautionary reasoning.', *Law, innovation and technology*, 4 (1). pp. 35-65.

Further information on publisher's website:

<http://dx.doi.org/10.5235/175799612800650644>

Publisher's copyright statement:

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Emerging Technologies, Extreme Uncertainty, and the Principle of Rational Precautionary Reasoning

Deryck Beyleveld* and Roger Brownsword*

I Introduction

Broadly speaking, emerging technologies give rise to two kinds of concern: one is that the application of a particular technology might present risks to human health and safety, or to the environment (as is the case, for example, with much of the concern about both synthetic biology and nanotechnologies);¹ and the other is that the technology might be applied in ways that are harmful to moral interests (as is the case with much human biotechnology and neurotechnologies, as well as with information technologies where interests in privacy and confidentiality, and the like, are recurrent concerns).² Those who harbour such concerns demand that regulators should take protective action. However, in many cases, the context in which such demands are made is both deeply contested and clouded by uncertainty—for example, it might be uncertain which types of impact (whether an impact on human health, on the environment, on human rights or human dignity, or whatever) a particular technology might have; or, if so, how likely it is that the impact will eventuate; or, indeed, whether an impact (such as the destruction of human embryos) involves any kind of moral harm.

In such a context, how should regulators respond to calls for action? Quite reasonably, it might be suggested that regulators should strive to maintain a responsible and

-
- University of Durham.
 - King's College London.

¹ Compare the Presidential Commission for the Study of Bioethical Issues, *New Directions: The Ethics of Synthetic Biology and Emerging Technologies* (Washington, December 2010). Having identified, five key principles (namely, public beneficence, responsible stewardship, intellectual freedom and responsibility, democratic deliberation, and justice and fairness), the Commission states that the "principle of responsible stewardship rejects two extreme approaches: an extreme action-oriented [proactionary] approach that pursues technological progress without limits or due regard for public or environmental safety, and an extreme precautionary approach that blocks technological progress until all possible risks are known and neutralized" (p. 26). Instead, as a middle way between proaction and precaution, the Commission advocates "the development of agile, measured oversight mechanisms..." (ibid.). In other words, "[r]esponsible stewardship calls for *prudent vigilance*, establishing processes for assessing likely benefits along with safety and security risks both before and after projects are undertaken" (p. 27). For further elaboration of prudent vigilance as an articulation of responsible stewardship, eschewing both extreme proaction and precaution, see op cit at 123-124.

² In this set of concerns, we should also include important issues relating to the equity of distributing risk in a particular way: whilst it is one thing for a community to accept (all things considered) the risk of, say, air travel, it is another to be a member of that community who lives in the vicinity of an international airport. See, further, Maria Lee, "Beyond Safety? The Broadening Scope of Risk Regulation" (2009) 62 *Current Legal Problems* 242.

rational approach. As the Appellate Body at the WTO put it in the *Hormones* dispute, “responsible, representative governments commonly act [and should act] from perspectives of prudence and precaution where risks of irreversible, e.g., life-terminating, damage to human health are concerned.”³ However, precautionary approaches are frequently accused of being irrational because they focus in a one-eyed way on the need to avoid a particular set of adverse consequences at the expense of considering the probability of the consequences eventuating (as well as ignoring the adverse consequences of making a precautionary intervention).⁴ Nevertheless, in this paper, we argue that, in a carefully defined context of “extreme uncertainty”, it is both responsible and rational for regulators to employ a special form of precautionary reasoning, which we call the Principle of Rational Precautionary Reason (the PRPR).

By “extreme uncertainty”, we refer to a cognitive condition in which regulators believe that it is possible (or, not impossible) that X might “bring about” (cause, result in, or lead to) Z. which is to say that regulators are neither certain that X brings about Z nor certain that X does not bring about Z—that is, they (regulators and their expert advisors) can say only that the probability of X causing Z is in the range $> 0 < 1$. Assuming that regulators have a negative conative attitude towards Z (that is, they fear that X might cause Z), it surely would be irresponsible for them simply to gamble on their fears being misplaced. However, it is far from self-evident that regulators would act rationally and responsibly if they took precautionary measures to protect against Z when: (i) it is not certain that X will lead to Z—indeed, when the likelihood of X leading to Z could be anywhere in the range $> 0 < 1$; and (ii) the conative attitude towards X (sic) is positive (sic) (so that restricting or giving up X has a negative value). Yet, our position is that, in some such conditions, the PRPR may be so engaged.

Even in these introductory remarks, we should emphasise that the PRPR is not to be mistaken for the Precautionary Principle, the PP. The latter, according to critics such as Gary Marchant and Douglas Sylvester, is “an overly-simplistic and under-defined concept that seeks to circumvent the hard choices that must be faced in making any risk management decision.”⁵ Against the PP, the critics argue that it simply is not rational to take protective measures against Z, without weighing the loss of X and without knowing that X will otherwise bring about Z. So, for example, whilst the critics might concede that it is possible that nanoparticles in tennis balls or in cosmetics or in surgical dressings might kill or injure us, or degrade the environment, they insist that, before regulators prohibit or restrict such uses of nanoparticles, it is essential to weigh the value of what would be given up as well as factor in the degree of scientific uncertainty about the relationship between nanoparticles and the harms in question. We think that this line of criticism against the PP is sound; and, to this extent, we are with the critics. However, we also think that there is good sense in such precautionary maxims as “If in doubt, play safe” and “Better safe than sorry”; and, in

³ EC Measures Concerning Meat and Meat Products (Hormones), Report of the Appellate Body WT/DS26/AB/R, WT/DS48/AB/R, 16 January, 1998, at para 124.

⁴ See, e.g., Cass R. Sunstein, *Laws of Fear* (Cambridge: Cambridge University Press, 2005).

⁵ Gary E. Marchant and Douglas J. Sylvester, “Transnational Models for Regulation of Nanotechnology” (2006) 34 *Journal of Law, Medicine and Ethics* 714, at 722.

the special context of extreme uncertainty, we see the PRPR as reflecting this good sense.

These qualifying remarks notwithstanding, the PRPR seems to invite precisely the charge of irrationality that is made against the PP, prompting the obvious question: under what conditions can it be rational to prioritise prevention of Z against continuing to enjoy the benefits of X even though it is not agreed that X is causing or will cause Z? What kind of scenario would it have to be for it to be rational for regulators to proceed in such a purely precautionary way? Stated shortly, the answer is that, given extreme uncertainty, regulators act rationally by taking preventive measures when they judge that the eventuation of (preventable) Z is a worse outcome than the (unnecessary) loss of the benefit of X. In such circumstances, the PRPR indicates that it would be rational to privilege the prevention of Z even though the likelihood of X causing Z cannot be specified (beyond saying that the probability lies somewhere in the range between 0 and 1).

In the paper, we start (in Part II) by sketching the way that the notions of hope and fear fit in our practical reasoning where the context is one of uncertainty. Then we work our way towards the PRPR by following two tracks, one (in Part III) with a focus on concerns about human health, safety, and the environment (which is where we typically find the PP advanced and criticised) and the other (in Part IV) with a focus on moral concerns. In both tracks, the critical context is that of extreme uncertainty and the resort to precaution is presented, in that setting, as a rational regulatory response. Finally (in Part V), we check out the application of the PRPR relative to three test-case technologies in respect of which precaution has been urged: namely, particle accelerators such as the Large Hadron Collider (LHC), nanofoods, and the use of neuroscience and novel neurotechnologies in the criminal justice system.

Before we embark on our discussion, one terminological clarification is in order. As the paper unfolds, it will become clear that the PRPR, *however it is applied*, instantiates a precautionary methodology. If regulators apply the PRPR to introduce a prohibition or a moratorium on X, we might describe this as a precautionary or as a preventive intervention (to avoid the eventuation of Z); and if regulators, having applied the PRPR, decide not to prohibit or restrain X, we might also talk about this as a precautionary or protective response (to preserve or to protect the benefit of X). In other words, where the precautionary methodology is employed, there is a sense in which the regulatory outcome (whether for or against intervention) is precautionary.

II Hope, Fear, and Rational Action in a Context of Uncertainty

In our introductory remarks, we said that the PRPR may be engaged in a context of extreme uncertainty, where regulators believe that there is a possibility that, say, X causes Z; but this means only that they believe that the probability of X causing Z is somewhere in the range of $> 0 < 1$. The attitude that regulators have towards this possibility depends on how they view Z. In principle, regulators may (i) *hope* that X causes Z (where their conative attitude towards Z is positive), or (ii) *fear* that X causes Z (where the conative attitude towards Z is negative), or (iii) be *indifferent* whether X causes Z (where their attitude towards Z is neither negative nor positive). So much is uncontroversial.

However, there is disagreement about the kind of possibility that grounds the state of uncertainty in which it is appropriate to hope or to fear that something is (or is not) the case. So, for example, in debates about the rationality of the PP, critics will sometimes ask: just how much room is there for “speculative” claims (such as the claim that there is a possibility that X causes Z)? In other words, the critics ask, is it enough that those who press for regulatory intervention (or regulators who so intervene) consider that there is a logical possibility that X might cause Z, or must they consider the claim to be supported by some empirical evidence or by theoretical reason (ultimately backed by empirical evidence)—which is to say that there is some inductive probability of X causing Z, without regulators considering that this evidence is conclusive?⁶ In this part of the paper, this is one of the questions that we address.

