Distributions of observed death tolls govern sensitivity to human fatalities

Christopher Y. Olivola[a,b,1] and Namika Sagara[c,d,e]

[a]Department of Cognitive, Perceptual, and Brain Sciences, University College London, London WC1H 0AP, United Kingdom; [b]Economic and Social Research Council, Centre for Economic Learning and Social Evolution, London WC1H 0AN, United Kingdom; [c]Lundquist College of Business and [d]Department of Psychology, University of Oregon, Eugene, OR 97403; and [e]Decision Research, Eugene, OR 97401

Edited by Dale Purves, Duke University Medical Center, Durham, NC, and approved October 28, 2009 (received for review August 7, 2009)

How we react to humanitarian crises, epidemics, and other tragic events involving the loss of human lives depends largely on the extent to which we are moved by the size of their associated death tolls. Many studies have demonstrated that people generally exhibit a diminishing sensitivity to the number of human fatalities and, equivalently, a preference for risky (vs. sure) alternatives in decisions under risk involving human losses. However, the reason for this tendency remains unknown. Here we show that the distributions of event-related death tolls that people observe govern their evaluations of, and risk preferences concerning, human fatalities. In particular, we show that our diminishing sensitivity to human fatalities follows from the fact that these death tolls are approximately power-law distributed. We further show that, by manipulating the distribution of mortality-related events that people observe, we can alter their risk preferences in decisions involving fatalities. Finally, we show that the tendency to be risk-seeking in mortality-related decisions is lower in countries in which high-mortality events are more frequently observed. Our results support a model of magnitude evaluation based on memory sampling and relative judgment. This model departs from the utility-based approaches typically encountered in psychology and economics in that it does not rely on stable, underlying value representations to explain valuation and choice, or on choice behavior to derive value functions. Instead, preferences concerning human fatalities emerge spontaneously from the distributions of sampled events and the relative nature of the evaluation process.

decision-making | psychophysics | risk preferences | value of human lives | decision by sampling

Every year, millions of human lives are lost to accidents, disasters, armed conflicts, and other deadly causes. Our reactions to these tragic events—including our willingness to provide aid and demand that our governments intervene—depend largely on the extent to which we are moved by the size of their associated death tolls (1). Research in psychology (1–4) and economics (5–7) has demonstrated that people tend to show a diminishing sensitivity to the number of human fatalities and, equivalently, a preference for risky (vs. sure) alternatives in decisions under risk involving human losses. As a result, policy makers charged with responding to humanitarian crises or preventing unnecessary deaths may inadvertently fail to maximize the number of lives saved (1). Although well documented, and despite its grave implications, the reason for this tendency is not well understood.

The standard way of explaining valuation and choice, even for human fatalities, has been to assume the existence of underlying utility functions (i.e., stable representations of value) that drive observed preferences (8). Although utility-based approaches have provided important constraints on the set of possible theories that can describe people's preferences, they are nonetheless limited in several ways. Most importantly, utility functions do not really explain preferences but merely redescribe them in mathematical terms (unless one assumes that such functions are directly represented in our mental architecture, which seems computationally implausible). Our understanding of valuation and choice concerning human fatalities would therefore greatly benefit from a process-level theory, which moves past the limitations of utility-based theories while still being able to accurately predict people's preferences.

In this article, we provide such an account and report empirical evidence to support it. Our account capitalizes on parallels between perception and decision-making (2, 9, 10). In particular, sensitivity to changes in a perceptual stimulus (e.g., brightness) or decision outcome (e.g., wealth) generally diminishes as the stimulus's initial magnitude increases (2, 9). This diminishing sensitivity implies that, in perception and decision-making alike, the relationship between objective magnitude and subjective evaluation is often a concave function that is monotonically increasing but marginally decreasing. As we noted earlier, a number of studies have shown that this psychophysical relationship even extends to evaluations of human fatalities: as an event's death toll increases, our sensitivity to the loss of life decreases, so that each additional death has a diminishing affective impact (1, 3, 4). Other studies have shown a general preference for risky alternatives in decisions under risk involving human losses (2), which also implies a concave disutility function for human fatalities (2, 8, 9).

Model Description. The model we describe builds on a recently developed theoretical framework (11–13) that uses fundamental psychological principles to explain how we evaluate relatively abstract magnitudes such as money (11, 12), time (11), probability (11, 12), color (14), and luminance (15). According to this framework, the evaluation process is governed by a few simple cognitive operations (11, 12): To evaluate the death toll associated with a specific target event (or "event-associated death-toll"; EADT), people first draw upon a sample of comparable events from their memory. Specifically, they sample from a mixture of previously observed events (i.e., long-term memory sampling) and events in their recent or immediate context (i.e., short-term memory or working memory sampling) to obtain a set of comparison death tolls. Then they compare the target EADT with all those in the sampled set. For example, a person might compare a target EADT with other EADTs that he or she has recently learned about from watching the news, reading a newspaper, or conversing with family, friends, or colleagues. The disutility or "shock" associated with a target EADT is simply the proportion of pair-wise comparisons in which it dominates or ties, which is its percentile rank among the sampled events (i.e., the proportion of sampled EADTs that are smaller than or equal

PSYCHOLOGICAL AND COGNITIVE SCIENCES

Author contributions: C.Y.O. designed research; C.Y.O. and N.S. performed research; C.Y.O. analyzed data; and C.Y.O. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: c.olivola@ucl.ac.uk.

This article contains supporting information online at www.pnas.org/cgi/content/full/0908980106/DCSupplemental.

