DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

# Underprotection of Unpredictable Statistical Lives Compared to Predictable Ones

**Citation**
Lipsitch, Marc, Nicholas G. Evans, and Owen Cotton-Barratt. 2016. "Underprotection of Unpredictable Statistical Lives Compared to Predictable Ones." Risk Analysis 37 (5) (July 9): 893–904. doi:10.1111/risa.12658.

**Published Version**
doi:10.1111/risa.12658

**Permanent link**
http://nrs.harvard.edu/urn-3:HUL.InstRepos:33731744

**Terms of Use**
This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

## Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

Underprotection of unpredictable statistical lives compared to predictable ones

Marc Lipsitch (corresponding author)
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
mlipsitc@hsph.harvard.edu

Nicholas G. Evans
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
owen.cotton-barratt@philosophy.ox.ac.uk

**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* is 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 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 $2\times10^{-5}$ too low (i.e. the correct value may be $3\times10^{-5}$), but it cannot be $2\times10^{-5}$ 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

hsphit 6/4/16 8:43 AM
**Comment [1]:** I put the reference here so it would not clash with the footnote

[1] The 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.

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 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 firms[2]. 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.

[2] In 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.

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 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.