Imagine that an agent, A, claims to be certain that Q, or claims that there is uncertainty about whether Q is the case. Either way, the issue is whether the claimed certainty or uncertainty is to be understood as phenomenological, psychological or epistemic. For A to be phenomenologically certain that Q is for it to be settled in A’s mind that Q (i.e., for A not to doubt that Q). In simple terms, it is for A to believe that Q. For A to be psychologically certain that Q is not merely for A to believe that Q. A feels that nothing could make A change A’s mind. In other words, A *feels sure* that A’s belief that Q is correct. For A to be epistemically certain that Q is for A to consider that there is conclusive evidence or reason for A to believe that Q, which must be distinguished from A considering that there is sufficient reason for A to believe that Q. For A to consider the evidence for Q to be conclusive but not to believe that Q would be irrational. On the other hand, A might consider that there is

⁶ For example, according to Wolfgang van den Daele, Alfred Pühler, and Herbert Sukopp, *Biotech Herbicide-Resistant Crops: A Participatory Technology Assessment* (Berlin: Federal Republic of Germany Ministry for Research and Technology, 1997) [reported in Stuart J. Smyth, A. Bryan Endres, Thomas P. Redick, and Drew L. Kershen, *Innovation and Liability in Biotechnology* (Cheltenham: Edward Elgar, 2010) at 82-83], we can identify the following three categories of risk:

The first category, probabilistic risks, is defined as those that involve theoretically grounded and empirically demonstrated risks related to the product or its technology. For example, the EPA advisory panel that rejected StarLink maize found troubling aspects of the protein (for example, resistance to digestion) that led to a precautionary decision based on the probability of allergenicity. The methods and much of the evidence about probabilistic risks is available in peer-reviewed journals or public records. The second category, hypothetical risks, involves those possibilities grounded in accepted theory but lacking in empirical experience or evidence that can establish probabilities. A good example of this would be the question of whether antibiotic resistant DNA or other viral proteins from biotech crops could ever merge with gut bacteria (making pathogens resistant to antibiotics)...[This hypothetical risk is] undergoing extensive testing....

The third category, speculative risks, has neither a credible hypothesis nor empirical experience to indicate the actual existence of these risks. When using the precautionary approach to [assess] biotech crop risks, theories range from Bt-pollen impacting bees...the “Canola Myth”, and semi-spiritual ideas about “biodynamic” agriculture that are very hard to explain as scientific hypotheses....These thought exercises illustrate that almost any correlation can be made to show the potential for risk, irrespective of whether there is any qualified theoretical basis for the speculative possibility.

sufficient evidence to justify A in believing that Q without considering that the evidence is so strong that A would be irrational not to believe that Q.

It is not possible here to justify fully the position we adopt on these matters. That is altogether too large a task for a single paper. But since our main objective is merely to outline an approach to the rationality of precautionary reasoning, it will suffice if the position we present has at least some degree of plausibility.

In our view, when A hopes (or fears) that Q, A is phenomenologically uncertain that Q. *With the question whether that Q or that not-Q in mind*, A neither believes that Q nor believes that not-Q (disbelieves that Q); A is in doubt about whether or not that Q. It is surely a *minimum* requirement that A considers that Q is possible (the question is whether it is enough). If A neither believes nor disbelieves that Q then A cannot consider Q to be impossible. However, does it follow that A considers Q to be possible? Well, not if A is not thinking about Q at all. But if A has Q in mind it surely follows that A must consider that Q to be possible. However, we will not continually emphasise this. It is to be taken as understood.

Why must the uncertainty be phenomenological? Let us suppose that Jack has said goodbye to his girlfriend Jill at Durham railway station when she got on a train to London. Suppose that, a couple of hours later, he hears on the radio that this train has been involved in an accident and that many passengers were seriously hurt. Jack would be very upset if Jill has been hurt. He wants her not to be hurt. Does he hope that she has not been hurt (i.e., fear that she has been hurt)? On our account, only if he neither believes nor disbelieves that she has been hurt. Why? Well, what would lead Jack, supposing that he fears that Jill has been hurt, to cease to do so? For Jack to receive a communication informing him that Jill has been hurt, or equally, a communication that she has not been hurt, would surely be sufficient *provided that* Jack believes what he is told. Such a communication that she has been hurt would extinguish his hope (realise his fear), whereas such a communication that she has not been hurt, would realise his hope (extinguish his fear). Suppose, then, that Jack receives a call from the hospital to say that Jill has been seriously hurt. This is really bad news. Jack does not want it to be true. Perhaps he will refuse to believe it. He will not believe it until he goes to the hospital and sees for himself. It might, he thinks, be a case of mistaken identity. If so then Jack will continue to hope that Jill has not been hurt. But suppose, despite the fact that he is not epistemically certain (it could be a case of mistaken identity), Jack believes that Jill has been hurt. Then he will not hope that Jill has not been hurt, though he might hope that the hospital has made a mistake. That Jack can and does hope that his belief that Jill has been hurt is mistaken does not translate into Jack hoping that Jill has not been hurt.

Therefore, we contend that “A hopes that Q” is equivalent to “A values that Q, and *having in mind the question whether or not that Q*, neither believes nor disbelieves that Q”. As we have already indicated in our Jack and Jill example, we hold that “A hopes that Q” is equivalent in its truth conditions to “A fears that not-Q”. “A fears that Q” is, therefore, to be analysed as “A values that not-Q, and neither believes nor disbelieves that not-Q.”

We are not, however, directly interested in statements of the form “A hopes that Q”, but in statements of the form “A hopes to bring it about that Q”, where A acts to bring

it about that Q.⁷ Specifically, we are interested in cases where Q is the prevention of some adverse event, Z, where bringing it about that Q amounts to preventing the occurrence of Z. For it to be rational for A to act to prevent Z, A must either believe that Z will occur without preventive action or at least fear that Z will occur without such action. It is irrational for A to act to prevent Z if A believes that action to prevent Z is unnecessary because Z will not occur even if no preventive action is taken. It follows that the rationality of action to prevent Z rests on the rationality, *inter alia*, of A (at the very least) not believing that Z will not occur if nothing is done to prevent Z occurring.

To set this in the standard context for *regulatory* precaution, the case is one in which (i) the regulator, R, fears that some activity X might cause some adverse event that involves some harm or damage Z, and (ii) R believes that regulatory intervention Q will prevent the occurrence of Z. Applying the above analysis, this standard case involves the following: first, R values that Z should not occur (or negatively values the occurrence of Z); secondly, R neither believes nor disbelieves that doing X will produce Z; thirdly, R does not believe that Z will not occur if no intervention is made; and, fourthly, R believes that intervention Q will, or at least might, serve to prevent Z.

Yet, how can it be rational for R to take some precautionary intervention, Q, simply on the basis of a fear that X might be “unsafe” relative to Z? Given that R’s fear, on our analysis of fear, entails neither believing nor disbelieving that X is safe, it follows that the justificatory burden rests on R’s judgment that occurrence of Z, were it to eventuate, would be unacceptable. If (i) Q were costless, (ii) it were agreed that X might cause Z, and (iii) Z were disvalued, then R might be able to plead that it is rational to act on the maxim of “better to be safe rather than sorry”. However, regulators will rarely find it so easy to justify precautionary intervention. For, in many cases, Q, the particular precautionary intervention, will involve some restriction on X, an activity that is valued. Typically, it is not only a matter of showing that it is rational to take precautionary action simply on the basis of fear, but that it is rational to do so despite this involving some sacrifice of a valued activity. In short, how can it be rational for regulators to act in a way that certainly gives up X, or some part of X, which is judged to have a positive value, simply because R thinks that it is possible that X might cause Z (Z being valued negatively)?

Of course, if R believes that Z certainly will occur if X is permitted, then it is subjectively reasonable for R to judge whether or not to permit (or prohibit) X simply by weighing the positive value R attaches to X against the disvalue of Z in order to decide whether or not to act against X. But if R neither believes nor disbelieves that Z will occur as a result of doing X (merely fears such an outcome), then this is compatible with R considering damage to Z to be logically possible or inductively possible (i.e., inductively having any probability $> 0 < 1$). One of the claims made by those who are sceptical about the rationality of precautionary reasoning is that more than logical possibility is required. However, as we will argue later, we consider that

⁷ It should be clear that if A acts to bring it about that Q, then A values that Q (disvalues that not-Q) and either believes that A’s act will bring it about that Q or hopes that A’s act will bring it about that Q. “A hopes to bring it about that Q” is to be analysed as “A values that Q; A does not believe that Q is already the case; and A neither believes nor disbelieves that A will bring it about that Q.”

whether or not this is so depends very much on the nature of uncertainty about the object of hope or fear.

III The PRPR as a Prudential Precautionary Principle

Where an emerging technology (X) elicits concerns about possible harm to human health, safety, or the environment (Z), responsible regulators will undertake an expert risk assessment that seeks to establish (i) which, if any, of the alleged harms might be relevant and (ii) how likely they are to eventuate.⁸ Relying on the expert assessment to balance the benefits of X against the risk of X causing Z, regulators will strive to set the regulatory environment in such a way that risk is tolerated only so far as this is acceptable. Of course, the notion of “acceptable risk” is eminently contestable: some regulatees might disagree with the regulators’ determination of what constitutes an acceptable risk; but, insofar as any exercise of regulatory prudence is straightforward, this is one such case. In practice, however, it is often more difficult for regulators to discharge their prudential responsibilities because the risk assessment is compounded by uncertainty.

In this part of the paper, we discuss the way in which rational regulators should (and should not) respond to uncertainty. We start with the kind of “scientific uncertainty” that is usually a cue for the invocation of the Precautionary Principle (PP); and then we deal with the case of extreme uncertainty that is focal for our present purposes.

(i) Precaution and scientific uncertainty

According to Principle 15 of the Rio Declaration of the UN Conference on Environment and Development (1992) (arguably, this being the foremost articulation of the PP):

In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.

In this context, “lack of full scientific certainty” (or “scientific uncertainty”) signals that, in the expert scientific community, there are different views about whether X causes Z or about the likelihood of Z eventuating. Clearly, regulators need to factor

⁸ Risk analysis involves assigning a disvalue to an undesirable possible consequence of an activity (referred to as a hazard). The greater the disvalue of the hazard, the more severe the hazard is said to be. The risk of the hazard incurred by the activity is the probability of the hazard eventuating multiplied by the severity of the hazard. An ability to attribute a probability or range of probabilities to a hazard eventuating is absolutely essential for a risk analysis, as is the ability to identify a possible hazard. Standard cost-benefit analysis is a strategy for determining whether or not it is worth taking a risk on something. The risk of X (which is now referred to as the expected cost of X) is weighed against the expected benefit of X. The expected benefit of X is a product of the desirability of the hoped for outcomes of X and the probability of the benefit. In a cost-benefit analysis the expected cost of X is compared with the expected benefit of X (either by the benefit minus the cost or by dividing the benefit by the cost, the latter yielding the utility of X if positive or the disutility of X if negative). The utility of X can be compared with the utility of Y, and so on.

into their calculations the fact that such differences of expert opinion exist. The question is: in such a context, is it rational for regulators to take a precautionary approach?