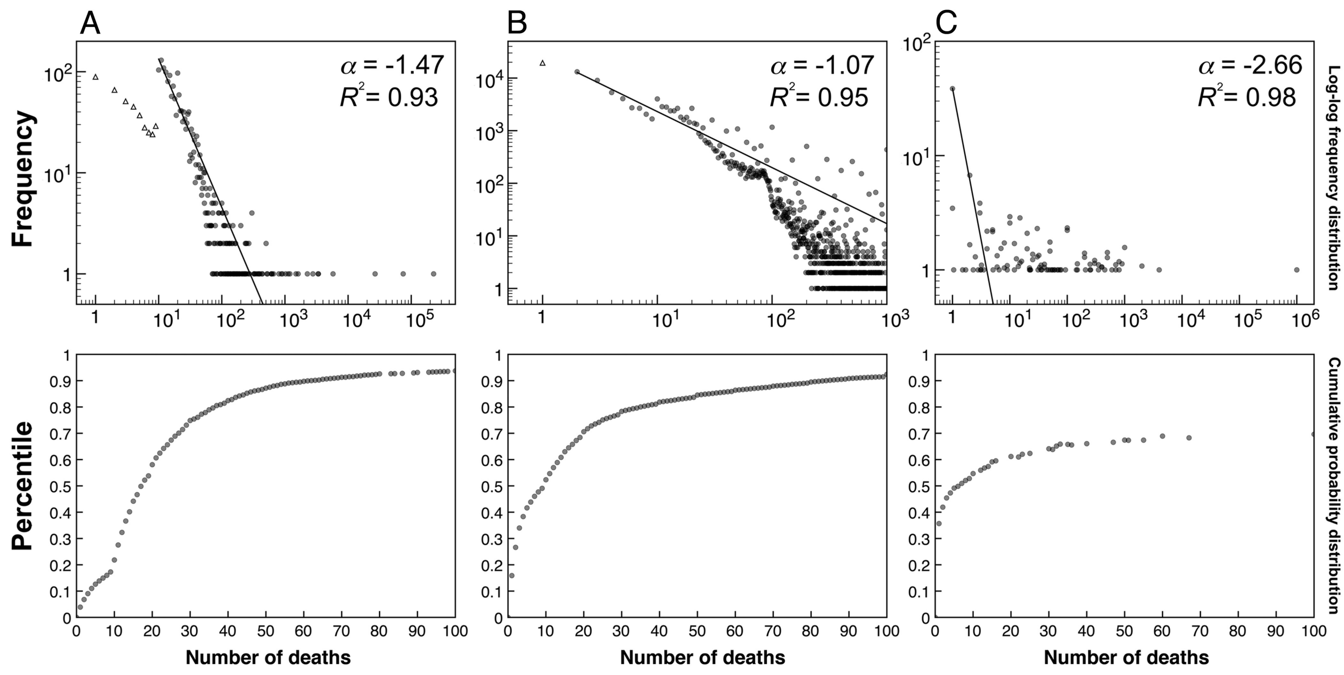

Fig. 1. Log-log plots of EADT frequency distributions (*Top*) and their corresponding cumulative probability distributions shown up to 100 deaths (*Bottom*). Solid lines (*Top*) are best-fitting power-functions of the frequency data, with the associated power parameter estimate (α) and model fit (R^2) displayed (*Top Right*). (*A*) Fatalities from natural and industrial disasters occurring in 2003 through 2007. The power-law function is fitted to events involving 10 or more deaths (dots) because frequencies are underestimated for events involving fewer than 10 deaths (triangles; see *SI Text*). The cumulative probability distribution plot is based on all of the data (triangles and dots). (*B*) Media attention (in 2000–2007) to mortality-related events. The power-law function is fitted to events involving two or more deaths (dots) because media attention to events involving a single death (triangle) is likely underestimated (see *SI Text*). The cumulative probability distribution plot is based on all the data (triangle and dots), as well as estimated frequencies for events involving more than 1,000 deaths (see *SI Text*). (*C*) Mean recalled EADTs occurring in a person's lifetime. Mean frequencies and mean percentile-ranks were obtained by repeatedly and randomly sampling the events that respondents recalled (see *SI Text*).

to the target). (A formal description of the model is provided in the *SI Text*.) Thus, a target EADT is considered large if it ranks above most sampled EADTs (and small if it ranks below most of them) regardless of its absolute magnitude. The sampling process is psychologically plausible given evidence that humans and other animals efficiently encode and recall frequencies (16). And the relative nature of the evaluation process is supported by extensive evidence that people are much better able to provide relative, rather than absolute, judgments (17, 18). For our current purposes, this model can remain agnostic about whether the memory sampling and pair-wise comparison operations are fast, unconscious, and automatic or slow, conscious, and deliberate processes (future research will be needed to answer this question). For simplicity, we start by assuming that the sampling process is uniformly random, so that every EADT in memory is equally likely to be selected for comparison.

This model implies that the disutility associated with a target EADT will be determined by the distribution of comparison EADTs from which a person can draw, which will be a function of his or her accumulated experiences and the recent or immediate context surrounding the evaluation. This allows the model to parsimoniously explain individual and contextual differences in reactions to human fatalities. Specifically, as an EADT's disutility is equal to the probability that it is larger than or equal to another randomly sampled EADT, the psychophysical (or disutility) function relating human fatalities to their associated shock values is equivalent to the cumulative probability distribution of EADTs that one has observed (see *SI Text*). And because risk preferences follow from these disutility functions (2, 8, 9, 19) (see *SI Text*), the model should be able to predict people's choices in decisions under risk involving human losses, as long as the distribution of sampled EADTs is approximately known.

This led us to make three predictions. First, diminishing sensitivity to human fatalities and risk-seeking preferences in the domain of human losses should be reflected in the distribution of EADTs that people generally observe. Specifically, for the model to explain existing results, the cumulative probability distribution of observed EADTs must be concave. In addition, the model makes two unique predictions: (*i*) that altering the distribution of observed EADTs will lead to predictable changes in people's risk-preferences and (*ii*) that risk preferences will vary between countries with different EADT distributions. We tested these predictions in three separate studies.

Study 1. First, we hypothesized that the cumulative probability distribution of EADTs that people typically observe is concave. To test this prediction, we analyzed data from three sources: (*i*) a global survey of disasters and their associated death tolls (Study 1A), (*ii*) a measure of the frequency with which EADTs are mentioned in the media (Study 1B), and (*iii*) the examples provided by a sample of respondents who were asked to recall EADTs from memory (Study 1C). This first prediction is well supported: all three cumulative probability distributions are monotonically increasing and marginally decreasing (Fig. 1 *A–C*). This concavity is a direct result of the fact that EADT frequencies are reasonably well approximated by a power-law distribution (Fig. 1 *A–C*). Power-laws are characterized by a number of interesting properties, such as scale invariance (20, 21), and may provide important clues about the underlying generation process (22, 23). With regard to the psychophysics of human life valuation, they have one particularly useful property: integrating over them yields cumulative probability functions

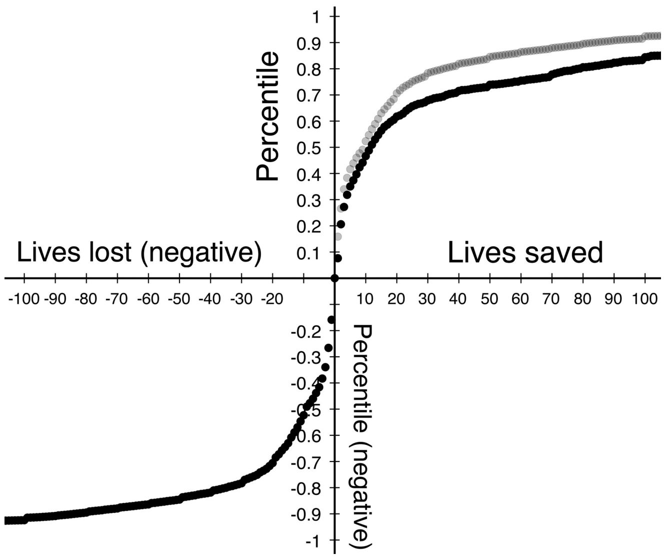

Fig. 2. Black dots represent the cumulative probability distributions of media attention (in 2000–2007) to events involving human lives saved (*Top Right*) and lives lost (*Bottom Left*). The gray dots (*Top Right*) represent a 180° rotation of the data for lives lost (*Bottom Left*) obtained by multiplying the coordinates by −1, and show that the curve is steeper for lives lost than for lives saved.

that are equivalent to the constant relative risk aversion (CRRA) utility functions often used in economics to describe preferences (19, 24) (see *SI Text*). In fact, one result of the model we describe is that the best-fitting power-law parameter provides a reasonably good measure of disutility curvature, which quantifies both the rate at which sensitivity to human life decreases and the resulting level of risk preference (see *SI Text*). The best fitting estimate of the power-law parameter for each dataset is negative (Fig. 1 *A*–*C*) and therefore implies a concave disutility function and a preference for risky alternatives (see *SI Text*). Furthermore, an examination of media attention to the number of human lives that are saved (rather than lost) reveals that the cumulative distribution of human "gains" is also concave but less steep than the cumulative distribution of lives lost (Fig. 2). We would therefore predict that the utility associated with lives saved is concave, leading to risk-averse choices (2, 9), and that human losses receive more weight than equivalent gains (9). In fact, both predictions are empirically supported (2, 25, 26).

