Published in final edited form as:

Risk Anal. 2017 May ; 37(5): 893–904. doi:10.1111/risa.12658.

Underprotection of unpredictable statistical lives compared to predictable ones

Marc Lipsitch,

Center for Communicable Disease Dynamics, Department of Epidemiology and Department of Immunology and Infectious Diseases, Harvard T.H. Chan School of Public Health, 677 Huntington Avenue, Boston, MA 02115

Nicholas G. Evans, and

Department of Medical Ethics and Health Policy in the Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA

Owen Cotton-Barratt

Future of Humanity Institute, Oxford Martin School, University of Oxford

Abstract

Existing ethical discussion considers the differences in care for identified versus statistical lives. However there has been little attention to the different degrees of care that are taken for different kinds of statistical lives. Here we argue that for a given number of statistical lives at stake, there will sometimes be different, and usually greater care taken to protect *predictable* statistical lives, in which the number of lives that will be lost can be predicted fairly accurately, than for *unpredictable* statistical lives, where the lives are at stake because of a low-probability event, such that most likely no one will be affected by the decision but with low probability some lives will be at stake. One reason for this difference is the statistical challenge of estimating low probabilities, and in particular the tendency of common approaches to underestimate these probabilities. Another is the existence of rational incentives to treat unpredictable risks as if the probabilities were lower than they are. Some of these factors apply outside the pure economic context, to institutions, individuals, and governments. We argue that there is no ethical reason to treat unpredictable statistical lives differently from predictable statistical lives. Moreover, lives that are unpredictable from the perspective of an individual agent may become predictable when aggregated to the level of a societal decision. Underprotection of unpredictable statistical lives is a form of market failure that may need to be corrected by altering regulation, introducing compulsory liability insurance, or other social policies.

1. INTRODUCTION

An ongoing ethical debate concerns whether it is justifiable to take more care to protect identified lives than to protect statistical lives.⁽¹⁾ For an agent facing a decision, identified lives that will be lost/saved by the decision are those of individuals whose identity is known to the agent, while statistical lives are lives of individuals whose identities are unknown to

the agent, but will be lost/saved by that agent's decision. A canonical treatment of the distinction is given by Thomas Schelling: "Let a 6-year-old girl with brown hair need thousands of dollars for an operation that will prolong her life until Christmas, and the post office will be swamped with nickels and dimes to save her. But let it be reported that without a sales tax the hospital facilities in Massachusetts will deteriorate and cause a barely perceptible increase in preventable deaths—not many will drop a tear or reach for their checkbook."⁽²⁾ In many situations, people are less inclined to bear a particular cost or exert a particular effort to protect statistical lives than to protect the same number of identified lives. Hereafter, when we speak of care taken to protect lives, we mean the amount of money or effort an agent is willing to expend to prevent a particular threat to those lives.

In economics, Schelling's work established the initially controversial proposition that the value of life could be quantified, leading ultimately to a number of refinements about how best to quantify the value of life and craft policies that used such valuations. In ethics, however, the debate about identifiable versus statistical lives concerns the justifiability of this differential care. Here we draw attention to a further, unexplored instance of differential care taken to protect two kinds of statistical lives. We claim there are reasons to expect agents to take a lower level of care for what we call unpredictable statistical lives—statistical lives whose loss is probabilistically very small—relative to the same expected number of predictable statistical lives, in much the same way as there is a lower level of care for statistical lives *simpliciter* relative to the same number of identified lives. There is thus an analogous ethical debate to be had on whether this is justified; moreover, we claim that regardless of whether individuals may be ethically justified in some cases in taking reduced care for unpredictable lives, there may be sound reasons for social policy to discourage such behavior.

We begin by defining predictable and unpredictable statistical lives. We then describe two different reasons why agents may take lesser care for unpredictable statistical lives: (1) difficulties in estimating the probability of rare events; and (2) rational incentives to reduce care for unpredictable statistical lives, relative to the care an agent would take for equivalent, predictable statistical lives. We make two lines of argument about the consequences of these differential levels of care. First, we suggest that reduced care for unpredictable statistical lives is ethically unjustified. Second, we argue that regardless of whether an individual agent can ethically justify taking a lower level of care for unpredictable than predictable statistical lives, society has a legitimate interest in discouraging this reduced care.

2. DEFINITIONS: UNPREDICTABLE VERSUS PREDICTABLE STATISTICAL LIVES

For an agent facing a decision that could affect statistical lives, let us distinguish between two cases. The agent faces a decision about *unpredictable statistical lives* in the case where unless the agent acts in a certain way there is a low probability $p < 50\%$ that all of these individuals' lives will be lost, and with the remaining probability $1-p$ their lives will be unaffected by the decision. A case of unpredictable statistical lives at stake would be the bystanders who die in an explosion at a factory, and the decision for the firm that owns the

factory is whether to install a safety system that reduces the probability of such an explosion (we call the decision to install such a system “mitigation.”) Most likely the number of lives lost will be 0, regardless of the firm’s decision, because an accident is improbable with or without mitigation. However, without mitigation there may be an explosion with L lives lost. The key point for unpredictable statistical lives is that, more likely than not, the number of lives affected by the decision is 0 (because $p < 50\%$), while the expected number of lives affected is pL , which could be large. The effect of the decision on the fate of the L people is unpredictable in the sense that either it will kill all of them or it will affect none of them. To simplify exposition, we make the following assumptions, none of them necessary for our argument: (i) we assume that the number of lives to be lost as a result of an explosion is fixed at L , rather than having some uncertainty in magnitude; (ii) we consider only death and no other harms of the explosion; (iii) we consider completely effective mitigation, which eliminates the risk that the L lives will be lost.