In many international instruments,⁹ given scientific uncertainty, a precautionary approach has crystallised into a version of the PP. However, critics of the PP argue that it is not a rational basis for regulatory intervention. In some small part, the problem is that the PP can be articulated and interpreted in many different ways;¹⁰ but, the major objection is that precaution is apparently urged not only without taking any account of the cost of the intervention (in particular, the loss of whatever value or benefit X has)¹¹ but also without it being certain that X will cause (or is already causing) Z. There is no need for us to rehearse this well-worn debate. Rather, we will make four short comments that will help to orient readers to our position.

First, there is no obvious improvement in the argument for precaution when the context changes from one of scientific certainty to one of scientific uncertainty. Suppose a situation of scientific certainty: the experts agree, let us suppose, that the use of mobile phones causes headaches. Before responsible regulators take precautionary measures, they will count the cost of giving up or restricting the use of mobiles. Let us suppose that regulators decide against a precautionary intervention; the benefit of mobiles is high, and headaches are a minor problem. If regulators so decide in a context of expert agreement, it is not obvious why a change to scientific uncertainty should improve the case for precautionary intervention—it will only do so if some experts now suggest that the harm to users is more serious than the occasional headache. Similarly, if regulators initially decide in favour of precautionary intervention, the change from scientific certainty to uncertainty does not necessarily strengthen the initial case for precaution—again, it will do so only if some experts now suggest that the harm to users is more serious than the occasional headache.

Secondly, appeal to the PP might be designed to challenge regulators who routinely respond to scientific uncertainty by procrastinating. Although, as Jonathan Zittrain has argued, regulatory procrastination has served us well in the development of information technologies,¹² it does not pass muster as an across-the-board rational response.¹³ We agree, therefore, with those who contend that it is irrational for regulators routinely to eschew precautionary intervention until there is full scientific certainty.

⁹ See EC Commission's Communication on the Precautionary Principle, Annex II, COM(2000) 1, Brussels 02.02.2000.

¹⁰ See, e.g., Neil Manson, "Formulating the Precautionary Principle" (2002) 24 *Environmental Ethics* 263; Elizabeth Fisher, Judith Jones, and René von Schomberg (eds), *Implementing the Precautionary Principle: Perspectives and Prospects* (Cheltenham: Edward Elgar, 2006); and Elizabeth Fisher, *Risk Regulation and Administrative Constitutionalism* (Oxford: Hart, 2007).

¹¹ See, e.g., Cass R. Sunstein, *Laws of Fear* (Cambridge: Cambridge University Press, 2005).

¹² Jonathan Zittrain, *The Future of the Internet* (London: Penguin, 2008).

¹³ Compare the comments of the Presidential Commission for the Study of Bioethical Issues, note 1 above.

Thirdly, as we have already intimated, we agree with those critics of the PP who argue that it cannot be rational for regulators routinely to respond to scientific uncertainty by making a precautionary intervention. This is not to say that a precautionary intervention is never rationally justified; but such an intervention will only be a prudent and responsible response where it follows from an all things considered judgment (including taking into account the loss of benefit occasioned by the intervention).

Fourthly, we would not rule out the possibility that regulators can operate with a rationally defensible presumption in favour of precaution where Z belongs to a class of extremely harmful outcomes. For example, if some experts take the view that X, without precautionary intervention, could have catastrophic effects,¹⁴ or might endanger the human species,¹⁵ or might destroy the essential infrastructure for human existence,¹⁶ or the like, regulators might treat the burden of justification as now lying on those who oppose such precautionary measures. To put this another way: we suggest that no one would object to a principle saying, “Protect against (avoid, restrict) activities posing real serious/irreversible threats to human health and safety, or to the environment, to the extent possible and proportionate.” The real question is to determine under what conditions, if any, it is reasonable to impose a burden of proof on those who wish to conduct an activity that poses only a potential threat to the human health and safety, or to the environment to show that the threat is not real if action against the activity is to be called off.¹⁷

¹⁴ See, e.g., Cass R. Sunstein, *Worst-Case Scenarios* (Cambridge, Mass.: Harvard University Press, 2007), esp Ch. 3. At 167-168, Sunstein develops the following precautionary approach:

In deciding whether to eliminate the worst-case scenario under circumstances of uncertainty, regulators should consider the losses imposed by eliminating that scenario, and the size of the difference between the worst-case scenario under one course of action and the worst-case scenario under alternative courses of action. If the worst-case scenario under one course of action is much worse than the worst-case scenario under another course of action, and if it is not extraordinarily burdensome to take the course of action that eliminates the worst-case scenario, regulators should take that course of action. But if the worst-case scenario under one course of action is not much worse than the worst-case scenario under another course of action, and if it is extraordinarily burdensome to take the course of action that eliminates the worst-case scenario, regulators should not take that course of action.

¹⁵ Compare, e.g., George J. Annas, “The ABCs of Global Governance of Embryonic Stem Cell Research: Arbitrage, Bioethics and Cloning” (2004) 39 *New England Law Review* 489.

¹⁶ Compare, Roger Brownsword, “Responsible Regulation: Prudence, Precaution and Stewardship” (2011) 62 *Northern Ireland Legal Quarterly* 573.

¹⁷ We might pursue this question by focusing on the seemingly contrasting regulatory cultures in Europe (where the emphasis is on ex ante precaution, together with the assumption that technologies are to be presumed to be unsafe unless shown to be safe) and the United States (where the emphasis is on ex post precaution, together with the assumption that technologies are to be presumed to be safe until shown to be unsafe). No doubt, this is a caricature of the respective regulatory cultures and practices. Nevertheless, if we were to conduct, so to speak, an “innovation audit”, would we be able to defend European regulatory precaution as rational? Are there any good reasons why regulators in Europe and the United States should make either of the presumptions that underpin their regimes? Against Europe, and in support of the United States, it might be claimed that: (i) innovation is intrinsically good (so it should be presumed that new technologies are safe unless there are good reasons to think otherwise); (ii) nothing is completely safe (if we were to insist on zero-risk, we would be living, for example, without electricity); and (iii) attitude optimism is more prudentially better than attitude

There is, of course, much unfinished business in these short remarks about the PP; however, our focal interest in this paper is the PRPR and the context of extreme uncertainty. With prudential responsibilities in mind, it is to this focal point that we now turn.

(ii) Extreme uncertainty and the PRPR

In conditions of extreme uncertainty, where the question is whether X might cause Z, the risk assessors will advise regulators that they cannot say that X certainly will cause Z, nor that it certainly will not; all that they can say is that this is a possibility with a likelihood in the range $> 0 < 1$. Under these conditions, all that regulators have to go on are the X (with whatever benefit X represents) and the Z that is feared. This is not enough for regulators to undertake anything resembling a standard balancing of benefit and risk; it is futile to seek out an acceptable level of risk. Under these conditions, as we have said, regulators are in the realm of hope and fear. And, the question is whether it can be rational for regulators to respond to their fears by applying precautionary measures.

Essentially, in a context of extreme uncertainty, regulators have only two options, each of which divides into two possibilities (getting it right or getting it wrong). Where regulators fear that X might cause Z, while one option is to make an intervention (Q) that is designed to prevent X causing Z (let us assume, by restricting X), the other is to make no such intervention (not-Q) in the hope that X will not after all cause Z. Regulators must, so to speak, place their bets: to intervene or not to intervene.

If regulators take the first option (intervention), they might get it right or wrong in the following sense: they will have got it right if X would otherwise have caused Z and the intervention, Q, has prevented this happening; and they will get it wrong if X would not have caused Z so that the intervention and restriction on X is unnecessary. Of course, if X has no positive value, the restriction on X is not a negative and this is an easy case—one in which regulators can rest on the aphorism that it is better to be safe than sorry.

If regulators take the second option (no intervention), they might again get it right or wrong as follows: they will have got it right if X does not cause Z and the benefit of X (assuming that X has a positive value) continues to be enjoyed; and they will get it wrong if, absent a preventive intervention (not-Q), X does cause Z.

Which is the rational response? We suggest that regulators should ask themselves the following question: if we get it wrong, which of the options would get it wrong in the least acceptable way? Stated formally, the PRPR provides:

If, under conditions of extreme uncertainty (where there is irresolvable doubt), regulators have to choose between Q and not-Q (here, intervention or no

pessimism. However, it is clear that each of these arguments raises more questions than it answers. In general, compare Stuart J. Smyth, A. Bryan Endres, Thomas P. Redick, and Drew L. Kershen, *Innovation and Liability in Biotechnology* (Cheltenham: Edward Elgar, 2010).

intervention), then choose whichever option, in the event of error (of getting it wrong), avoids the least acceptable outcome.

Where the frame of regulatory reference is prudential, then this means that whichever (in error) outcome is least acceptable relative to the interests of the regulatees (for whom regulators act as prudential proxies) should be avoided. If the (preventable) occurrence of Z (for example, if Z is catastrophic) is the least acceptable outcome, then precautionary intervention is rationally called for. If the (unnecessary) loss of the benefit of X is the least acceptable outcome, then precautionary reasoning indicates that intervention is not rationally justified.

To avoid any misunderstanding, two points should be noted. First, as we flagged up in our introductory remarks, although it is tempting to think of the intervention option as the precautionary option (because it involves adoption of measures to prevent Z), strictly speaking, *both* options are precautionary; for, even when regulators opt for non-intervention, they do so for the reason that they fear unnecessarily sacrificing the benefit of X. Secondly, even though regulators might get it wrong (by choosing the wrong option), they can still (and at the same time) act rationally and get it right relative to the PRPR.