Study 2. Second, we hypothesized that altering the distribution of EADTs that people observe would lead to predictable changes in their risk preferences. Although we have assumed, for simplification, that comparison EADTs are uniformly sampled from memory, a more realistic assumption is that recently observed EADTs are sampled with higher likelihood (27). This would allow them to have a nonnegligible impact on risk preferences even though they represent a tiny fraction of all of the EADTs in memory. We therefore predicted that exposing decision-makers to a highly concave cumulative distribution of EADTs would increase the appeal of risky options whereas exposing them to an S-shaped cumulative distribution of EADTs would have the opposite effect (see *Methods*). To test this prediction, we conducted an experiment in which we had some human participants read about a set of EADTs and rate how negatively they felt about each one. These EADTs were selected to form either a concave or S-shaped cumulative distribution (Fig. 3*A*). Following this subtle manipulation, participants were asked to choose between a risky option and a sure option in a hypothetical decision scenario involving human lives at risk (2). The rating task had a significant impact on participants' risk preferences (Fig. 3*B*). Participants who first rated a concave cumulative distribution of EADTs were subsequently more likely to choose the risky alternative than those who first rated an S-shaped cumulative distribution of EADTs [$n = 101$, $\chi^2(1) = 3.94$, $P = 0.047$, $\phi = 0.20$].

Study 3. Finally, we hypothesized that risk preferences would vary between countries with different EADT distributions (we assume that people are especially likely to observe and sample EADTs occurring in their own country). In particular, we predicted that a preference for risky options would be more prevalent in countries in which the cumulative distribution of EADTs is highly concave than in countries in which it is less concave. To test this prediction, we measured risk preferences in four countries that differ substantially in terms of their EADT distributions: Japan, the United States, India, and Indonesia. High-magnitude EADTs are relatively rare in Japan and the United States, but much more frequent in India and Indonesia [Emergency Events Database (EM-DAT); see www.em-dat.net]. As a result, the cumulative distributions of EADTs are more concave for the former two countries than for the latter two (Fig.

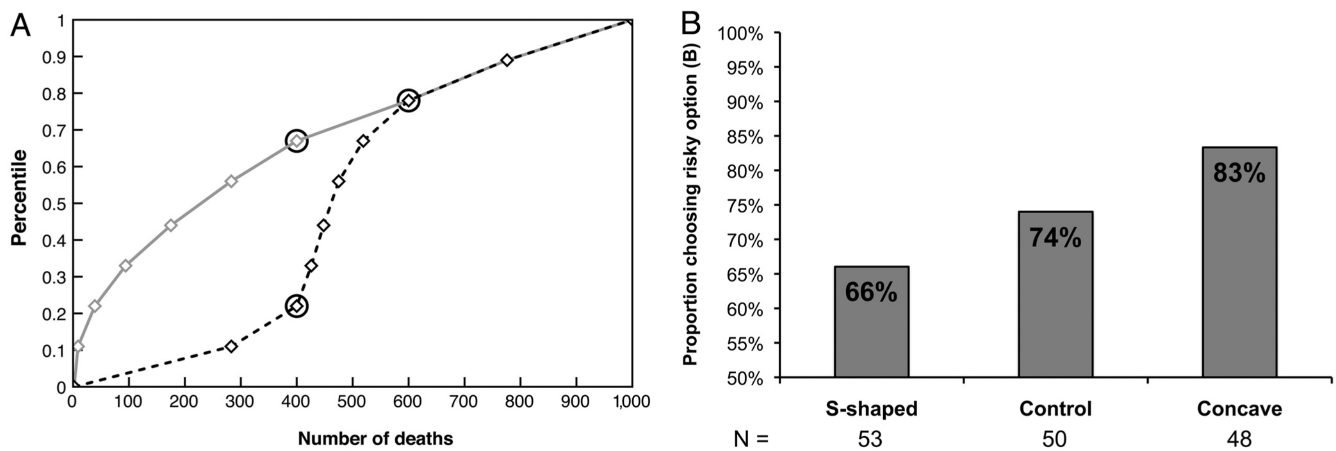

Fig. 3. Experimental manipulation and results for Study 2. (*A*) Cumulative probability distributions of EADTs (diamonds) obtained by ranking the EADTs that participants rated and the decision scenario's two target EADTs (circled). The S-shaped distribution curve (dashed black line) is superimposed on the concave distribution curve (gray line). (*B*) The proportion of participants in each condition who chose the risky option (program B). Numbers below the bars are sample sizes.

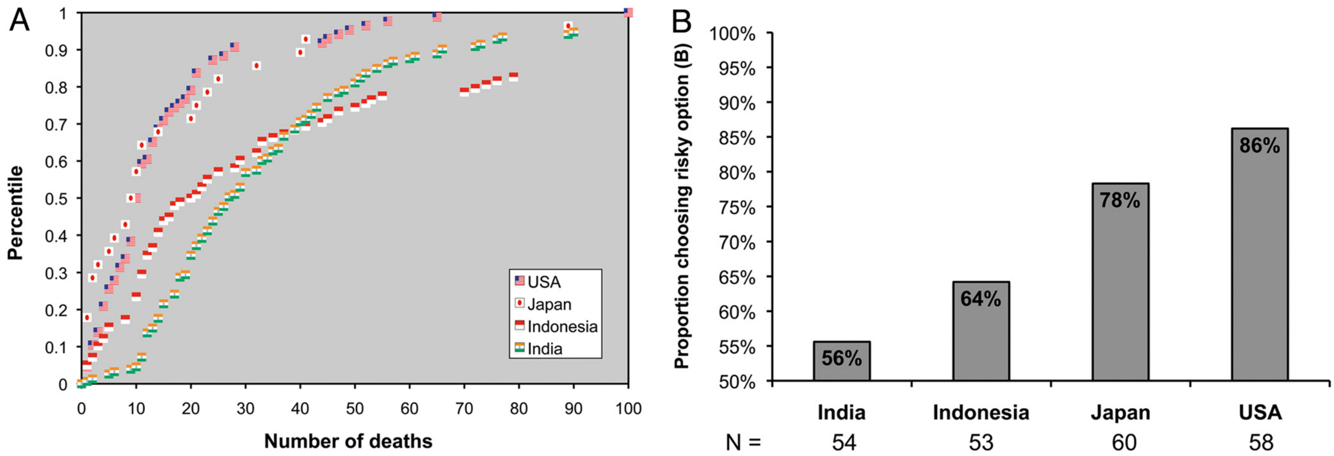

Fig. 4. Cross-national disaster deaths and risk preferences in Study 3. (*A*) Cumulative probability distributions of disaster-related fatalities in 2003 through 2007 (shown up to 100 deaths) experienced by India (best-fitting power parameter for ≥10 deaths, $\alpha = -0.59$, $R^2 = 0.33$), Indonesia ($\alpha = -3.07$, $R^2 = 0.78$), Japan ($\alpha = -4.56$, $R^2 = 0.91$), and the United States ($\alpha = -6.64$, $R^2 = 0.78$). (*B*) The proportion of respondents in each country who chose the risky option (program B). Numbers below the bars are sample sizes.