In contrast, many other kinds of decisions involve *predictable* statistical lives. Preventing the daily release of a highly toxic effluent from a factory will reduce the risk of death to those in the neighborhood of the factory, and although the exact number of lives to be saved cannot be predicted, it is safe to say that (if one has an adequate understanding of the toxicity and exposure of the population) the actual number of deaths prevented will be comparable in magnitude to the expected number predicted by an appropriate statistical model. This occurs because, in the case of toxic exposures, the number of people exposed is (approximately) fixed, and each person’s probability of dying given that they are exposed, is (approximately) constant and independent of whether the others die from the exposure. The law of large numbers – which applies to large numbers of independent events – makes it very likely that the number actually affected will be close to the expected or average number affected.

Our definition of *unpredictable* has two particular features. First, the predictability of a set of statistical lives is defined from the perspective of a particular *agent* whose action will affect whether the lives are lost or not. For most of this paper that agent will be a firm, which either will or will not have an accident at the factory it runs. As we note later, another agent that could make a decision affecting these lives is the national government, for example through legally requiring safety measures at all factories of a certain type, including the one belonging to this firm. In such a case, the lives that are unpredictable from the firm’s perspective may become more predictable from the government’s perspective, because across a whole country the expected number of accidents may approach or even exceed 1. The law of large numbers may apply at larger scales of aggregation, such as a country, even when it does not to an individual firm. In this situation, the most likely outcome is that there will be an accident in *some* factory during the year, even if the most likely outcome for any *individual* factory is no accident. We explore the consequences of this difference in perspective below.

Second, unpredictability in our sense does not require that the probability the agent’s decision will affect lives is unknown, only that it is low. It may be universally known and agreed by all parties that the probability of an accident at a factory in any given year without mitigation is 1%. We believe this situation is unlikely in practice, but we emphasize that the key point of *unpredictability* as we use it is that the probability of the harmful event is low.

3. REASONS WHY UNPREDICTABLE STATISTICAL LIVES MIGHT RECEIVE DIFFERENT LEVELS OF CARE FROM PREDICTABLE ONES

The degree of care taken to protect unpredictable statistical lives will depend on the capacity of agents to estimate accurately the risks of low-probability events, and on their incentives to act on these estimates. Therefore a systematic tendency to underestimate low probabilities, or rational incentives to act as if the probabilities associated with harming unpredictable statistical lives were lower than an agent's own best estimate of these probabilities, could induce lower levels of care for unpredictable statistical lives. We argue that both are often operative.

3.1. Tendency to underestimate low probabilities

When an event happens rarely, it is hard to estimate the probability that such an event will occur a defined time period. This is intuitively clear, because there will typically be small amounts of data available for rare events (with the exception of events, such as earthquakes, for which long-timescale geologic or written records exist), and there may be legitimate uncertainty about the relevance of the data that do exist. In the case of a factory explosion, well-informed experts may differ on the question of whether the history of such explosions can be used to estimate a probability for an explosion in a particular factory, which may differ from those in the historical record in many ways, including design, maintenance, staffing and the like. While some might argue that an estimate of the probability of explosion in the factory of interest should be based on the rate of explosions in all factories in the country in question over the last decade, others might argue that only the record of the last three years for factories built by the same contractor should be relevant. This tradeoff of direct relevance against sample-size may, in the extreme, lead to a shrinking of the relevant historical record to include only very few factories, at which point we cannot trust the law of large numbers to ensure that the observed rate is close to the true rate – indeed it may be that none has experienced an explosion.

The estimation of probabilities from events that have never happened raises particular problems, which we discuss below. For now, we note that disagreements about the relevant historical experience may lead either to overestimates or underestimates of the probability of a rare event. A recent such controversy that exemplifies this problem is a debate over the probability of an accidental influenza pandemic by experiments to enhance the transmissibility of avian influenza viruses: critics have estimated the probability at around 1 in 1000 to 1 in 10,000 for a single year of research in a single laboratory⁽³⁾, while one of the scientists who performs such experiments argues that the true probability is 1 in 33 billion⁽⁴⁾. This figure has been disputed by those who provided the original estimate⁽⁵⁾ and by another commentator.⁽⁶⁾ At least one of these estimates must be far from correct.

While disagreement about data sources – as well as a number of other cognitive biases we discuss below – may lead to errors in either direction in estimating rare-event probabilities, several factors specifically tend to produce *underestimates* of these probabilities. The first, which is independent of the approach used, is the problem of model misspecification. When estimating the probability of a very unlikely event, the probability of an inaccurate

calculation leading to a substantial underestimate of the risk (due to an error in model or in arithmetic) may exceed the probability of the event estimated by the analyst, making the estimate unreliably low in a way that may not be recognized by the analyst.⁽⁷⁾

Other factors are particular to the method used to estimate such probabilities. Logistic regression, a commonly-used statistical method for estimating the probability of rare events from large datasets, has been shown to systematically underestimate such probabilities.⁽⁸⁾ Moreover, the use of point estimates to represent probabilities tends to lead risk analysts to underestimate low probabilities. Hansson writes:⁽⁹⁾ “Consider, for instance, an estimate that the probability of an explosion in a certain pressure vessel in the next year is 10^{-5} . This probability may be 2×10^{-5} too low (i.e. the correct value may be 3×10^{-5}), but it cannot be 2×10^{-5} too high (since it cannot be negative). Due to this asymmetry, a risk-benefit analysis based exclusively on the central, most probable estimate can be expected to be more risk-prone than the ‘risk-neutral’ ideal of consistently maximized expected utility.”¹ This problem is particularly acute when estimating probabilities of events that have not yet occurred.⁽¹⁰⁾ If in x factory-years of experience there have been no explosions, the maximum-likelihood estimate of the probability of an explosion in any given year is zero. This estimate is uncertain, but all of the uncertainty lies to the right of the maximum-likelihood estimate. Thus, use of the maximum-likelihood point estimate in this case may very well underestimate the true risk, and cannot overestimate it. As noted above, debates about which historical context is directly relevant to estimating a probability can lead to whittling down of the historical record to such an extent that there are indeed zero events of the sort whose probability is being estimated.