Before we move on to precautionary reasoning where the concerns are of an explicitly moral nature, we can give an example of how the PRPR might be used by regulators who are called upon to account for their prudential precautionary actions.¹⁸ Let us suppose that, when the Icelandic volcano, Eyjafjallajökull, erupted in Spring 2010, air traffic regulators found themselves operating in conditions of extreme uncertainty. Actually, we do not think that this was the context because there was evidence of the hazards presented by volcanic ash in the atmosphere; nevertheless, let us suppose that this was a case of extreme uncertainty. On that basis, regulators might have justified their precautionary measures in the following terms: “We know virtually nothing about the hazards presented by clouds of volcanic ash. Our best advice is that it is possible that the ash could cause aircraft to crash. Given this possibility, we judged that it would be irresponsible to do anything other than restrict air traffic. We understand that the disruption to air travel was inconvenient; and, no doubt, some will think that the restrictions were unnecessary—and, who knows, they might be right. However, if we had not taken precautionary action and, if the lives of passengers and crew had been lost when aircraft were brought down, we take it that this would have been judged to be a worse outcome than any unnecessary disruption and inconvenience that might have been caused by the temporary restrictions on air travel.”¹⁹

¹⁸ A similar kind of regulatory response might have been given by the WHO when, in the context of swine flu (H1N1 virus), it was criticised for having caused unnecessary panic and disruption. For discussion of how far the quite different value of solidarity might have shaped the actions of individuals and authorities, see Barbara Prainsack and Alena Buyx, *Solidarity* (London: Nuffield Council on Bioethics, 2011) Ch 7.

¹⁹ Suppose, however, that regulators are faced with a proposal to do Q, which offers some chance of benefits B, but also runs risks of hazards H in the following conditions. The context is not one of extreme uncertainty; however, experts can only suggest a range of probabilities for B and H. They estimate the probability of B as 30%-70%, and that of H as 5%-20% and have no idea about what figures are more likely within these ranges. Assume too that we have a way of measuring the disvalue (seriousness) of H and the value of B and that we have a

In sum, the PRPR provides that, in conditions of extreme uncertainty, it is rational for regulators to make a precautionary intervention if this is required in order to prevent what would be the least acceptable outcome of getting the decision wrong. Of course, regulators will hope that they get their decision right; but, in the special case of extreme uncertainty, the critical point is that the PRPR can commend an intervention as rational even if (as the omniscient would know) it is actually wrong.

IV The PRPR as a Moral Precautionary Principle

Emerging technologies, as we have said, elicit a range of concerns. Whilst some concerns relate to questions of human health, safety, and the environment, there are also many questions about whether new technologies are applied in ways that are compatible with human rights as well as about their relationship to the much-contested value of human dignity.²⁰ One of the most contested questions of this latter kind has concerned the status of a human embryo. If, as many believe, it is a grave moral wrong to use human embryos as research tools, then regulators will be pressed to prohibit the creation of human embryos for research as well as to exclude from patentability the products or processes of such research.²¹ Famously, regulators have been so pressed in relation to the permissibility and patentability of human embryonic

scale that renders the disvalue of H and the value of B commensurable. If we could put a price on everything (as cost-benefit analyses often assume) then we would be able to do this in principle. So, suppose that B would gain us €10,000, but H would lose us €20,000. Should we do Q? On cost-benefit analysis principles, the benefit of B ranges from €3,000-€7,000, depending on which probability estimate we use. The cost of H ranges from €1,000-€4,000. If we are optimists we will be inclined to assume that the probability of B is 70% and the probability of H is 5%. We will look at doing Q as producing a probable nett benefit of €6,000 and do Q. If we are pessimists we will be inclined to assume a probability of B of 30% and a probability of H of 20%. We will look at Q as producing a probable nett cost of €1,000 and not do Q. If we are just inclined to take the average of the estimates, we will calculate the probable benefit as €5,000 and the probable cost as €2,500 and again do Q. This will be the rational thing to do if the range is based on a statistically random distribution where the average will in fact be the most likely to be correct; but we have assumed that this is not the case. Can we use precautionary reasoning? If we suppose, as we have, that we really have no rational way of deciding what estimates are most likely to be correct, then we can: we are in total (empirical) doubt about what probability estimates to use. So what does the PRPR tell us to do? It does not here tell us to compare assuming we do Q and are wrong against assuming we do not do Q and are wrong discounting any thought of probabilities (which would produce a nett cost figure of €10,000), because we are not in total doubt about the probabilities of B and H. We are in total doubt about the correct probability estimates within ranges of probability. Here, it asks us to calculate the difference it makes between being an optimist and being wrong and being a pessimist and being wrong, where being an optimist and a pessimist is defined by our probability ranges. If we are optimistic we go for a gain of €7,000 and stand to lose €4,000 if we are wrong. If we are pessimistic we act to save ourselves a cost of €4,000 and stand to lose a gain of €7,000. We do Q, because we have more to lose by being pessimists and being wrong than we have to lose by being optimists and being wrong (or more to gain by being optimists and being right than we have to gain by being pessimists and being right!). Erring on the side of caution means avoiding the worst error, and making the pessimistic assumption, acting accordingly, and being wrong is the worst error in this case.

²⁰ See Deryck Beyleveld and Roger Brownsword, *Human Dignity in Bioethics and Biolaw* (Oxford: Oxford University Press, 2001).

²¹ For an overview, see Samantha Halliday, "A Comparative Approach to the Regulation of Human Embryonic Stem Cell Research in Europe" (2004) 12 *Medical Law Review* 40.

stem cell research;²² and, in Europe, there have been highly controversial decisions on patentability at both the European Patent Office and at the European Court of Justice.²³

Various, the key question in these regulatory debates might be put as: is the human embryo “a life”, or “a life in being”, or a “human life” or “an agent”, or “a bearer of rights”, and so on? If the answer to any of these versions of the question is “yes”, then this has major implications for the legitimacy of using human embryos as research tools. But, what if the answer is “possibly”? Is this an occasion for some kind of precautionary regulation? Our discomfort in relation to this question is not likely to be eased with further developments in biotechnologies (e.g., with hybrids and chimeras), robotics and artificial intelligence. Where we are uncertain about the moral status of an entity, what is the responsible regulatory response?

In the first section of this part of the paper, we set out the argument for believing in the existence of other agents based on a categorical imperative that requires other agents to be treated always as ends and never simply as means. This, we believe, is a paradigm of rational precautionary reasoning, driven not by prudential considerations but by a moral imperative. In the second section, we translate this principle of rational precautionary reasoning for individual moral agents into a general principle for the guidance of regulators in conditions of extreme uncertainty.

(i) The Rationality of Precautionary Reasoning

Consider the case of A, who aspires to act as a moral agent. Let us suppose that A accepts (as Kant and Gewirth both maintain can be demonstrated) that there is a categorically binding imperative requiring all agents to treat all agents never merely as a means, but always as ends-in-themselves.²⁴ The idea that this imperative is categorically binding is the idea that agents contradict that they are agents if they do not accept that they must act in accordance with this principle.²⁵ Suppose, however,

²² See, e.g., Roger Brownsword, *Rights, Regulation and the Technological Revolution* (Oxford: Oxford University Press, 2008) Ch 2.

²³ See Aurora Plomer and Paul Torremans (eds), *Embryonic Stem Cell Patents* (Oxford: Oxford University Press, 2009) (particularly on the *Wisconsin Alumni Research Foundation* (the WARF) case: Case G 0002/06, November 25, 2008). At the ECJ, the key case is *Brüstle v Greenpeace eV*: Case C 34-10 (judgment given 18 October, 2011).

²⁴ This is Kant’s Principle of Humanity, his second formula for what he calls the Categorical Imperative. (See *Groundwork of the Metaphysics of Morals* (4:428) (ed. Mary Gregor) (Cambridge: Cambridge University Press 1998) 37. The Principle of Generic Consistency (PGC) which Gewirth claims has the status that Kant claims for the Categorical Imperative, arguably entails the Principle of Humanity, but has a richer content (see Deryck Beyleveld and Roger Brownsword, *Human Dignity in Bioethics and Biolaw* (Oxford: Oxford University Press 2001) Ch. 5.

²⁵ According to Kant, the Categorical Imperative is “connected completely a priori with the concept of the will of a rational being as such” (*Groundwork* 4:426 (35)), which he claims is shown by the fact that the Categorical Imperative is the law of a free-will (see 4:447 (53)) and rational beings with a will (agents) must suppose that they have free-will if they are not to contradict the idea that they are acting (see 4:449 (54)). Gewirth’s argument for the PGC takes a different path but purports to provide the PGC with the same epistemological status as Kant claims for the Categorical Imperative.

that A nevertheless denies that “the Categorical Imperative” has any categorically binding effect on A *in practice*, on the grounds that A cannot be certain that there are any agents in the world other than A. While A, in accepting the Categorical Imperative, accepts that A contradicts that A is an agent by not accepting that A ought to treat all agents as ends-in-themselves, A does not accept that A contradicts that A is an agent by denying that any other being, B, is an agent, even if B behaves as though B is an agent. To be sure, A must concede that it is logically possible that B is an agent. But it is also logically possible that B is a mere figment of A’s imagination, and even if B is not, B might merely be a mindless automaton programmed to behave like an agent. To be an agent, B must be a self-conscious thinking subject who acts for ends that B feels that B has chosen; but the only being that A must, without making wholly unverifiable or unfalsifiable metaphysical assumptions, hold to be an agent on pain of contradicting that A is an agent is A. Hence, reasons A, A is not categorically bound to treat B as an end-in-itself, even if B behaves as though B is an agent.

A’s reasoning is sound if A’s acceptance of the Categorical Imperative does not commit A to accepting that A contradicts that A is an agent by not regarding B as an agent when B behaves like an agent. However, if A assumes that B is an agent, and treats B accordingly (which A can do, because B behaves like an agent), and happens (wholly unknowably) to be wrong (because B is, in fact, not an agent), then A will not have violated the Categorical Imperative. No agent will have been treated merely as a means to A’s ends. On the other hand, if A assumes that B is not an agent, and acts accordingly, and happens (again wholly unknowably) to be wrong (because B, in fact, is an agent), then A will have violated the Categorical Imperative by treating an agent merely as a means to A’s ends. For A to risk violating the Categorical Imperative when A can avoid doing so (which A can do because B behaves like an agent) is not to respect the imperative as categorically binding, and this, on A’s starting presumption, is to contradict that A is an agent. Consequently, A must accept that A categorically must treat all beings that behave as though they are agents as agents or give up the idea that the “Categorical Imperative” really is categorically binding. The practical effect of A’s presumption that “the Categorical Imperative” is categorically binding is that A must act as though it is certain that B is an agent, i.e., act as though there is a probability of 1 of B’s being an agent, when absent that presumption (or various wholly speculative metaphysical assumptions) that probability is completely indeterminate. It follows that A must act as though A believes that B is an agent. (Whether this means that A must actually believe that B is an agent is moot, as it makes no difference for practical purposes).