4*A*). We therefore predicted that a large majority of American and Japanese respondents would prefer the riskier option in a decision under risk involving human fatalities, but that Indian and Indonesian respondents would be less inclined to choose it. To test this prediction, we gave respondents in each country a hypothetical decision scenario similar to the one we used in our previous experiment and measured the proportion of risky option choices (Fig. 4*B*). The tendency to select the risky alternative differed across countries [$n = 225$, $\chi^2(3) = 15.69$, $P = 0.0013$, $\phi = 0.26$]. Specifically, Indian respondents were less likely to prefer the risky option than American respondents [$n = 112$, $\chi^2(1) = 12.87$, $P = 0.00033$, $\phi = 0.34$] and Japanese respondents [$n = 114$, $\chi^2(1) = 6.73$, $P = 0.0095$, $\phi = 0.24$]. Similarly, Indonesian respondents were less likely to choose the risky option than American respondents [$n = 111$, $\chi^2(1) = 7.32$, $P = 0.0068$, $\phi = 0.26$] and Japanese respondents [$n = 113$, $\chi^2(1) = 2.79$, 1-tailed $P = 0.047$, $\phi = 0.16$]. By contrast, Indian and Indonesian respondents did not significantly differ in terms of their risk preferences [$n = 107$, $\chi^2(1) = 0.82$, $P = 0.36$], nor did American and Japanese respondents [$n = 118$, $\chi^2(1) = 1.25$, $P = 0.26$]. In line with our prediction, we found that Indian and Indonesian respondents were less likely to choose the risky alternative than Japanese and American respondents (Fig. 4*B*). In fact, the relative prevalence of risk-seeking choices across countries (Fig. 4*B*) is perfectly predicted by the relative size of their EADT frequency distribution power parameter estimates (Fig. 4*A*; also see *SI Text*), which further supports the model we are testing. The possibility that these results reflect a domain-general tendency for Indian and Indonesian respondents to be less risk-seeking (or more risk-neutral) than American and Japanese respondents seems unlikely in light of evidence that Indian and American respondents have comparable risk preferences in the financial domain (24, 28), that risk preferences are highly domain-specific (29), and that the existence of cross-national differences in risk preferences is domain-dependent (30).

Discussion

In summary, we find support for a model of human life valuation based on memory sampling and relative judgment (11–13). This model departs from the utility-based approaches typically encountered in psychology (8) and economics (19) in that it does not rely on stable, underlying value representations to explain valuation and choice, or on choice behavior to derive value functions. Instead, preferences concerning human fatalities emerge spontaneously from the distributions of sampled events and the relative nature of the evaluation process. The model successfully explains a variety of findings in the existing literature (Study 1) and makes unique predictions that were also supported (Studies 2 and 3). Our results suggest that reactions to fatalities are fundamentally relative and dependent on personal history and context. These studies also provide support for the general idea that the distributional properties of our environment are reflected in our cognitive system (31–34), and that some apparent cross-cultural differences may actually reflect cross-national variations in these distributions (35).

The diminishing sensitivity we show to losses of human life could have a number of functional purposes. For example, if our capacity to distinguish EADTs is limited, it might make sense for this ability to be concentrated (i.e., sharpest) around high-frequency, low-death toll events, as these tend to be packed together in time and in magnitude, making them otherwise very difficult to distinguish. Unfortunately, this comes at the cost of blurring our ability to distinguish low-frequency, high-death toll events. Alternatively, and given evidence that people have limited resources for coping with negative events (36), a diminishing sensitivity to increasing fatalities may protect us from being emotionally overwhelmed by large death tolls. The notion that there is an upper bound on the disutility (or shock) experienced by a person who observes a tragic event is captured in our theoretical account by the asymptotic nature of cumulative probability functions.

Regardless of its possible functions, understanding the factors responsible for our diminishing sensitivity to human fatalities has the potential to save many lives by increasing public reactions to distant disasters, epidemics, and genocides, while also helping policy makers correctly prioritize their efforts to prevent and respond to these humanitarian crises and other deadly risks.

Methods

Study 1A: Centre for Research on the Epidemiology of Disasters/EM-DAT Data. Data on the occurrence of disasters and their associated death tolls were obtained from the EM-DAT (www.em-dat.net) maintained by the Centre for Research on the Epidemiology of Disasters (CRED) at Université Catholique de Louvain in Brussels, Belgium. The EM-DAT is the only publicly available global database on the occurrences and impacts of natural and industrial disasters. The EM-DAT data were also used to produce the country-specific disaster-mortality distributions in Study 3 (see Fig. 4*A*). See *SI Text* for details of Study 1A data collection.

Study 1B: Google News Archive Data. Data on media attention to events involving human deaths (i.e., EADTs) were obtained by iterative search of the Google News Archives (GNA; see http://news.google.com/archivesearch) for news articles whose titles contained keywords related to losses (e.g., ''10 people died'') or ''gains'' (e.g., ''10 people survived'') in human lives. For each search, the number of relevant articles returned (i.e., the number of ''hits'') was recorded, thus providing a measure of the total media attention allocated to events associated with a given loss (or gain) in human lives. See *SI Text* for details of Study 1B data collection.

Study 1C: Recalled EADTs. Data on recalled EADTs were obtained by administering a survey that asked respondents to recall events involving human deaths. We then repeatedly sampled these events to estimate the average frequency and cumulative probability distribution of recalled EADTs. See *SI Text* for details of Study 1C data collection.

This study and all future studies reported in this paper were approved by Princeton University's Institutional Review Panel for Human Subjects and/or the University of Oregon's Office for Protection of Human Subjects (in Study 3, we also obtained approval from the overseas institutions where we administered surveys). Written or oral consent was also obtained from all our participants.

Study 2: Experimental Manipulation of Observed EADTs. Participants were 157 adults (56% female; age range, 18–77 y) recruited from a large shopping mall in the northeastern United States, who were paid for their participation.

The entire experiment was conducted using a paper questionnaire that participants read and completed on their own. The questionnaire's cover page contained written instructions and the final page asked participants to report various demographic characteristics. The manipulation and the decision scenario were presented on separate pages (with the former directly preceding the latter).

To manipulate the distribution of comparison EADTs from which they would sample, participants were alternately assigned to one of three conditions (concave distribution, S-shaped distribution, and control). Participants in the two treatment groups were first presented with eight randomly ordered single-sentence descriptions of disasters and accidents whose death tolls ranged from two to 1,000 (see Table S2). The top of the page containing the manipulation explained that these events represented a random sample of all of the natural and industrial disasters that had occurred in the past year (in reality, these events were all fictional). Participants were instructed to first read all the event descriptions carefully, and then to indicate how each event made them feel, using a 10-point rating scale that ranged from ''neutral'' to ''very negative.'' The purpose of this exercise was to provide a subtle channel through which participants would encode these EADTs. The two distributions of EADTs were selected so that, if sampled, they would either increase or decrease the distance (relative to baseline), in percentile-ranks, between the two key magnitudes that participants had to evaluate in the subsequent decision scenario (400 deaths and 600 deaths; as discussed later).

Participants in the concave distribution condition were presented with eight EADTs that, if used as comparison magnitudes, would decrease the distance between the percentile-ranks of the two options presented in the subsequent decision scenario, thereby making the risky option more appealing (according to the model we are testing). Specifically, these EADTs were all selected to be either less than 400 or greater than 600, thus reducing the difference in ranks between the two target EADTs (Fig. 3*A*). Participants in the S-shaped distribution condition were presented with eight EADTs that, if used as comparison magnitudes, would increase the distance between the percentile-ranks of the two options presented in the subsequent decision scenario, thereby reducing the appeal of the risky option (according to the model we are testing). Specifically, these EADTs were selected so that many of them fell between 400 and 600, thus increasing the difference in ranks between the two target EADTs (Fig. 3*A*).

Following this manipulation, participants in both treatment conditions advanced to the decision scenario, which was presented on the next page of the questionnaire. To establish a baseline, a third, control group of participants was not exposed to either distribution manipulation and instead advanced directly to the decision scenario.