3.2. Rational incentives

Putting aside the difficulties in estimating a risk to unpredictable statistical lives, what are the rational incentives on that agent to take appropriate levels of care to mitigate that risk? Specifically, will the level of care taken for unpredictable lives is equivalent to that taken when an equal number of predictable statistical lives are at stake. In this section, we concentrate on the behavior of a risk-neutral, profit-maximizing firm as the agent in question, and consider its rational incentives. We identify two economic conditions that may provide an incentive to this firm to invest less to mitigate risks to unpredictable statistical lives than it would to protect the same number of predictable statistical lives. These conditions are: a) Limited liability, which reduces the amount of financial risk to firms in the event of a low-probability, high-consequence accident, and b) competition from firms that do not choose to mitigate the risk to unpredictable lives, which may make it unprofitable for any firm to compete in the marketplace if it does mitigate that risk.

3.2.1. Limited liability—The system of limited liability, according to which a firm cannot lose more than its net assets if it goes bankrupt, generates an incentive to firms to underinvest in measures to mitigate risks that might lead to bankruptcy, or equivalently to overinvest in risky activities that may lead to bankruptcy. In liability law^(11,12), this is called

¹The word “central” here is slightly misleading, because Hansson’s argument turns on the fact that the density of a nonnegative random variable that is very close to zero is asymmetric, with the central point (the median, or perhaps the mean) to the right of the most likely point (mode) so the most-likely value is not also the central value. With this emendation Hansson’s argument holds.

the “judgment-proof problem,”⁽¹³⁾ in which liability for the damage caused by an accident will harm the firm only up to a certain level, normally its total assets, leaving the victims or society to pay the remaining costs.⁽¹⁴⁾ For risks that involve a liability exceeding the firm’s assets, the firm has a financial incentive to treat the risks as if they involved a lesser amount of liability – equal to the firm’s net assets.

The judgment-proof problem arises in the context of accident risk only for low-probability, high-consequence events, which typically involve unpredictable statistical lives. It applies only to high-consequence events because it applies only when the liabilities incurred by an accident exceed the firm’s total net assets, which will for a firm of considerable size mean that many lives have been lost. It applies only to low-probability events because measures to reduce the risk of high-probability, high consequence accidents would be worthwhile to avoid a high probability of bankruptcy. The judgment-proof problem thus tends to arise in contexts surrounding measures considered by an agent to protect statistical lives that are unpredictable from the agent’s perspective.

3.2.2. Competition from nonmitigating firms—Suppose that a firm that is a monopolist faces a decision about spending money to mitigate a risk to a number of unpredictable statistical lives. Its risk analysis finds that its expected profits if it mitigates that risk are larger than if it does not. If it expects to make an adequate profit even after accounting for the costs of mitigating, it would mitigate.

Now suppose that the firm were competing against other firms that did not mitigate. Such firms might have fewer assets and the protection of limited liability, so they do not mitigate because of the judgment-proof problem. Alternatively, they may underestimate the probabilities due to one or more of the factors described in Section 3.1, and may therefore believe (incorrectly) that mitigation is not cost-effective in expectation. In theory, rational agents who observe that they are making different probability estimates would update until they agree; in practice such observation is difficult and such consensus is unlikely. In this case, the nonmitigating competitors will have lower costs and will be able to offer the product at a lower price, reducing the profits of our firm. Our firm might then face a situation where its expected return is negative whether it mitigates or not, because its gross profits at the lower price set by the nonmitigating competitor(s) are inadequate to support the cost of mitigation. It would then withdraw from the market, leaving the market to the nonmitigating firms². Notably, firms that choose not to mitigate would be very likely to survive and prosper for years, even decades, in a situation where the true probability of an accident is only (say) 1% per year, because on average an accident would happen to such a firm only once every hundred years.

The same competitive dynamics might occur within a firm. If two analysts within different divisions of a large firm differ in their estimates of a low probability – one estimating the correct figure of 1%, and leading her division to mitigate or withdraw from a market, while the other erroneously estimates the probability, thinking it is 0.1%, leading her division to

²In some markets, the lack of mitigation might be very visible, and the firm could try to compete on for example an ethical image. But if there is a substantial market share who are selecting on price, then the dynamic will hold at least for that share.

avoid mitigation and stay in the market, the second analyst's division will most likely outperform the first analyst's division for decades, the length of these analysts' (and their bosses') careers.⁽¹¹⁾ Performance bonuses, normally paid for annual performance with no clawback provisions if performance is disastrous in future years, provide incentives to maximize short term performance. This is another aspect of the difficulty of predicting rare events: an agent who systematically underestimates small probabilities will usually not be proved wrong in any short span of time, and indeed may be rewarded for these underestimates. Similar incentives have been identified in the financial sector where money rather than lives is at risk.⁽¹⁵⁾