It follows that if there is a wholly cogent argument for a categorical imperative requiring agents to treat all agents as ends-in-themselves, then both the probability that another being is an agent and the non-speculative use of behavioural evidence as evidence for agency can be grounded in morality and can only be grounded in morality.²⁶ But, what if there is no such argument? Well, then there is also no non-

²⁶ This argument, which is, in effect, an argument for other minds as well as for other agents, may also be extended to form a precautionary moral argument for the existence of a real world independent of the senses. Only if there is a real world existing independent of A’s senses can A harm any other agents and the Categorical Imperative have any practical application. A cannot know for certain whether there is a real world existing independent of A’s senses. However, if A supposes that there is no such world, acts accordingly, and is wrong, then A has acted contrary to the total prohibition of the Categorical Imperative. On the other hand, if A

speculative basis to treat behavioural evidence as evidence of agency. Those who happen to suppose that there is such a categorical imperative will be bound to treat apparent agents as agents, but it will not be unreasonable for those who do not accept this imperative not to treat apparent agents as agents unless they accept wholly speculative metaphysical assumptions.

The most important point, though, is that unless it can be shown that it is irrational to hold that there is a categorical imperative and not merely that belief in a categorical imperative is not rationally required, it will not be shown to be irrational to act according to the dictates of such an imperative.

The form of the precautionary argument for the recognition of other agents is in accordance with the following principle:

If in total irresolvable doubt (uncertainty) as to whether that Q or that not-Q, assume categorically that Q and act accordingly, because to err by assuming that not-Q (i.e., to assume that not-Q and be wrong) is categorically prohibited, while to err by assuming that Q is not categorically prohibited (provided only that it is possible and meaningful to assume that Q and act accordingly).

We contend that this statement is analytic; and can be more succinctly stated as:

If in total irresolvable doubt as to whether that Q or that not-Q, in deciding between that Q and that not-Q, categorically avoid making an error that constitutes doing something that is categorically prohibited (to the extent that it is possible and meaningful to do so).

This is a special case of the following principle:

If in doubt as to what to choose between that Q and that not-Q, guard against making the potential error that is worse

which can be expressed rather more neatly but more vaguely as

If in doubt, play safe!

This, in our opinion, is the very essence of precautionary reasoning. However, in the moral argument for other minds, special qualifications apply. These qualifications are:

- (1) The doubt (that Q v that not-Q) is of a totally irresolvable kind.
- (2) It is categorically better to err on the side of accepting that Q and acting accordingly.

assumes that there is such a world, acts accordingly, and is wrong, A does not act contrary to this prohibition. Therefore, A *morally* must suppose that there is a real world existing independently of A's senses. (See, Deryck Beyleveld and Shaun D. Pattinson, "Defending Moral Precaution as a Solution to the Problem of Other Minds: A Reply to Holm and Coggon" (2010) 23 *Ratio Juris* 258. On this basis, Kant was right to claim the primacy of practical reason over theoretical reason in his *Critique of Practical Reason* 120-121 [216-217].

- (3) It is possible and meaningful to accept that Q and to act accordingly.

When these qualifications apply, it may be concluded that

It is categorically required to accept that Q and to act accordingly.

It is very important to emphasise that playing safe means playing on the side that it is least bad—that is, least bad *morally*—to err on. That side is not decided before the seriousness of following Q and being wrong is weighed against the seriousness of following that not-Q and being wrong. If it turned out that (contrary to what is in fact the case in the example) following that not-Q and being wrong were the less heinous option as judged by whatever master moral principle is in play (in our case, the Gewirthian Principle of Generic Consistency), then playing safe would be constituted by accepting that not-Q and acting accordingly. “Playing safe” does not mean being inactive rather than active. It does not mean refraining from pursuing potential tangible benefits because of merely possible risks of a serious nature. “Playing safe” just means “acting to avoid the more serious moral wrong”.

It is equally important to note that while any argument exhibiting this form will be valid, the actual substantive result will depend on what the categorical imperative in play (for us, as we have said, it is the PGC) actually requires. If we substitute some other alleged supreme moral principle held to be categorically binding, a different result might be obtained.²⁷

In relation to the problem of “other agents”, the result justified (indeed, required) is a total prohibition on acting on that not-Q. This is a consequence of the doubt being totally irresolvable and it being categorically worse to be wrong having assumed that not-Q (that an apparent agent is not an agent) than to be wrong having assumed that Q (that an apparent agent is an agent). That the doubt is totally irresolvable means that nothing can be said about the independent (ex ante) cognitive status of that Q and that not-Q beyond that it is *logically* possible that either assumption could be true. Beyond this, the cognitive status (uncertainty) of that Q and that not-Q can play no part in rational deliberation about what to accept or do in relation to that Q or that not-Q. Its sole function is to make it rational to hope or fear that Q. Hence the only factor to consider is which possible error constitutes the worst/best outcome. Since the worst outcome here (not treating an agent as a generic rights bearer) attaches to being in error when acting on that not-Q, and this is categorically prohibited by the standard of evaluation in play, acting on that not-Q is categorically prohibited (which is equivalent to acting on that Q being categorically required).

Of course, all prescriptions are subject to the principle that “ought” implies “can”. It must be meaningful and practicably possible to follow the prescription. This is

²⁷ Clearly, in a moral frame of reference, precautionary reasoning can never be value-free. However, we have sketched a view elsewhere that holds that, in practical reason generally, precautionary reasoning can never be value-free. If that is thought to be a problem, then our short response is simply that cost-benefit analysis, which is frequently held up as the paradigm of scientific risk assessment, can never be value-free either. See, further, Deryck Beyleveld and Roger Brownsword, “Complex Technology, Complex Calculations: Uses and Abuses of Precautionary Reasoning in Law” in Marcus Duwell and Paul Sollie (eds), *Evaluating New Technologies: Methodological Problems for the Ethical Assessment of Technological Developments* (Springer, 2009) 175.

important because it is surely the case that it is logically possible, for example, not only for other humans to be agents, but also for cats and dogs, plants, and even rocks and tables to be agents. Without attention to the principle that “ought” implies “can”, the precautionary argument we have sketched would lead to the absurdity of requiring agents to treat everything as an agent. But, of course, it does not lead to any such requirement. By granting the generic rights to agents, the PGC also imposes reciprocal duties on agents, and, in any case, the generic rights can only be exercised by agents. Thus, it is only practicably possible and meaningful to treat as agents *those beings that behave as though they are agents* (i.e., those beings that we can characterise as apparent agents). Hence, the precautionary argument for other minds only categorically requires apparent agents to be regarded as agents (and, hence granted the generic rights). It does not, however, follow that the PGC does not grant any protection to beings that are not apparent agents (“apparent non-agents”). One effect of the precautionary argument for other minds is that evidence for apparent agency must be regarded, under precaution, as evidence of agency. Evidence that is sufficient for a being to be judged an apparent agent is sufficient for the being to be regarded as an agent for all practical purposes. Behaviour, capacities and features of a being that are necessary, though by themselves not sufficient for apparent agency are nonetheless to be regarded as necessary, albeit insufficient, evidence of agency under precaution.

Now, just because a being, say a dog, is not judged to be an apparent agent, does not mean that it is not an agent. It is (outside of precaution) just as uncertain (possible) that it is or is not an agent as it is uncertain (possible) that an apparent agent is or is not an agent (outside of precaution). Furthermore, apparent agents sometimes seem to lose the capacity to display their apparent agency. An important case to reflect on is where apparent agents, having lost that capacity, regain it and declare that they were fully conscious and retained all their mental capacities throughout the period that no agency-like interaction was possible with them. So, it is possible that some apparent non-agents might be “locked-in” apparent agents. If they are, they must (in theory) be treated as agents. If they never display the capacities of apparent agency, this can never be ascertained. We have argued, elsewhere, however, that the PGC requires agents to guard for the possibility that they are actually agents, in proportion to the extent to which they approach being apparent agents.²⁸ This does not mean that they must be granted generic rights in proportion to the degree to which they approach apparent agency. Unless they are apparent agents, they cannot be granted these rights at all. Instead, it means that agents may not do things to them or fail to do things that would involve a violation of their generic rights if they are (purely hypothetically) agents without a generic rights justification (i.e., without showing that to restrain one’s actions in this way would involve a violation of at least some of one’s own generic rights). However, this will only be meaningful to the extent that it is possible to link the capacities of these non-apparent agents to the capacities of agents in relation to which they have generic rights. This raises a number of complex issues about which we have written elsewhere.²⁹ Here, it suffices to say that the precautionary prescription to play safe if this is possible and meaningful is, in

²⁸ See Deryck Beyleveld and Roger Brownsword, *Consent in the Law* (Oxford: Hart, 2007).

²⁹ See, e.g., Deryck Beyleveld and Shaun D. Pattinson op cit n. 26 supra, and Deryck Beyleveld and Roger Brownsword, n 28 supra.

principle, subject to scaling. This is to say that precaution requires one to play safe not only in a way proportionate to the conative (qualitative) dimension of degree of importance to avoid error, but also according to the quantitative degree to which it is possible and meaningful to play safe.

(ii) The PRPR and Regulators

Where regulators are faced with competing arguments as to the moral status of a being, how do the considerations just outlined bear on the articulation of a rational precautionary response?

On the basis of the other agents' argument, we propose that uses of precautionary reasoning generally will be rationally justified within a particular evaluative framework if they conform to the following elaborated principle, the PRPR:

If in total irresolvable doubt (where the choice is between Q and not-Q), play safe insofar as it is possible and meaningful to do so, to the extent that it is better to err on the side of the least unacceptable outcome.