The decision scenario used to measure risk preferences was a modification of the ''loss frame'' version of the ''Asian disease'' problem (2, 8), which asks respondents to imagine a choice between two programs for combating the outbreak of a disease that is expected to kill 600 people. One program (the sure option) leads to 400 deaths with certainty, whereas the other (the risky option) offers a one-third chance that no one will die, but a two-thirds chance that all 600 will die. Our scenario referred to a real disease (the West Nile virus) rather than an imaginary ''Asian disease,'' but this was the only notable modification. One respondent did not provide a response to the disease scenario and was therefore excluded from the analysis. We also excluded data from five respondents who recognized the disease scenario (all five were university students). The exact scenario is presented in Fig. S1.

Although both programs have the same expected value (both lead to 400 deaths, on average), previous research (2) found that the majority (78%) of respondents preferred the risky option when these choices were framed in terms of lives lost, implying a diminishing sensitivity (i.e., marginally diminishing disutility) for the loss of human lives (a result that we replicated with our control group; see Fig. 3*B*). In particular, people seem to perceive the disutility associated with a sure loss of 400 lives to be greater than the expected disutility associated with a two-thirds probability of 600 deaths. In other words, for most respondents, the subjective shock produced by the knowledge that 400 people will die (or have died) seems to be greater than two thirds of the subjective shock associated with 600 deaths.

However, we expected that the share of respondents preferring the risky option in this scenario would decrease as the subjective distance (in terms of psychological shock) between 400 deaths and 600 deaths increased. Similarly, the share of respondents preferring the risky option was expected to increase as this subjective distance decreased. According to the model we are testing, participants assigned to the concave distribution condition would therefore be more likely to choose the risky alternative than participants assigned to the S-shaped distribution condition, which is what we found (Fig. 3*B*).

We should note that, strictly speaking, our manipulation varied not only the relative ranks of the target EADTs within the distributions that participants were exposed to, but also the means of these distributions. One might therefore attribute our results to a difference in means rather than ranks. This alternative explanation, however, is unlikely for at least three reasons. First, it seems implausible that participants calculated the means of these EADT distributions, let alone used this information to inform their subsequent choices. Second, research shows that rank-based accounts are better able to explain the effects of observed distributions on people's attitudes and preferences compared with mean-based accounts (37). Finally, an account based on differences between means fails to explain the results of Study 3 as, for example, Indonesia experienced a much higher mean EADT (1,813 deaths, on average, between 2003 and 2007) than the other three countries (including India, which was second with an average of 153 deaths), yet Indonesian respondents were neither the least, nor the most, likely to choose the risky alternative.

Study 3: Cross-National Differences in Risk-Preferences. Data on EADT frequencies in each country were obtained from the EM-DAT (www.em-dat.net), as described earlier (Study 1A in *Methods*). Responses to the risky-decision scenario were collected from 249 university students in four countries. The Indian sample consisted of students at the University of Delhi ($n = 60$; 55% female). The Indonesian sample consisted of students at Pelita Harapan University and Ciputra University ($n = 56$; 73% female). The Japanese sample consisted of students at Rissyo University ($n = 62$; 65% female). The American sample consisted of students at the University of Oregon ($n = 71$; 61% female).

The procedure was similar across countries: paper questionnaires containing the decision scenario were administered to respondents who were recruited in their classrooms (some American respondents instead completed a Web-based version of the questionnaire in exchange for course credit). The questionnaire's cover page contained written instructions and the final page asked participants to report various demographic characteristics, including whether they had lived their entire life in the country where they were recruited. Those respondents who reported having lived outside the country in which the data were collected ($n = 24$) were excluded from the analysis.

The decision scenario was similar to the one used in Study 2 except for three modifications: First, the disease in question was not a specific epidemic but was simply described as ''an unusual disease.'' Second, the disease was specifically described as affecting the country in which the data were collected (this was the only feature that differed across countries). Third, the outbreak in the scenario was expected to kill 40 people, and the sure option led to 20 deaths with certainty whereas the risky option offered a 50% probability that all 40 would die and a 50% chance that none would die. These numbers were chosen because they represented values for which the four countries differed strongly in terms of EADT percentile-ranks (Fig. 4*A*) and because a 50% probability is easier to understand and compute than the one-third and two-thirds probabilities used in the original scenario. A 50% probability can, for example, be conceptualized as equivalent to flipping a fair coin to determine the outcome of the risky option. As before, the expected value, in terms of lives lost, was the same for both programs (20 deaths on average). English, Indonesian, and Japanese versions of the scenario are presented in Figs. S2–S4.

PSYCHOLOGICAL AND COGNITIVE SCIENCES

The survey was written in English for the American and Indian samples and translated for the Japanese and Indonesian samples. Two steps were taken to ensure that translations were as similar as possible in terms of the information conveyed: First, the survey was translated into Japanese and Indonesian by one pair of translators, then back-translated into English by a separate pair of translators. This was done to identify any meaningful distortions produced by the translation process. In addition, the three versions (English, Japanese, and Indonesian) of the survey were iteratively modified to accommodate each language's unique constraints until they converged on a shared meaning. These steps were repeated until we were satisfied that the three versions were as semantically similar as possible, and that they properly communicated the scenario.

ACKNOWLEDGMENTS. We thank N. Chater, Y.C. Lim, D.M. Oppenheimer, D. Purves, P. Slovic, A.B. Sussman, and two anonymous reviewers for helpful suggestions; C. Hsu and N. Oosterhof for help with GNA searches; B. Barooah, M. Das, A. Deshpande, N. Gupta, M. Jain, J.V. Meenakshi, A. Nath, R. Ray, and the University of Delhi for help with data collection in India; S. Budiyuwono, D.K. Halim, I. Kurniawan, J.E. Kurniawan, J.L. Setiawan, K. Tanjung, Universitas Ciputra, and Universitas Pelita Harapan for help with data collection in Indonesia; S. Matsubara and Rissyo University for help with data collection in Japan; and J. Harris, H. Ishii, and J. Zemla for help with data collection in the United States. C.Y.O. was supported by funding from Princeton University's Global Network on Inequality and a Princeton University Centre for Human Values Laurance S. Rockefeller Graduate Prize Fellowship. N.S. was supported by a Calvin Reed Smith Research Grant and a Merle King Smith Scholar Award.