In summary, those firms that overestimate or correctly estimate the risk may be driven out of the industry by those that underestimate or discount the risk, because the latter firms will set the lowest price in the market and will remain profitable, potentially for many years, before facing the consequences of their error. This phenomenon shows some similarity to the "winner's curse" in auction theory⁽¹⁶⁾ and to the related "unilateralist's curse."⁽¹⁷⁾

3.2.3. Numerical Illustration—The example in this section illustrates with numbers the operations of these economic incentives. We consider three cases in which a self-interested, risk-neutral firm might make decisions about statistical lives. The firm operates a factory that, each year the factory is in operation, creates a risk (to be specified further below) in which each of L people are exposed to a probability p of death. The family of each of the people who die as a result of the factory's operation will be able to successfully sue the firm, costing the firm C dollars for each death. Throughout our examples, $p=1\%$, $L=2000$ people, and $C=\$1$ million.

The firm faces a choice of whether to operate a safety device, at a cost of M per year, which completely prevents the risk of harm to the L people. The gross profits the firm makes from the products of the factory, if it is a monopolist and can set a price for its goods, will be G , excluding the cost of mitigation and of liability. If, however, there are other firms in the marketplace that do not mitigate, competition from these firms will cause the price of goods to fall such that the firm only makes gross profits (before liability and mitigation cost) $G' < G$. The firm has total net assets W . Throughout our examples, $M=\$5$ million, $G=\$10$ million, and $G'=\$4$ million. W varies as described below.

Table I shows three cases, which differ in the mechanism by which the factory's operations lead to deaths (Case I: predictable, Case II and III: unpredictable), and the firm's net assets (Cases I and II equal, Case III higher).

In Case I, the lives are put at risk through leakage of an effluent that will poison the water in a nearby community of 2000 people, causing the death of approximately 20 members of the community. Whether the effluent kills any individual is independent of whether it kills other individuals. Thus the deaths are predictable, in the sense that approximately 20 lives will be lost as the result of operating the factory for a year, if the safety device is not installed. In Cases II and III, there is instead a 1% probability of a massive explosion at the factory that would kill 2000 people; their deaths are unpredictable in that either all 2000 will die (with probability 1%) or none will (with probability 99%). The difference between them is that in

Case II, the accident would result in liability claims that exceed the firm's assets, leading to bankruptcy. The firm would lose all its assets $W=\$300$ million, but no more, in line with the limited liability that prevails in most developed countries.⁽¹⁴⁾ In Case III, by contrast, the firm has more assets ($W=\$2.5$ billion) and would not be bankrupted by the claims resulting from the accident.

We show in what follows that a risk-neutral firm would run the mitigation system in Case I, but it would not in Case II. In Case III, the firm would choose to mitigate system if it were a monopolist, but in a competitive market might choose not to, or might leave the industry, leaving other firms to run similar factories without mitigation. These conclusions depend on the values of the particular parameters in the example, and our examples amount to an "existence proof" that there are circumstances in which these different choices would be rational. In the Appendix we give the general conditions which suffice to produce this behavior.

In Case I, the firm's assets are \$300 million, and running the factory for a year leads to release of an acutely toxic effluent, which is expected to result in the deaths of 20 exposed people who live downstream from the factory. The expected costs of compensation are \$1 million per death. The firm thus predictably faces around \$20 million per year in legal liability if it does not mitigate, which it can reduce to 0 by mitigation. Here it is easy to see that mitigation at a cost of \$5 million to avoid \$20 million in liability costs is a good investment, so the firm will mitigate.

Similar arithmetic applies for Case I if, instead of being a monopolist, the firm is in a competitive industry competing against some firms that do not mitigate. Gross profits are lower, leading to lower net profits (in fact, net losses) whether or not the firm mitigates. Even so, the firm does better with mitigation than without. The situation with a competitive industry and nonmitigating competitors of less interest here, as each firm facing the decision whether or not to mitigate will see that nonmitigation will lead to predictable large losses, so there might be no nonmitigating competitor in this scenario. Overall, Case I shows that when the statistical lives at stake are predictable, under a certain set of assumptions about the costs and benefits, the risk-neutral firm will mitigate.

Case II considers a firm that also has assets of \$300 million, but a different mechanism by which lives may be lost from the factory's activity. Here, lives are lost in a low-probability (1%), high consequence (2000 lives) accident, with the same expected lives at stake. Here the expected net profits for a mitigating firm are as in Case I, since the mitigation removes the accident risk and with it the liability risks. For a nonmitigating firm, expected net profits are a weighted average of losing the entire assets of the firm (with probability 1%) and making a profit of \$10 million (with probability 99%). Here, bankruptcy laws limit the firm's losses in the event of an accident to its net assets of \$300 million, much less than the \$2 billion in damage if the accident occurs. The limited liability system externalizes the risk above and beyond the firms assets onto society, thereby subsidizing risk-taking by firms.⁽¹²⁾ Here the subsidy is sufficiently large that, even in expectation, the firm will do better by not mitigating the risk than by mitigating it. Its expected losses from accident risk are not the expected legal liabilities $pCL = \$20$ million, but the expected amount it would lose, which is

equal to its assets times the probability of the accident, $pW = \$3$ million. These expected losses are not sufficient to offset the certain costs of mitigation. Thus a risk-neutral firm would not mitigate.