But must the doubt be total and irresolvable before precautionary reason can validly apply? To ask this is essentially to ask whether the doubt must be total and irresolvable before the conative dimension of what evaluatively must be desired/undesired (hence, within uncertainty, hoped for/feared) is to be permitted to determine by itself what may or ought to be done.

The doubt must surely be total. If it is not then at least some estimate of degree of probability can be made and must be factored into the decision-making process. There are difficulties with how this is to be done, but, as we have already indicated, the disvalue of two, *ceteris paribus*, equally undesirable objects is affected by the probability of being confronted with them. Once one is factoring in degrees of uncertainty, the decision-making process becomes a standard (albeit complex) risk assessment or cost-benefit calculation.³⁰

So must the doubt also be irresolvable? Being totally, but not irresolvably, uncertain means being only contingently, not necessarily, totally uncertain. Things can be done to reduce the degree of uncertainty or even to pinpoint it precisely. However, for so long as it is total, precautionary reasoning still applies. The difference that resolvability makes is that precautionary reasoning cannot, without more, lead to a total time-unlimited prohibition on an activity. A moratorium rather than a ban is indicated, the specific features of which must be assessed.

Further to our discussion in Part II of the paper, it is sometimes said that the risks in a precautionary argument must not merely be logically possible, they must at least be "real" (meaning based on empirical evidence).³¹ This is too sweeping. As we have

³⁰ For our understanding of risk analysis, see note 8 above.

³¹ See, e.g., Fritz Allhoff, "Risk, Precaution, and Emerging Technologies" (2009) 3 *Studies in Ethics, Law and Technology* Art 2 at [18]

already seen, in our other agents problem, logical possibility is enough because no other kind of possibility even has any application. Whether or not an apparent agent is an agent is not something amenable to empirical investigation at all in a non-speculative manner. What is important is that mere logical possibility is insufficient if a kind of possibility that applies less widely is applicable.

Thus far, our discussion of the moral track to the PRPR has centred on a relatively easy case, one in which, by employing precautionary reasoning to treat B as an agent, there is no violation of the categorical imperative, there is no moral loss. In other words, this is just the kind of case which, in the typical regulatory context, the critics of the Precautionary Principle insist cannot be generalised. Typically, the critics insist, there is a price to be paid for precaution and it is irrational not to take this into account. What, then, does the PRPR say about the kind of case where precaution comes at a price, where there is a moral cost in treating B as an agent? For example, what if treating B as an agent means that the moral interests of C, an ostensible agent, are less well served or even harmed? If the moral version of the PRPR is to engage with such a case, there seem to be two requirements, namely: (i) that there is extreme uncertainty as to the status of B (B is a possible agent); and (ii) that failure to treat an (actual) agent as an agent is a fundamental moral harm. Where these requirements are met, regulators must proceed on the basis that B is to be treated as an agent, without having established the likelihood of B being an agent. To this extent, the impact on C of treating B as an agent does not weigh against treating B as an agent; but, of course, once B is treated as an agent, the importance of B's particular moral interests has to be judged against whatever moral interests C has. In some cases, B's interests might prevail; in others, it will be C's interests that prevail. However, whatever the regulatory outcome, it will be the result of a comparative assessment of the interests of beings who, for the purposes of the assessment, are treated as agents.³²

Needless to say, there will be many more hard cases for regulators as they strive to apply their best interpretation of the community's moral principles in a way that guards against making the worst kind of error (that is, the error which, from a moral standpoint, is the one to be avoided). In the final part of our paper, one of the test-cases (neuroscience and novel neurotechnologies in the criminal justice system) will give a further indication of how regulators might employ precautionary reasoning where they find themselves on the horns of several moral dilemmas.

V Three Test-Cases for the PRPR

So far, our discussion has been highly schematic: our purpose has been to show that, in conditions of extreme uncertainty, there is still scope for rational precautionary reasoning. It might be, as we have already hinted, that there are relatively few occasions when the PRPR is directly relevant. Accordingly to indicate more clearly the practical applicability and application of the PRPR, we now consider three test-cases of uncertainty. These three cases are, respectively, particle accelerators such as the Large Hadron Collider (LHC), nanofoods (and nanomaterials in food packaging), and the use of neuroscience and novel neurotechnologies (particularly fMRI brain imaging technologies) in the criminal justice system.

³² Arguably, if the moral interests of B and C are identical, then the fact that C is an ostensible agent might tip the balance in C's favour.

(i) Particle accelerators and the Large Hadron Collider

On the day that the Large Hadron Collider (LHC)³³ was ready to function, Europeans held their collective breath. Some maintained that the LHC should not have been built, because, for all we know, its use might destroy the world. This was (and still is) logically possible. It is also logically possible that it will not. Was this a case for regulators to let us sleep easy by applying the PRPR?

The rational starting point for regulators would be to consider (i) their options (put simply, to intervene or not to intervene) and (ii) the worst consequences of being wrong about the supposed harm that the LHC might cause relative to their options. If regulators, fearing that the LHC might destroy the world, opt to intervene, and if their fear is misplaced, what would be lost? To which the answer is: the benefits of much improved knowledge of fundamental physics that use of the collider is likely to give (if it does not destroy the world). On the other hand, if regulators, hoping that the LHC will not destroy the world, do not intervene, and if they are wrong, what would be lost? To which the answer is: everything that matters (though some fanatics might take a different view)! Ergo, because it is (categorically) worse to err having supposed that the LHC will not destroy the world (and going ahead with the LHC's use) than to err having supposed that it will destroy the world (and preventing its use) regulators should intervene to prohibit the construction and use of the LHC.

If this is what the PRPR prescribes for the LHC, some will argue that this is unreasonable.³⁴ Let us suppose that the argument is (i) that there is no good scientific reason to suggest that the LHC will destroy the world; (ii) so it is extremely unlikely that the LHC will destroy the world;³⁵ and (iii) it is therefore unreasonable to categorically ban the use of the LHC on the precautionary model we are advocating. Consequently our precautionary reasoning is invalid.

Not so! If we accept the idea that the scientific evidence shows that it is very unlikely that the LHC will destroy the world, then we are not in a position of total irresolvable doubt at all (which is where logical possibility is the appropriate standard). Indeed, we

³³ The LHC is a facility over one hundred metres below the Alps built to, amongst other things, simulate conditions shortly after the "Big Bang" that scientists believe marks the beginning of the known universe. Use of the LHC involves causing protons accelerated to close to the speed of light colliding with each other. It is hoped that the experiments will answer some important questions in fundamental physics. The LHC began operation in 2009 and full power operations were deployed in November 2010. Thus far fears about its use (e.g., that it might cause mini-black holes to be created) have not been realised. So, there is now some inductive evidence that the LHC will not have the catastrophic effects feared. This paper, however, considers the position before that evidence became available.

³⁴ Compare the discussion of precautionary responses to remote risks in Fritz Allhoff, Patrick Lin, and Daniel Moore, *What is Nanotechnology and Why Does it Matter?* (Chichester: Wiley-Blackwell, 2010).

³⁵ The move from "no good reason to think that it will destroy the world" to "it is very unlikely that it will destroy the planet" when discussing just this example (though not the precautionary reasoning we apply to it) is made by Fritz Allhoff, note 31 above, at 18. The move is obviously unsound but we will let it pass for the moment.

are not even in a state of total empirical (resolvable) doubt (where empirical possibility is the appropriate standard). The objection tells us nothing about the validity of our model of rational precautionary reasoning, because that reasoning, according to the principle itself, does not apply when there is enough information about the less than total degree of uncertainty to apply a standard risk analysis or cost-benefit analysis.

If the context for making a regulatory decision about the LHC is one of extreme uncertainty (or total irresolvable doubt), and if the destruction of the world is a fundamental harm (and, even allowing for fanatics, responsible regulators must surely treat it as such), then this is a paradigmatic case for the application of the PRPR. If this seems unreasonable and irrational, it is not because the PRPR is unreasonable or irrational but because it is not accepted either (i) that the context is one of extreme uncertainty or (ii) that unnecessarily restricting research into the Higgs Boson particle, or some such, is a worse option than erroneously permitting a “strangelet” disaster that reduces the planet to a tiny inert hyperdense sphere.

(ii) Nanofoods

In its report on *Nanotechnologies and Food*,³⁶ the House of Lords Science and Technology Committee recommended that regulators should take a precautionary approach to the authorisation of nanofoods for human consumption. The Committee noted the potential benefits that nanotechnologies might realise in the food sector, but it was particularly concerned about the risks presented by nanoparticles that lodge in the human gut. Thus, in the Summary, we read:

Nanomaterials have a range of potential applications in the food sector that may offer benefits to both consumers and industry. These include creating foods with unaltered taste but lower fat, salt or sugar levels, or improved packaging that keeps food fresher for longer or tells consumers if the food inside is spoiled. At present the number of food products that contain nanomaterials is small, but this may well change over the next five years or so as the technology develops. For these reasons, we make a series of recommendations that are intended to support the responsible development of nanotechnologies in the food sector and to ensure that potential benefits to consumers and society are supported, where appropriate, by Government.

Nanotechnologies may also present new risks, as a result of their novel properties, as well as potential benefits to consumers. There are a wide variety of nanomaterials, and while many types of nanomaterials may well prove to be harmless, others may present a higher risk. Our current understanding of how they behave in the human body is not yet advanced enough to predict with any certainty what kind of impact specific nanomaterials may have on human health. Persistent nanomaterials are of particular concern, since they do not break down in the stomach and may have the potential to leave the gut, travel throughout the body, and accumulate in cells with long-term effects that cannot yet be determined.

³⁶

1st Report of Session 2009-2010, HL Paper 22-I, January 8, 2010.

Although there has been some research on the way in which nanotechnologies might impact on human health, the Committee emphasised that “knowledge of the risks associated with the use of nanomaterials is incomplete. Significant gaps remain.”³⁷ One of the most worrying gaps concerns the impact of nanomaterials in the human gut. As the Committee put it:

4.21. To date, little research has been undertaken into the impact, behaviour and interaction of nanomaterials in the GI tract, including their effect on natural gut flora. In contrast, a significant amount of research has taken place into the effects of nanomaterials on the lung—according to Dr Knowles, “most research has been, and continues to be, on inhalation” (Q 170). But this work may not assist in understanding the effect of nanomaterials that enter the body through ingestion because, as Professor Donaldson told us, you cannot generalise from the effects of particles in the lungs or on the skin to the effects in the gut: “The gut is a wholly different environment to me to these other situations in terms of the extremity of conditions” (Q 215).