1. Slovic P (2007) ''If I look at the mass I will never act'': Psychic numbing and genocide. *Judgment Decis Making* 2:79–95.
2. Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211:453–458.
3. Fetherstonhaugh D, Slovic P, Johnson SM, Friedrich J (1997) Insensitivity to the value of human life: A study of psychophysical numbing. *J Risk Uncertain* 14:283–300.
4. Friedrich J, et al. (1999) Psychophysical numbing: When lives are valued less as the lives at risk increase. *J Consum Psychol* 8:277–299.
5. de Blaeij A, Florax RJGM, Rietveld P, Verhoef E (2003) The value of statistical life in road safety: A meta-analysis. *Accid Anal Prev* 35:973–986.
6. Viscusi WK, Aldy JE (2003) The value of a statistical life: A critical review of market estimates throughout the world. *J Risk Uncertain* 27:5–76.
7. Strömberg D (2007) Natural disasters, economic development, and humanitarian aid. *J Econ Perspect* 21:199–222.
8. Kahneman D, Tversky A (2000) *Choices, Values and Frames* (Cambridge Univ Press, New York).
9. Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–291.
10. Galanter E (1962) The direct measurement of utility and subjective probability. *Am J Psychol* 75:208–220.
11. Stewart N, Chater N, Brown GDA (2006) Decision by sampling. *Cogn Psychol* 53:1–26.
12. Stewart N (2009) Decision by sampling: The role of the decision environment in risky choice. *Q J Exp Psychol* 62:1041–1062.
13. Howe CQ, Lotto RB, Purves D (2006) Comparison of Bayesian and empirical ranking approaches to visual perception. *J Theory Biol* 241:866–875.
14. Long F, Yang Z, Purves D (2006) Spectral statistics in natural scenes predict hue, saturation, and brightness. *Proc Natl Acad Sci USA* 103:6013–6018.
15. Yang Z, Purves D (2004) The statistical structure of natural light patterns determines perceived light intensity. *Proc Natl Acad Sci USA* 101:8745–8750.
16. Sedlmeier P, Betsch T (2002) *Etc. Frequency Processing and Cognition* (Oxford Univ Press, New York).
17. Stewart N, Brown GDA, Chater N (2005) Absolute identification by relative judgment. *Psychol Rev* 112:881–911.
18. Stewart N, Chater N, Stott HP, Reimers S (2003) Prospect relativity: How choice options influence decision under risk. *J Exp Psychol Gen* 132:23–46.
19. Wakker PP (2008) Explaining the characteristics of the power (CRRA) utility family. *Health Econ* 17:1329–1344.
20. Chater N, Brown GDA (1999) Scale-invariance as a unifying psychological principle. *Cognition* 69:B17–B24.
21. Enquist BJ, Niklas KJ (2001) Invariant scaling relations across tree-dominated communities. *Nature* 401:655–660.
22. Chater N, Brown GDA (2008) From universal laws of cognition to specific cognitive models. *Cognit Sci* 32:36–67.
23. Reed WJ, Hughes BD (2002) From gene families and genera to incomes and internet file sizes: Why power laws are so common in nature. *Phys Rev E* 66:067103.
24. Holt CA, Laury SK (2002) Risk aversion and incentive effects. *Am Econ Rev* 92:1644–1655.
25. Guria J, Leung J, Jones-Lee M, Loomes G (2005) The willingness to accept value of statistical life relative to the willingness to pay value: Evidence and policy implications. *Environ Resour Econ* 32:113–127.
26. McDaniels TL (1992) Reference points, loss aversion, and contingent values for auto safety. *J Risk Uncertain* 5:187–200.
27. Greene RL (1986) Sources of recency effects in free recall. *Psychol Bull* 99:221–228.
28. Binswanger HP (1980) Attitudes toward risk: Experimental measurement in rural India. *Am J Agricultural Econ* 62:395–407.
29. Blais AR, Weber EU (2006) A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgment Decis Making* 1:33–47.
30. Hsee CK, Weber EU (1999) Cross-national differences in risk preference and lay predictions. *J Behav Decis Making* 12:165–179.
31. Shepard RN (2001) Perceptual-cognitive universals as reflections of the world. *Behav Brain Sci* 24:581–601.
32. Anderson JR (1990) *The Adaptive Character of Thought* (Lawrence Erlbaum Associates, Hillsdale, NJ).
33. Purves D, Lotto RB (2003) *Why We See What We Do: An Empirical Theory of Vision* (Sinauer, Sunderland, MA).
34. Howe CQ, Purves D (2005) *Perceiving Geometry: Geometrical Illusions Explained by Natural Scene Statistics* (Springer, New York).
35. Miyamoto Y, Nisbett RE, Masuda T (2006) Culture and the physical environment: Holistic versus analytic perceptual affordances. *Psychol Sci* 17:113–119.
36. Linville PW, Fischer GW (1991) Preferences for separating or combining events. *J Pers Soc Psychol* 60:5–23.
37. Brown GDA, Gardner J, Oswald AJ, Qian J (2008) Does wage rank affect employees' well-being? *Industrial Relations* 47:355–389.

Supporting Information

Olivola and Sagara 10.1073/pnas.0908980106

SI Text

Formal Theoretical Framework. The theoretical framework that underlies our account of the diminishing sensitivity to human fatalities proposes that people evaluate a target "event-associated death-toll" (EADT) by comparing it with other EADTs sampled from memory. The final subjective value assigned to the target EADT is simply the proportion of pair-wise comparisons in which it dominates or ties.

Formally, this can be expressed as follows: consider a standard ranking system that ranks all the EADTs in a set by assigning a rank of 1 to the largest EADT, the same rank-values to identical EADTs, and progressively larger ranks to smaller EADTs. Within this ranking system, if the rank of the target EADT is r_t when it is compared with the n_s other EADTs that were sampled, its subjective value ψ_t will be:

$$\psi_t = \frac{n_s + 2 - r_t}{n_s + 1} \quad \textbf{[S1]}$$

Thus, a target EADT is considered large if it ranks above most sampled EADTs (small if it ranks below most of them), regardless of its absolute magnitude.

The proportion of pair-wise comparisons in which a target EADT (x_t) dominates or ties is equivalent to its percentile-rank, which is the proportion of sampled EADTs that are smaller than or equal to the target. Equivalently, this can be expressed as the probability that x_t is larger than or equal to a randomly drawn comparison EADT: $p(x_t \geq X_s)$. Mathematically, this is represented by the cumulative distribution function $F(x_t)$.

The sampling process implies that the value assigned to a target EADT will be determined by the distribution of comparison EADTs from which a person can draw, which will be a function of the EADTs (s)he has previously observed. If we assume, for simplification, that people draw uniformly random samples from the entire set of events they have observed, the sample of events under consideration will be a representative subset of all events in memory, and an EADT's expected percentile-rank within the sample will be equal to its percentile-rank within the entire population of observed events (1, 2). This implies that the psychophysical function relating an EADT's magnitude to its subjective evaluation can be approximated by the cumulative distribution function of all relevant EADTs that one has observed.

Mathematically, the cumulative distribution function is obtained by integrating the probability density function. The probability density function is simply the frequency distribution function divided by the total number of EADTs in the sampled population:

$$f(x) = \frac{freq.(x)}{N} \quad \textbf{[S2]}$$

Because EADTs seem to roughly follow a power-law distribution (see Study 1), their frequency distribution is reasonably well approximated by a power function:

$$freq.(x) = qx^{\alpha} \quad \textbf{[S3]}$$

with power α. Combining Eqs. **S2** and **S3**, we obtain the probability density function:

$$f(x) = bx^{\alpha}\left(\text{where } b = \frac{q}{N}\right) \quad \textbf{[S4]}$$

By integrating Eq. **S4**, we obtain the cumulative distribution function:

$$F(x) = b\left(\frac{x^{1+\alpha}}{1+\alpha}\right) \quad \textbf{[S5]}$$

As b is simply a normalizing constant with no real empirical meaning, we can simplify Eq. **S5** by setting $b = 1$ while still conserving the main features of the relationship between an EADT's magnitude (x) and its cumulative frequency (or percentile):

$$\psi(x) = \frac{x^{1+\alpha}}{1+\alpha} \quad \textbf{[S6]}$$

According to our account of the way people evaluate human fatalities, Eq. **S6** approximately describes the relationship between an event's death toll and its associated disutility or shock value. In fact, a common assumption in economics is that preferences follow a utility function characterized by constant relative risk aversion (CRRA) (3, 4). CRRA utility functions often take the form:

$$U(x) = \frac{x^{1-\gamma}}{1-\gamma} \quad \textbf{[S7]}$$

where γ describes the degree of concavity of the utility function. [Note that when $\gamma = 1$, $U(x) = \ln(x)$.] When x represents a desirable gain, the CRRA parameter γ describes the degree of relative risk aversion for an individual. Here, however, x represents the number of human deaths and $U(x)$ the disutility (i.e., negative utility) associated with this undesirable loss, so the relationship is reversed (5–7): γ describes the degree of relative risk-seeking for an individual (i.e., the tendency to prefer risky choice options). An individual is risk-averse if $\gamma < 0$, risk-neutral if $\gamma = 0$, and risk-seeking if $\gamma > 0$.