Now consider Case III, where all assumptions are as in Case II, except that the firm has much larger assets of $W = 2.5$ billion. These assets exceed the liability in the event of the accident, so the firm will not go bankrupt if the accident occurs. The firm will therefore face the full cost of its accident liability, unsubsidized by limited liability laws.⁽¹⁴⁾ Without such a subsidy, if the firm is a monopolist, it will face higher expected net profits from mitigating than not, as in Case I, and will mitigate.

If the firm is in an industry with non-mitigating competitors, however, it will expect to lose money whether or not it mitigates. In such a setting, a risk-neutral firm would withdraw from the industry because it was not profitable in expectation – in effect, it would be driven out of business by its non-mitigating competitors, leaving only nonmitigators in the industry.

We have compared three cases in which an expected 20 lives are at risk from the activities of a firm. In Case I the firm has an incentive to spend money to prevent the risk to these lives, and this incentive occurs because the risks are *predictable* and thus subject to the law of large numbers, which ensures that with near-certainty the number of lives lost without mitigation will be approximately 20. With the particular assumptions we have made about the costs and benefits of mitigation, the firm will choose to mitigate rather than suffer the financial losses resulting from those 20 deaths. In Cases II and III the lost lives are *unpredictable*: with 1% probability, 2000 lives are lost, and with high probability none are; in expectation the number lost is 20. The costs of mitigation in these cases remain the same. In Case II, the firm's limited assets, combined with the bankruptcy laws that limit liability to the assets of a firm, create a subsidy for taking the risk of an accident, and the firm chooses not to mitigate, a phenomenon well-known in liability law⁽¹²⁾. In Case III, we have increased the assets of the firm so that the subsidy from limited liability does not operate, and we find that the firm will likely withdraw from the market, as its expected profits are negative whether or not it mitigates. Above, we described several reasons why some other firms might underestimate the (difficult-to-estimate) probability of an accident, and based on that estimate (or on other variations in the economics of those firms) will choose to stay in the industry and not to mitigate. Even if they are wrong, they will most likely prosper for years or decades before an accident occurs. Thus the marketplace will be left to non-mitigators. Thus competition in Case III, or the subsidy from the liability system in Case II, both create incentives to undervalue unpredictable statistical lives, relative to the same number of predictable ones.

4. ASSUMPTIONS

Considering the decisions of a self-interested firm permits straightforward calculations about the consequences of various decisions and the resulting incentives. However, profit-seeking firms are not the only agents that take low-probability, high-consequence risks. Such risks may also be taken by scientists and their non-profit research institutes⁽¹⁷⁾ and by national governments in the areas of defense, environmental engineering⁽¹⁷⁾, and science policy.⁽¹⁸⁾

Agents in these fields might face unduly weak incentives to invest in hedges against events that have low probability during any particular term of office – an influenza pandemic, an asteroid impact, or the like, because the same resources could be put into activities that are more likely to provide short-term benefits. In these fields too, low-probability events are, by definition, rare, and therefore do not provide a strong check on the actions of agents who, for whatever reason, choose to take on risks of low-probability accidents or to forego preparation for low-probability catastrophes.

This discussion has suggested that there are circumstances in which rational agents will care less for unpredictable statistical lives than for predictable ones. For the limited-liability scenario, the legal literature suggests the problem is real, and that indeed corporations make decisions about their structure precisely to make the most hazardous subsidiaries the most asset-poor, and to use debt financing to further reduce the assets of corporations that risk unpredictable accidents.^(12,14) We are unaware of systematic empirical evidence about the extent to which competitive pressures prompt such undervaluing.

We have focused largely on factors that tend to induce agents to take less care to protect unpredictable statistical lives than predictable ones. In particular cases, there may be other factors whose effects on levels of care taken by an agent could produce greater protection for either predictable or unpredictable lives. Behavioral economists have cataloged a long list of ways in which humans tend to estimate risk relative to some objective standard.

Cognitive biases may lead to either overestimation or underestimation of any particular risk, and thus could have complex effects on the degree of care taken for predictable or unpredictable statistical lives. For example, the availability heuristic,⁽¹⁹⁾ in which risks that are more easily accessed in memory are considered more serious, may lead to overestimation of highly available risks and underestimation of those less available to the agent. The tendency to greater concern about risks that are more inspiring of dread or more unfamiliar⁽²⁰⁾ may similarly function in either direction, depending on the characteristics of the particular risk in question. Overall, some such cognitive biases may point more often in one direction than another, but to the extent that many such biases are operative, they will increase the range of risk estimates made by individuals, and (unless firms can systematically avoid the biases of individuals) also by firms.

Firms might take into account costs from harming statistical lives beyond the liability costs, such as reputational effects. Firms with a reputation for causing damage to bystanders might lose sales, while those with “clean” reputations might be able to charge a price premium. Particularly in industries with strong brand identity, association with catastrophes could cause long-term harm to a brand, and this harm might be much smaller for the smaller events associated with predictable statistical lives. For example, an airline official stated recently that painting over an airplane’s logo after the plane was involved a crash (and while it is on the runway exposed to media attention), in order to reduce the association of the brand with a disastrous event, is “routine practice used all over the world”⁽²¹⁾. It would be interesting to consider further how such considerations would affect the care taken for unpredictable lives. On one hand, the brand itself has finite value, so the limited liability effect would still hold for very large catastrophes. On the other, for medium-sized

unpredictable events, not large enough to cause bankruptcy but large enough to cause major damage to the brand above and beyond liability costs, it is possible that such reputational considerations would increase the degree of care taken for unpredictable lives. It would also be interesting to further investigate whether there is a general pattern such that low-probability, high consequence disasters have greater reputational effects than more routine forms of damage to statistical lives, such as pollution. We have established here that there are certain conditions under which rational firms will underprotect unpredictable statistical lives, but empirical and theoretical work remains to be done to identify how frequently these conditions obtain and how wide is the explanatory power of this consideration for firms' actual behavior.