4.22. It appears that a great deal of work still needs to be done on the effect of nanomaterials in the gut. Dr Powell, for example, said: “more work needs to be done in terms of both nanoparticles and the larger nanoparticles or microparticles, those larger than 100nm in diameter, in terms of what happens inside the gut” (Q 215). Professor Depledge argued: “the amount of evidence available with regard to the effects of nanomaterials, delivered through food or in food, is very, very small indeed and there is an urgent need to conduct more studies” (Q 215). Other witnesses agreed (QQ 123, 232, 256). The EFSA stated: “the understanding of the potential toxicity after oral intake of ENMs is in its infancy. Only a very limited number of ENMs have been studied after oral administration ... The ENMs used in the toxicity studies were often characterised only to a very limited extent”.[\[25\]](#)

Given this uncertainty, the Committee recommended that any regulatory gaps should be closed, that government should take steps actively to increase our understanding of the risk profile of nanotechnologies, that cooperative data sharing should be encouraged during the pre-competitive period, and that consumers should be properly engaged and informed with regard to the use of nanotechnologies in the food sector. More particularly, the effect of the Committee’s recommendations 5 and 6 is that the Research Councils should establish more pro-active forms of funding to encourage the submission of research bids to address the severe shortfalls in research required for risk assessment of nanomaterials, particularly relating to the harm caused by such materials that lodge in the gut; and recommendation 9 is that the Government should work more closely with European and international partners on research related to the health and safety risks of nanomaterials to ensure that knowledge gaps are filled quickly and without duplication.

From a regulatory perspective, although the Committee does not recommend special labelling requirements for nanofoods, it does recommend that the Government should work within the European Union to promote the amendment of current legislation to ensure that all nanomaterials used in food products, additives, or supplements fall

³⁷

Para. 4.15.

with the scope of current legislation (recommendation 11). This is modestly precautionary but would it pass muster as a rational response?

On one reading of the expert evidence, we might say that, in some respects, so little is known about the health risks associated with nanofoods that we are in a condition of extreme uncertainty. If so, applying the PRPR, the question is whether regulators should proceed on the assumption that nanofoods will not cause harm to human health, knowing that they might be wrong, or proceed on the assumption that nanofoods might present such risks, knowing that whatever intervention they make might be unnecessary. Clearly, putting the choice this way, the Committee finds the former option the less acceptable; and, having taken the latter option, the interventions that it proposes are not likely to be seen as disproportionate. After all, it is not as though the Committee is recommending that there should be a prohibition of the development of nanofoods.

When the gaps in our knowledge and understanding of nanomaterials are serious but do not amount to extreme uncertainty, the Committee's recommendation that nanomaterials used in food products, additives, or supplements should be put through the usual safety checks for novel foods seems unproblematic. In other words, if we can conduct a standard risk assessment on nanofoods, the Committee is saying what any responsible regulator surely would accept, namely that they should be subject to the usual expert assessment.³⁸

(iii) Neuroscience, Neurotechnologies, and the Criminal Justice System

With the development of fMRI scanning technologies, neuroscientists have been able to offer new insights into the operation of the brain. Some believe that, in due course, neuroscience will so transform our understanding of the relationships between brain, mind, and behaviour that we will see the notion of treating humans as morally responsible for their actions as inappropriate. This implies a fundamental rethink for many of our social practices (for instance, simple praising and blaming), not least for criminal justice systems insofar as they are predicated on holding convicted defendants as morally accountable and as appropriately punished (and stigmatised) for their crimes.³⁹ Against this view, however, the headline message in a recent report from the Royal Society is that the translation of neuroscientific findings into the practice of law should be treated with great caution⁴⁰; and, in several papers,

³⁸ This, however, is not the same as saying that the risk-assessment processes are already geared for nanotechnologies. For doubts about the preparedness of regulatory regimes in relation to nanotechnologies, see, e.g., Albert C. Lin, "Size Matters: Regulating Nanotechnology" (2007) 31 *Harvard Environmental Law Review* 349, at 361-374 (for the view that US regulatory provisions are inadequate); Giorgia Guerra, "European Regulatory Issues in Nanomedicine" (2008) 2 *Nanoethics* 87 (for the view that EC regulation does not fit very well with potential nanomedical applications); and, generally, Trudy A Phelps, "The European Approach to Nanoregulation" in Nigel M de S Cameron and M Ellen Mitchell (eds), *Nanoscale* (Hoboken, NJ: Wiley, 2007) 189.

³⁹ Joshua Greene and Jonathan Cohen, "For the Law, Neuroscience Changes Nothing and Everything" (2004) 359 *Philosophical Transactions of the Royal Society B: Biological Sciences* 1775.

⁴⁰ Royal Society, *Brain Waves Module IV: Neuroscience and the Law* (London: December 2011).

including one on this journal, Stephen Morse has entered a plea for “neuromodesty”.⁴¹ If we follow these latter reservations, we will take a hard look at the claims made by neuroscientists—or, more accurately perhaps, we will take a hard look at the scenarios sketched by futurologists who draw on neuroscience to anticipate a world in which we can identify “genes for criminality”, use fMRI technology for lie detection, scan one another’s minds to reveal our innermost thoughts, and so on. As for giving up on the criminal law, Morse has repeatedly argued that, unless neuroscience can show that we are automatons or lack practical rationality, there is no reason why it should fundamentally destabilise current practices in the criminal justice system.⁴²

Nevertheless, even if neuroscience is not about to provoke a radical rethink of criminal justice, there are many questions of policy and principle on which it might impact.⁴³ For example, does it make sense to reject the kind of defence presented by Stephen Mobley (to the effect that his MOMA activity was dysfunctional)⁴⁴ and yet allow neuroscientific evidence to save teenage killers from the death penalty⁴⁵ or even excuse acts in some cases⁴⁶? Or, if Bert-Jaap Koops is correct in seeing a movement towards “the crime society”,⁴⁷ where individuals are increasingly profiled, how far should neuroscience be employed in assessing the risk presented by individuals (i) who have not yet committed a crime and (ii) who are awaiting release from custodial sentences?

One of the more urgent questions flagged up in the Royal Society report concerns the age of criminal responsibility. Historically, the English common law took a relatively

⁴¹ Stephen J. Morse “Avoiding Irrational Neurolaw Exuberance: A Plea for Neuromodesty” (2011) 3 *Law, Innovation and Technology* 209.

⁴² See, e.g., Stephen J. Morse, “Uncontrollable Urges and Irrational People” (2002) 88 *Virginia Law Review* 1025; “Moral and Legal Responsibility and the New Neuroscience” in Judy Illes (ed), *Neuroethics* (Oxford: Oxford University Press, 2006) 33; and “Lost in Translation? An Essay on Law and Neuroscience” in Michael Freeman (ed), *Law and Neuroscience* (Oxford: Oxford University Press, 2011) 529.

⁴³ See, e.g., Henry T. Greely, “The Social Effects of Advances in Neuroscience: Legal Problems, Legal Perspectives” in Judy Illes (ed), *Neuroethics* (Oxford: Oxford University Press, 2006) 245; and “Law and the Revolution in Neuroscience: An Early Look at the Field” (2009) 42 *Akron Law Review* 687 (where the focus is on prediction, mind reading, responsibility, treatment, and enhancement); and Jeffrey Rosen, “The Brain on the Stand” *The New York Times*, March 11, 2007.

⁴⁴ For discussion, see Eliezer J. Sternberg, *My Brain Made Me Do It* (New York: Prometheus Books, 2010) Ch. 1.

⁴⁵ The highest profile example is *Roper v Simmons* 543 US 551 (2005); see, for an impressive survey and analysis, O. Carter Snead, “Neuroimaging and the ‘Complexity’ of Capital Punishment” (2007) 82 *New York University Law Review* 1265.

⁴⁶ Compare, the well-known case of Oft (the teacher who developed paedophilic tendencies): for discussion, see Stephen J. Morse, “Lost in Translation? An Essay on Law and Neuroscience”, note 42 above, at 559-56.

⁴⁷ Bert-Jaap Koops, “Technology and the Crime Society: Rethinking Legal Protection” (2009) 1 *Law, Innovation and Technology* 93.

nuanced approach to whether children were sufficiently mature to be held criminally responsible for their acts. Up to the age of 7, children were regarded as too young to be held responsible; at 14, they were treated as responsible; and, in relation to children from 7 up to 14, there was a rebuttable presumption that they did not understand the difference between right and wrong and, thus, they were not yet responsible (this was the so-called defence of *doli incapax*). In the Twentieth Century, the minimum age was raised, first to 8 and then to 10, leaving the *doli incapax* defence available for children of 10 or over but not yet 14. However, the defence fell into disrepute—the media highlighting some notorious cases in which youngsters under the age of 14 seemed all too knowing about their criminality—and it was abolished by section 34 of the Crime and Disorder Act, 1998. In *R v JTB*⁴⁸, a unanimous House of Lords removed any doubt that the legislative intent was to abolish the defence rather than simply to reverse the burden of proof (so that it would be for the defence to prove incapacity rather than for the prosecution to prove capacity). This leaves English law with a bright line rule for criminal responsibility; but, in setting the threshold at 10, the English position is at variance with many other European countries (including Scotland) where the age of criminal responsibility is significantly higher and where a much less punitive culture prevails in relation to young offenders. The question now is whether the English position can be defended in the light of a broad consensus amongst neuroscientists that, by the age of 10, a child's brain has not only a long way to go before it is mature but also has yet to develop the mechanisms (particularly in the pre-frontal cortex) that restrain impulsivity.