Notice that Eq. **S7** is obtained from Eq. **S6** simply by setting α as $-\gamma$. Thus, according to our account, the curvature (γ) of the disutility function for human deaths is approximately equal to the negative value of the power parameter (α) that governs the distribution of EADTs.

A potential issue with the framework we present concerns the treatment of ties between EADTs. In the current account, a target EADT's disutility is equal to the proportion of pair-wise comparisons in which it dominates or ties. One advantage of having ties counted in favor of the target EADT is that it allows individual deaths ($X = 1$) to produce relatively large amounts of disutility, in line with the empirical evidence (8). Conversely, this specification could also generate some counterintuitive predictions in certain cases, notably when the target EADT is equal to comparison EADTs in the middle of the sampled range. For example, a target event involving 5 deaths would, in the current framework, obtain a disutility value of 0.8 when compared with the set [1, 5, 5, 10], whereas 0.5 might seem like a more intuitive value as the target event falls in the middle of the distribution in this case. However, the likelihood of drawing a sample of this sort is extremely small, given the distribution of EADTs we observe. In these power-law-like distributions, only the very lowest EADTs have a nonnegligible probability of being sampled more than once (or at all, for that matter). Therefore, it is highly unlikely that ties will occur in the middle of the sampled range (ties occurring in the very beginning or very end of the sampled range do not pose such an issue). Nevertheless, we examined how different specifications concerning the treatment of ties might impact our results. We found that our predictions and results are qualitatively unchanged if we adopt specifications that separate ties from inequalities between EADTs. In fact, even adopting a "strictly-greater-than" definition of percentile ranks, whereby a target EADT's disutility is equal to the proportion of pair-wise comparisons that it strictly dominates,* was found to have a negligible impact on the shapes of the cumulative probability distributions reported in this article.

Methods

Study 1A: Centre for Research on the Epidemiology of Disasters (CRED)/ Emergency Events Database (EM-DAT) Data. Data on the occurrence of disasters and their associated death tolls were obtained from the EM-DAT (www.em-dat.net) maintained by CRED at Université Catholique de Louvain in Brussels, Belgium. The EM-DAT is the only publicly available global database on the occurrences and impacts of natural and industrial disasters. EM-DAT data are compiled from a variety of reliable sources, including United Nations agencies, governments, nongovernmental organizations, insurance companies, research institutes, and press agencies. When new data are added to the dataset, they undergo a validation process to minimize error before they become public.

The types of disasters included in the EM-DAT dataset mostly fall into two broad groups: natural disasters, which include droughts, earthquakes, epidemics, extreme temperatures, floods, insect infestations, slides, volcanoes, waves/surges, wildfires, and windstorms; and industrial disasters, which include industrial accidents, miscellaneous accidents, and transport accidents. Only events that meet specific criteria are classified as disasters and recorded in the EM-DAT dataset: an event is added to the dataset if at least one of the following conditions is met: (*i*) at least 10 people were killed; (*ii*) at least 100 people were affected, injured, or homeless; (*iii*) considerable damage was incurred; (*iv*) a declaration of a state of emergency and/or an appeal for international assistance was made; or (*v*) the event is considered noteworthy for some other reason. One consequence of the first selection criterion is that the frequency of events involving fewer than 10 deaths is underestimated. Conversely, low-impact events are generally less salient and less likely to be reported in the media than high-impact events. As a result, they may be rarely observed, encoded in memory, or sampled during the evaluation process.

In the raw data we consider, the unit of analysis is a disaster. The EM-DAT dataset contains a number of useful variables associated with each disaster, including the country or countries affected, the number of people killed, and the starting and ending dates of the event.

Starting in 2003, the CRED decided to alter the process for entering disasters into the EM-DAT database, in an effort to improve its methodology. As a result of this shift, there is a discontinuity in the way disasters are compiled before and after 2003. We therefore only considered disasters that occurred between 2003 and the time the data were downloaded.

The data were downloaded on October 24, 2007, from the EM-DAT Web site. Only disasters that caused at least one (human) death were considered in our analyses. Disasters for which there was a mistake in the recording of the start date and/or end date (e.g., the recorded end date occurred before the recorded start date) were removed. Disasters for which no starting month or ending month was available were removed. Disasters for which the classification year did not correspond to the starting and/or ending year were removed. As the deaths associated with a disaster that unfolds over an extended period are spread out in time, it is unclear whether this loss of life is coded as a single, high-mortality event or as a series of multiple, low-mortality events. To avoid this potential ambiguity, we considered only events that occurred over a period of 10 d or fewer (events lasting >10 d represented 12% of the sample). Finally, when the same disaster affected multiple countries, the death toll was aggregated across those countries and the resulting total was coded as a single event. Overall, we selected 2,282 individual events.

The EM-DAT data were also used to produce the country-specific disaster-mortality distributions in Study 3 (see Fig. 4*A* and Study 3 in *Methods*). For that analysis, however, the death tolls were not aggregated when multiple countries were affected by the same disaster. In addition, only disasters affecting India, Indonesia, Japan, and/or the United States were considered, and separate analyses were carried out for each country. The data selection process was identical in every other respect. Of the 2,282 disasters that were selected (as detailed earlier), 153 affected India, 98 affected Indonesia, 28 affected Japan, and 86 affected the United States.

Study 1B: Google News Archives (GNA) Data. Data on media attention to events involving human deaths (i.e., EADTs) were obtained by iteratively searching the GNA (http://news.google.com/archivesearch). The GNA allows users to search for news articles (using key words) across a large collection of historical archives from many countries, including major newspapers and magazines, news archives, and legal archives. GNA searches draw from a large variety of different sources, and include content that is publicly accessible as well as content that requires a fee.

We searched the GNA for news articles whose titles contained keywords related to losses (e.g., "10 people died") or gains (e.g., "10 people survived") in human lives. For each search, the number of relevant articles returned (i.e., the number of hits) was recorded, thus providing a measure of the total media attention allocated to events associated with a given loss (or gain) in human lives. The search process was limited to English-language pages only and to articles published between 2000 and 2007 (all searches were conducted in 2008). To minimize the number of articles about nonhuman losses (or gains), stories were not counted if the keyword "animal" appeared anywhere in the article. An iterative search process was carried out by an automated search algorithm, which sequentially implemented searches and recorded the number of hits produced by each one.

Two general types of searches were carried out:

(I) "*X* [keyword]," where "*X*" represents the number of lives lost or saved and "[keyword]" represents the specific word used to signify a loss or gain. This search yielded articles with titles of the form "3 killed in car crash," in which no words appear between the number "*X*" and the keyword.

(II) "*X* * [keyword]," where the asterisk is used to signify any words appearing between the number "*X*" and the keyword. This search yielded articles with titles of the form "3 people killed in car crash." Because this approach also counts titles of the form "3 million killed" (thus yielding false alarms), the search was designed to ignore articles with titles containing the words "*X* hundred," "*X* thousand," or "*X* million."