In summary, there are many sources of “noise” in the estimation of probabilities and the incentives to act on risks. How to address these is an open question, but in many cases we find it hard to identify a systematic direction of bias induced by such noise. The mechanisms that form the focus of this paper, by contrast, have a clear direction in favoring care for predictable over unpredictable statistical lives.

5. ETHICAL AND POLICY IMPLICATIONS

We have claimed that unpredictable statistical lives may often be protected less than predictable ones. If we stipulate that there are cases where some of the mechanisms we describe lead to underprotection of unpredictable statistical lives, relative to predictable ones, we see no reason why this is ethically justifiable. As long as an agent's decisions are guided by the *expected* value of their actions—a key tenet for most, if not all modern forms of consequentialist ethics³—and the value of a number of lives scales linearly with the number affected—they are committed to not treating statistical persons differently by virtue of their predictability.⁽²⁴⁾ Our definition of unpredictability is related to Keeney's definition of catastrophes, in a paper noting the opposition between risk equity and risk aversion;⁽²⁵⁾ the literature on catastrophe aversion vs. risk equity^(26–29) considers issues related to but distinct from those considered in this paper.

The ethical justification for reasoning based on expected number of statistical lives at stake has been questioned or qualified by the argument that agents have justifiable reasons to take greater care for some statistical lives than others. An agent may have greater, or stronger, duties of care to some persons than to others;⁽³⁰⁾ prevention of risks that engender fear⁽³¹⁾ or arise from morally culpable causes of harm (e.g. fires from arson)⁽³¹⁾ may deserve greater care. There may be an ethical difference between jeopardizing statistical lives by acts of commission and by acts of omission or by acts that do or do not involve violation of an individual's autonomy,⁽³²⁾ even if the same number of statistical lives would be jeopardized.

Given these complexities, we do not claim that all statistical lives should receive equal weight under all conditions. However, the predictable statistical lives that receive greater care in the example of this paper do not seem to meet any of the conditions enumerated

³We note the existence of arguments that the numbers of lives at stake in a decision should not, even if all else were equal, determine the right course of action⁽²²⁾, but we do not pursue this line of reasoning⁽²³⁾.

above—special obligation, greater fearsomeness, greater culpability, commission vs. omission, or violation of autonomy – that have been argued to justify differential care. In the absence of such a distinction, the mere fact of unpredictability *per se* cannot justify a lower degree of care for unpredictable statistical lives. It has become a central principle of risk-benefit analysis that risk can be quantified as probability times consequence; indeed, this principle was first established in the context of nuclear power accidents, classic low-probability, high-consequence events.⁽³³⁾ In the absence of reasons to depart from this approach, the principle of equal treatment would seem to dictate equal treatment of predictable and unpredictable statistical lives.

Whether or not an agent is ethically justified in providing lesser care for unpredictable than predictable statistical lives, society as a whole is often ethically and indeed prudentially justified in inducing agents within society to treat them equally. This is because lives that are unpredictable from the perspective of any individual agent may be predictable from the perspective of a society or a national government. To continue our example, an accident that has a probability of 1% at any given factory is unlikely in each factory, but if there are enough factories within one jurisdiction, an accident at *some* factory might be quite likely. If there were 300 factories within a jurisdiction, each incurring a risk of an accident killing 2000 people with a probability of 1%, then the expected number of deaths within the jurisdiction would be 6000, and the probability of at least one accident in a given year would be large – $1 - (.99)^{300}$ or approximately 95%. If society values each of these lives at \$1 million, to continue our example, it should be willing to pay \$1 million to protect each life, or \$6 billion, more than the $300 \times \$5 \text{ million} = \1.5 billion that mitigation would cost. In short, in a situation that produces a calculation under Case II or Case III for an individual firm and may induce underprotection of unpredictable lives, the equivalent calculation for the society would be similar to Case I, where the lives are predictable, and would result in a judgment that mitigation is worth the cost. While mitigation might not be in the interest of a risk-neutral firm, it would be cost-effective for the society as a whole.

In this sense, an agent's decision not to mitigate a risk to lives that are from its perspective unpredictable is a form of market failure, because the agent has incentives to act as if the risk were smaller than it truly is. This can produce behavior that is privately advantageous to the agent but socially disadvantageous. Since the agent cannot compensate society fully in the case of an accident, running the factory without mitigation is imposing an (unpredictable probabilistic) externality. As with externalities that are the result of predictable statistical behavior—such as pollution—these are appropriate targets for social regulation.^(34,35)

In cases where such underinvestment occurs, society (or the state) has reason to design interventions and regulation that account for the special, and often overlooked status of these lives in decision making. A number of approaches could be considered for such interventions. Each of the following approaches has been discussed as a potential solution to the judgment-proof problem, and in principle each could also be applied to other situations of underinvestment in mitigating low-probability risks including protection of unpredictable statistical lives.

One approach is regulation (for example environmental, health and safety regulation), in which the state requires investments in risk mitigation that are socially efficient but may not be efficient for each firm in the absence of regulation. Alternatively, firms engaged in certain risky activities that might be subject to these incentives for underinvestment might be compelled to purchase liability insurance⁽³⁶⁾. This solution has the advantage that it places the burden of estimating the risk on the insurer and/or reinsurers, who (a) are specialists in risk estimation who should be aware of the pitfalls described above and (b) have incentives to estimate risks correctly. Since they usually have more capital and (like governments) are exposed to more independent risks, a risk which is unpredictable from the perspective of the firm may become predictable from the perspective of the insurer.