The question just posed is a variation on the other agents problem that we have already discussed. In the classic version of the problem, precautionary reasoning suggests that, where a being behaves as if it is an agent (as an apparent agent), it should be treated as an agent (even though this might be to make an error); in the present version, the question is whether precautionary reasoning might indicate that we should treat a youthful apparent agent as though he or she were not yet a fully responsible agent. If regulators were to bring the PRPR to bear on this particular question, their reasoning should run along the following lines:

- (i) For the most part, 10 year olds behave in ways that are consistent with their being agents; *prima facie*, they are apparent agents.
- (ii) However, there is a broad consensus amongst neuroscientists that, at 10, the levers of agency are a long way from being developed and, crucially, that the usual brain-based restraints on impulsivity are not yet developed.
- (iii) In setting the age of criminal responsibility at 10, we (regulators) have been guided by the perception that children of that age seem to understand what they are doing and that, where the conduct is a crime, to appreciate that what they are doing is wrong.
- (iv) In the light of the evidence from the neuroscientists, are we making a moral mistake in treating 10 year olds as criminally responsible?

⁴⁸

[2009] UKHL 20.

(v) If the evidence from neuroscientists is that children develop at different speeds and that we cannot be sure about the state of development in any particular 10 year old, then the most we can say is that it is possible that a particular child is sufficiently developed (but also possible that it is not). (For the purposes of this illustration, we will set aside an alternative scenario where neuroscientists can advise regulators as to the percentage of 10 year olds who are mature in the way that is required for criminal responsibility).

(vi) On the basis of the neuroscientific evidence, it must be recognised that mistakes might be made—notably that children who should not be held criminally responsible are held responsible. However, if the age of criminal responsibility is raised, we might also make mistakes—notably that children who should be held criminally responsible are not held responsible.

(vii) If there are to be errors, which is the worse error to make? Is it worse to treat 10 year olds as mature and to hold a child who is not actually mature as criminally responsible; or, to treat 10 year olds as immature and not hold a child criminally responsible even though the child is sufficiently mature? It is a hard case. However, for communities that are strongly committed to the principle that the innocent should not be convicted of criminal offences, the background principles might well indicate that the worse moral error is wrongly to hold to account (as criminally responsible) those whose brains are not yet sufficiently developed.

Having got to his point, regulators might wish to revisit the question of whether neuroscientific examination of particular 10 year old defendants might throw light on the state of their brain development. If so, the now repealed principle of *doli incapax* (as a rebuttable presumption either for or against responsibility) might seem like a suitably precautionary response. Or, regulators might decide to follow the example of their counterparts elsewhere and simply raise the age of criminal responsibility.

As we have said, in some communities, background principles will treat it as a particularly grave moral wrong to convict an innocent person of a crime. This plays out in the detailed design of the criminal justice system so that, at every point, the imperatives of crime control are set against due process values. In such a context, neuroscientific evidence will be admitted at a trial only where it can cross a high threshold of reliability⁴⁹—so, we are unlikely to see evidence based on fMRI lie detectors in criminal trials for a very long time (although, worryingly, their use elsewhere in the criminal justice system might be less stringently regulated). At all events, for a community with such moral commitments, one of the golden threads of the criminal justice system is the legal doctrine that in a criminal trial the defendant must be shown to be guilty beyond a reasonable doubt (while in a civil action the defendant need only be shown to be guilty on the balance of the probabilities). Might such a doctrine might be justified relative to the PRPR?

In response to this question, we can start by noting that it is rarely, if ever, possible to be absolutely certain that a defendant D is guilty of an offence. And certainly, this is not something that can be judged before presentation of the evidence relating to any

⁴⁹ In England, the Law Commission has quite recently recommended a more stringent approach to the admissibility of expert scientific evidence: see The Law Commission, *Expert Evidence in Criminal Proceedings in England and Wales* (LAW COM No. 325, March 21, 2011).

alleged offence. So, in general, there is room for error when judging guilt on the basis of the evidence that will be presented in court. An innocent person might be convicted, while a guilty person might be acquitted. The situation here might seem to put the scenario outside the ambit of the PRPR because evidence can definitely be brought to bear on the question of guilt or innocence. It is perfectly meaningful and applicable to talk about evidence establishing guilt or innocence to a degree of probability. This, however, is not the focus of the problem. The problem here is not whether it is possible to ascertain the probability that D is guilty or not. The problem, rather, is what level of probability needs to be established in order to return a verdict of “Guilty!” as against a verdict of “Not guilty!” This is not a matter of probability. We are in total doubt about this and this doubt is one that cannot be resolved by applicable cognitive information. This immediately puts the problem into the frame of the PRPR.

We have a choice between being in error if we convict an innocent person (returning “guilty” (Q) and being wrong) and being in error if we acquit a guilty one (returning “not guilty” (not-Q) and being wrong). Which is the worse error, how bad is it, and how much worse is it than its contradictory? If the two errors are judged to be equally bad, then the standard of proof should be to show that D is guilty on the balance of the probabilities (to show that it is merely more likely that D is guilty than that D is not guilty). That the test is that D must be shown to be guilty beyond a reasonable doubt indicates that erring on the side of Q is being judged to be much worse than erring on the side of not-Q. It is only if erring on the side of Q is judged to be much worse than erring on the side of not-Q that such an imbalanced standard of proof can be justified. At the same time, erring on the side of Q is not judged to be categorically bad, because then the burden should be to show that D is guilty beyond any possible doubt (“zero risk” really would be the order of the day). Note that erring on the side of Q would be judged to be worse than erring on the side of not-Q if the standard of proof were anything greater than showing that D is more likely to be guilty than not (e.g., “more likely to be guilty than not by a fair margin”, “more likely to be guilty than not by a handsome margin”, “very much more likely to be guilty than not”, “almost certainly guilty”, etc., would all qualify). So what does “guilty beyond a reasonable doubt” imply? Well, acquitting a guilty person is not something that we should be happy about. We should be very unhappy about it. This is a very serious matter. For that reason, there needs to be some good reason for an acquittal. Placing the burden at “almost certainly guilty” looks too high. Placing it at “very much more likely to be guilty than not”, however, suggests that there is no real qualitative difference between the two errors. It is just a matter of quantity. Placing it at “guilty beyond a reasonable doubt” (which we interpret to mean “there remains at least one scenario under which D is not guilty which the prosecution has not been able to render so implausible that a reasonable person could not entertain it”) seems about right if one thinks that there is a qualitative difference but at the same time the alternative error is serious.⁵⁰

⁵⁰ We should not discount the possibility that the burden and the level of justification required in a criminal trial are an attempt to guard against wrongful convictions that are brought about by the imbalance of power and resources as between the state and the accused. This does not mean that a precautionary argument of the kind that we have outlined is not sound given the evaluative assumptions that we have made.

The general lesson we draw from this is that, while it is true that when the doubt is not total the PRPR cannot be used to make the decision, even if we know or can judge accurately what the probabilities are, it can be used to interpret the significance of a level of probability.⁵¹

VI Concluding Remarks

In this paper, we have focused on the special context of extreme uncertainty, a context in which we claim that it is rational and responsible for regulators to apply precautionary reasoning. The reasoning so applied, the PRPR, is a methodological principle, which lays down rules that substantive risk management policies, rules, principles or actions in specific contexts must satisfy if they are to be reasonable. By contrast, the PP is definitely a risk-averse policy (in typical articulations, imposing an asymmetrical burden and standard of proof). As such, the PP stands to be judged by the PRPR, rather than being something that is deducible from the PRPR. What this means is that it must be interpreted and applied in a manner consistent with the PRPR. If it is not capable of being so interpreted and applied it must be rejected as unreasonable.

While conditions of extreme uncertainty might be exceptional, precautionary reasoning is an everyday feature of regulatory decision-making—witness the constant appeals to acceptable risk, the concerns (for example, in relation to biometric technologies) about false positives and false negatives, and of course the pressing of the precautionary principle itself. If precaution begins where straightforward risk/benefit calculations run out, and if the PRPR is to govern all aspects of precautionary reasoning in relation to emerging technologies, then the central question will be whether it is better or worse (and to what degree) to err in one way (involving an unnecessary restriction on a technology that promises some benefits) rather than another (involving a failure to prevent a harmful effect resulting from the development and use of the technology). A question of this form will also be central to precautionary reasoning that we find elsewhere in the law—for example, in setting the groundrules for pre-trial criminal process, for criminal trials, for reviews of allegedly mistaken convictions, for decisions about the release of prisoners back into the community, and so on. In all these criminal justice cases, the dilemma is whether it is worse to err by curtailing the freedom of the innocent or by not curtailing the freedom of the guilty.

Before we can answer these precautionary questions, we need to have a yardstick for judging what is better or worse. Where the frame for regulatory judgment is prudential, the yardstick is the interests and preferences of those regulatees for whom regulators act as proxies. However, this frame is always subject to the impingement of overriding moral values and great needs to be taken to ensure that the pronouncement that a technology is “safe” does not foreclose or collateralise discussion of moral concerns.⁵² Where the frame is moral, a particular ethical yardstick needs to be

⁵¹ In Deryck Beyleveld and Roger Brownsword, op cit, note 28 above, we said that all evidential requirements that place different standards of proof on different sides of a case are viewable functions of precautionary reasoning.

⁵² See Roger Brownsword, “Human Dignity and Nanotechnologies: Two Frames, Many Ethics” (2011) 19 *Jahrbuch für Recht und Ethik* 429.

specified; and, as we have indicated in the paper, particularly when we discussed the problem of other agents, we will call into service Gewirth's PGC. However, a detailed application of the PGC to a range of precautionary puzzles is a task for another occasion.⁵³

In spelling out the PRPR we are conscious that, even at a merely formal level, what we have said in this paper is only an outline. We suspect that fine adjustments and qualifications might need to be made when we attempt to apply the PRPR to a large number of case studies. However, we will have succeeded in our present aim if what we have identified as the PRPR is the analytic principle we claim it to be, if it will enable us to recognise cases of precautionary reasoning when we come across them, and if it enables us to identify the key questions that need to be asked to assess the rationality of precautionary interventions (or non-interventions) taken by regulators as they face a succession of emerging technologies and novel applications.

⁵³ As we have indicated elsewhere, see op cit note 27 above, we are treating the PP (and now the PRPR) as pieces in a much larger mosaic of precautionary reasoning in the law. See, too, Roger Brownsword, "Nanoethics: Old Wine, New Bottles?" (2009) 32 *European Journal of Consumer Policy* 355.