Requiring liability insurance is controversial because holding liability insurance itself creates moral hazard by externalizing risk to the insurer, and may itself incentivize inadequate care-taking in particular if the insurer has limited ability to observe the actions of the insured.⁽¹³⁾ For example, if mitigation involves buying and installing a technological system, it may be relatively easy for the insurer to observe and thereby charge a premium reflecting the actual probability of harm. If mitigation is harder to observe, the firm may take less care than the insurer thinks it is taking, and not be caught doing so. This could lead to the firm's taking much less care to mitigate than if its own assets were on the line – a classic example of moral hazard. Thus the feasibility of requiring liability insurance depends on the particular kind of risk and the types of mitigation involved, which help to determine the level of concern about moral hazard. Other possible responses to the judgment-proof problem include taxation, prohibiting the purchase of liability insurance (the opposite of compelling it), and criminal liability⁽¹²⁾. Each of these approaches could in principle be extended to other situations in which some factor other than the judgment-proof problem prompts inadequate care. Each, however, has particular social and economic implications, and choosing from among them is beyond our scope.

In some cases, at least, the very agents that would have an incentive not to mitigate risks to unpredictable statistical lives might, with equal rationality, favor regulation that requires all agents to undertake such mitigation. This seems particularly plausible for some kinds of accidents (eg nuclear meltdown, or a high-consequence accident from biological research), where the agents performing the risky activity anticipate that such an accident, even if caused by a competitor, would lead to a crackdown on the whole industry including one's own activity (e.g. governments might defund nuclear power or curtail biological research). This may provide an incentive for participants to call for regulation, since with a large enough industry, the collective externality on each agent from the other agent's risky activity would be considerable, even if each agent were willing to accept the small incremental contribution he made to the total risk of an accident. This is reminiscent of Hardin's slogan of "mutual coercion, mutually agreed upon" to solve the tragedy of the commons⁽³⁷⁾. Notably, regulation requiring mitigation by all participants would (if compliance were assured) also remove the incentive not to mitigate because of competition from a nonmitigating competitor.

We have argued that agents with the power to make decisions affecting unpredictable statistical lives will do so differently, in many situations, from similarly situated agents

facing decisions affecting predictable ones. Factors prompting these differences include statistical challenges in estimating improbable risks as well as rational incentives to neglect low-probability risks. While some of these factors may induce either over- or under-protection of unpredictable lives, many of them specifically spur underprotection of unpredictable lives, relative to the same number of predictable lives. Ethically, there appears to be no justifiable basis for such differential care. Moreover, lives that are unpredictable from the perspective of any individual agent may be predictable from the perspective of a society, because the law of large numbers applies across many low-probability risks, although it does not apply to any individual one. Further consideration is needed to address what means of redress may be appropriate to address the ethical and policy issues raised here.

Acknowledgments

We thank Carl Bergstrom, Ted Bergstrom, Nir Eyal, Jim Hammit, Meira Levinson, Daniel Markovits, and Annette Rid for helpful comments on earlier drafts. While working on this project, OCB received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 669751). This article reflects only the view of the authors. The ERCEA is not responsible for any use that may be made of the information it contains. ML received support from Award Number U54GM088558 from the National Institute Of General Medical Sciences. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of General Medical Sciences or the National Institutes of Health.

References

1. Cohen, IG., Daniels, N., Eyal, N. Identified Versus Statistical Lives. New York, NY: Oxford University Press; 2015.
2. Schelling, TC. The Life You Save May Be Your Own. In: Chase, SB., editor. Problems in Public Expenditure Analysis. Vol. 127. Washington, DC: Brookings Institution; 1968. p. 129-130.
3. Lipsitch M, Inglesby TV. Moratorium on Research Intended To Create Novel Potential Pandemic Pathogens. *mBio*. 2014; 5(6):e02366–14–e02366–14. [PubMed: 25505122]
4. Fouchier RAM. Studies on influenza virus transmission between ferrets: the public health risks revisited. *mBio*. 2015; 6(1)
5. Lipsitch M, Inglesby TV. Reply to “Studies on influenza virus transmission between ferrets: the public health risks revisited”. *mBio*. 2015; 6(1)
6. Klotz LC. Comments on Fouchier's calculation of risk and elapsed time for escape of a laboratory-acquired infection from his laboratory. *mBio*. 2015; 6(2):e00268–15. [PubMed: 25873376]
7. Ord T, Hillerbrand R, Sandberg A. Probing the improbable: Methodological challenges for risks with low probabilities and high stakes. *Journal of Risk Research*. 2010; 13(2):191–205.
8. King G, Zeng L. Logistic Regression in Rare Events Data. *Political analysis*. 2001; 9(2):137–163.
9. Hansson SO. Economic (ir)rationality in risk analysis. *Economics and Philosophy*. 2006; 22(02): 231.
10. Hanley JA, Lippman-Hand A. If nothing goes wrong, is everything all right? Interpreting zero numerators. *JAMA: the journal of the American Medical Association*. 1983; 249(13):1743–1745. [PubMed: 6827763]
11. Siliciano JA. Corporate Behavior and the Social Efficiency of Tort Law. *Michigan Law Review*. 1987; 85(8):1820.
12. Shavell S. Chapter 2. Liability for Accidents. *Handbook of law and economics*. 2007; 1:139–182.
13. Shavell S. The judgment proof problem. *International Review of Law and Economics*. 1986; 6(1): 45–58.
14. Hansmann H, Kraakman R. Toward Unlimited Shareholder Liability for Corporate Torts. *The Yale Law Journal*. 1991; 100(7):1879.

15. Taleb NN, Martin GA. The Risks of Severe, Infrequent Events. *The Banker*. 2007 Sep.:188–189.
16. Thaler RH. Anomalies: The Winner's Curse. *Journal of Economic Perspectives*. 1988; 2(1):191–202.
17. Bostrom N, Sandberg A, Douglas T. The Unilateralist's Curse: The case for a principle of conformity. 2013
18. Smithson, M. Unknowns in the Dual-Use Dilemmas. In: Smithson, M. Rappert, B., Selgelid, M., editors. *On the Dual Uses of Science and Ethics*. Canberra: ANU E Press; 2013. p. 165–184.
19. Tversky A, Kahneman D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*. 1973; 5(2):207–232.
20. Slovic P, Fischhoff B, Lichtenstein S. Behavioral decision theory perspectives on risk and safety. *Acta psychologica*. 1984; 56:183–203.
21. Coldwell, W. Thai Airways and that logo – just part of post-plane-crash etiquette?. *The Guardian*. <http://www.theguardian.com/world/2013/sep/09/thai-airways-logo-crash-etiquette>
22. Taurek JM. Should the Numbers Count? *Philosophy & Public Affairs*. 1977; 6(4):293–316. [PubMed: 11662404]
23. Parfit D. Innumerate ethics. *Philosophy & Public Affairs*. 1978; 7(4):285–301.
24. Pettit P, Brennan G. Restrictive consequentialism. *Australasian Journal of Philosophy*. 1986; 64(4):438–455.
25. Keeney RL. Equity and Public Risk. *Operations Research*. 1980; 28(3-part-i):527–534.
26. Adler MD, Hammitt JK, Treich N. The social value of mortality risk reduction: VSL versus the social welfare function approach. *Journal of Health Economics*. 2014; 35:82–93. [PubMed: 24699210]
27. Rheinberger CM. Experimental Evidence Against the Paradigm of Mortality Risk Aversion. *Risk Analysis*. 2010; 30(4):590–604. [PubMed: 20163558]
28. Covey J, Robinson A, Jones-Lee M, Loomes G. Responsibility, scale and the valuation of rail safety. *Journal of Risk and Uncertainty*. 2009; 40(1):85–108.
29. Jones-Lee MW, Loomes G. Scale and context effects in the valuation of transport safety. *Journal of Risk and Uncertainty*. 1995; 11(3):183–203.
30. Goodin RE. What is so special about our fellow countrymen? *Ethics*. 1988; 98:663–686.
31. Wolff J. Risk, fear, blame, shame and the regulation of public safety. *Economics and Philosophy*. 2006; 22(03):409–427.
32. Thomson JJ. A defense of abortion. *Philosophy & Public Affairs*. 1971; 1(1):47–66.
33. Rechar RP. Historical Relationship Between Performance Assessment for Radioactive Waste Disposal and Other Types of Risk Assessment. *Risk analysis: an official publication of the Society for Risk Analysis*. 1999; 19(5):763–807. [PubMed: 10765434]
34. DeMartino, G. *Global Economy, Global Justice*. Routledge; 2002.
35. Moss, DA., Cisternino, J. *New Perspectives on Regulation*. Cambridge, MA: The Tobin Project; 2009.
36. Jost P-J. Limited liability and the requirement to purchase insurance. *International Review of Law and Economics*. 1996; 16(2):259–276.
37. Hardin G. The tragedy of the commons. The population problem has no technical solution; it requires a fundamental extension in morality. *Science (New York, NY)*. 1968; 162(3859):1243–1248.

Table I

Three cases of statistical lives

		Case I: Predictable lives	Case II: Unpredictable lives, bankruptcy	Case III: Unpredictable lives, competition
Assumptions	W : Assets	300,000,000	300,000,000	2,500,000,000
	M : cost of mitigation	5,000,000	5,000,000	5,000,000
	p : probability of losing each life	1%	1%	1.0%
	L : number of lives at risk	2,000	2,000	2,000
	C : cost per life	1,000,000	1,000,000	1,000,000
	Liability - if accident	20,000,000	2,000,000,000	2,000,000,000
	Liability - no accident	20,000,000	0	0
	Liability - predictable	20,000,000	0	0
Choice for monopolist	G : Gross profits - monopolist	10,000,000	10,000,000	10,000,000
	Net profit, without mitigation, if accident occurs	-10,000,000	-300,000,000	-2,000,000,000
	Net profit, without mitigation, if no accident occurs	-10,000,000	10,000,000	10,000,000
	Expected net profit, without mitigation	-10,000,000	6,900,000	-10,100,000
	Net profit with mitigation	5,000,000	5,000,000	5,000,000
Choice for firm with nonmitigating competitor	G : Gross profits - nonmitigating competitor	4,000,000	4,000,000	4,000,000
	Net profit, without mitigation, accident occurs	-16,000,000	-300,000,000	-2,000,000,000
	Net profit, without mitigation, no accident occurs	-16,000,000	4,000,000	4,000,000
	Expected net profit, without mitigation	-16,000,000	960,000	-16,040,000
	Net profit with mitigation	-1,000,000	-1,000,000	-1,000,000