Articles on events related to losses in human lives were counted using the following keywords: "die," "died," "dead," "deaths," "killed," "fatalities," "homicides," "murders," "murdered," and "massacred." Articles on events related to gains in human lives were counted using the following keywords: "saved," "rescued," and "survive." Keywords were adjusted to the singular form for $X = 1$ whenever appropriate (e.g., "deaths" was replaced with "death"). These key words were specifically chosen because test searches showed that they seemed to yield the largest ratio of correct hits (i.e., relevant stories) to false alarms (i.e., irrelevant stories). As GNA allows only a limited number of words to be used in each search, we were required to divide the search process into multiple search strings. However, keywords were grouped (using the "OR" operator) to produce the fewest strings possible. This yielded seven search strings for loss-events and two search strings for gain-events. These search strings are provided in Table S1.

Searches were carried out for every integer-value of *X* between 1 and 1,000. Beyond $X = 1{,}000$, numbers were sampled up to 1,000,000 in a different manner: for each order of magnitude, 10^m (with $m = 3, 4, 5$), the first 10 values were sampled in increments of 10^{m-1} (e.g., 1,100; 1,200; […]; 2,000), and the next 16 values were sampled in increments of $5 \times 10^{m-1}$ (e.g., 2,500; 3,000; […]; 10,000). This led to the selection of 78 salient integers: values yielding larger numbers of GNA hits, as, for high death tolls, news articles are much more likely to report approximate values (e.g., "3,000 dead following attack") than exact values (e.g., "3,147 dead following attack").

To account for hits produced by nonsalient values of *X*, we randomly sampled 10 integers between each of the 78 salient values. If an integer appeared more than once, one of its occurrences was replaced with another randomly sampled value from the same range (eight replacements were made in total). The resulting 780 additional integers were then used as search values for *X*. For each range, we calculated the average number of hits returned by its 10 nonsalient values (based on a given search string). We then multiplied this average by the size of the range [either $(10^{m-1} - 1)$ or $(5 \times 10^{m-1} - 1)$], to yield an estimate of the total number of hits contained within that range. The resulting 78 estimates were added to the 78 salient values to produce, for each search string, an approximation of the total number of hits that GNA would produce for all values of *X* between 1,001 and 1,000,000. We also conducted a few searches (by hand) using values beyond 1,000,000, but these failed to produce any relevant hits (even for salient integers), suggesting that few if any

*Notice that any reasonable specification that explicitly accounts for ties must fall somewhere between a "greater-than-or-equal-to" [$p(x_t \geq X_s)$] and a "strictly-greater-than" [$p(x_t > X_s)$] definition of percentile ranks, in terms of the subjective value it assigns to EADTs. The "strictly-greater-than" specification, in which ties do not contribute to disutility, therefore provides the most stringent test of robustness to the treatment of ties that we could use.

events reported in 2000 through 2007 were associated with more than 1 million deaths. We therefore decided to stop searching beyond this point.

The number of hits from each search string were added up, separately for losses and gains, to produce the total number of hits associated with each value of $X \leq 1{,}000$, as well as the estimated total number of hits for larger values of X (up to 1,000,000). For events involving lives lost, we counted a total of 119,769 search hits ($X \leq 1{,}000$) and estimated another 2,776 hits ($1{,}000 < X$). For events involving lives saved, we counted a total of 3,160 search hits ($X \leq 1{,}000$) and estimated another 144 hits ($1{,}000 < X$).

It is worth noting that the number of hits obtained for $X = 1$ was likely underestimated. Articles in which a single person dies are more likely to have titles such as ''man dies in car accident'' than ''1 man dies in car accident.'' In addition, the deaths of famous persons, although also constituting an individual death, are likely to be missed as the relevant article titles usually refer to the person by name, without a quantity indicator. However, single-death events are probably less memorable, on average, than higher death toll events, and thus less likely to be sampled during the evaluation process. It is also possible that certain types of single-death events (e.g., the death of a friend, family member, or celebrity) are categorized differently from events typically encountered in the news, which involve the deaths of strangers. These types of single-death events might not, therefore, be sampled in the evaluation process.

Study 1C: Recalled EADTs. Data on recalled EADTs were obtained by administering a survey that asked respondents to recall events involving human deaths. We then repeatedly sampled these events to estimate the average frequency and cumulative probability distribution of recalled EADTs.

Respondents were 160 university students in the United States (43% female) who participated for course credit.

The survey asked respondents to recall specific nonfictional events involving human deaths, and to report the first eight examples that came to mind. They were encouraged to use real events that they had previously heard about, read about, or seen on television, as long as these events had occurred in their lifetime. They were asked to provide a brief description of each event and their best estimate of the number of deaths involved (as a single number rather than a range of values). Respondents were given the option of completing an alternate questionnaire of similar length (about recalling temperatures) if they felt uncomfortable with the survey's topic. Only one respondent requested this option. Another respondent who was noticeably distracted was also removed from the sample. Many respondents reported fewer than eight events and, of those reported, some events were excluded from the analysis because they occurred before the respondent's lifetime, referred to general causes of death rather than specific events (e.g., ''all deaths from cancer''), or referred to nonhuman deaths. Finally, events were excluded if their estimated death tolls were missing, equal to zero, or reported as a range of numbers or some other ambiguous indicator of quantity (e.g., ''thousands''). Using data from those respondents who recalled at least 6 valid events ($n = 108$), we randomly sampled one recalled event from each person and calculated the frequency and percentile-rank distributions based on these 108 sampled death tolls. This sampling process was repeated 1,000 times (with replacement), and the resulting output was used to calculate mean frequencies and percentiles for each death toll (Fig. 1*C*). Including all participants and events with unambiguous, nonzero death tolls into our analyses produced qualitatively similar results.

It should be noted that explicitly asking participants to recall events involving human deaths could have generated a memory search process that differs somewhat from how they might spontaneously recall EADTs when trying to evaluate a target event. In particular, our task may have led them to focus heavily on the loss of a close other (thereby producing many single-person death tolls) and on extremely large death tolls. They might, for example, have considered these two classes of events to be especially worth reporting, even if they initially sampled more broadly. This tendency could have been further reinforced by the instructions, which required respondents to not only recall the number of deaths associated with each event but also to provide brief descriptions. This may help to explain why the distribution of recalled EADTs differs somewhat from the other two distributions we obtained in Study 1 (see Fig. 1). Despite this potential limitation, however, the distribution of recalled EADTs still makes qualitatively similar predictions concerning sensitivity to human fatalities and risk preferences concerning human losses.

1. Stewart N, Chater N, Brown GDA (2006) Decision by sampling. *Cogn Psychol* 53:1–26.
2. Stewart N (2009) Decision by sampling: The role of the decision environment in risky choice. *Q J Exp Psychol* 62:1041–1062.
3. Wakker PP (2008) Explaining the characteristics of the power (CRRA) utility family. *Health Econ* 17:1329–1344.
4. Holt CA, Laury SK (2002) Risk aversion and incentive effects. *Am Econ Rev* 92:1644–1655.
5. Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211:453–458.
6. Kahneman D, Tversky A (2000) *Choices, Values and Frames* (Cambridge Univ Press, New York).
7. Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–291.
8. Slovic P (2007) ''If I look at the mass I will never act'': Psychic numbing and genocide. *Judgment Decis Making* 2:79–95.

Imagine that the U.S. is preparing for an outbreak of West Nile virus, which is expected to kill 600 people. There are two alternative programs. If Program A is adopted, 400 people will die. If Program B is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

Which do you prefer, Program A or Program B? (please circle one)

Program A Program B

Fig. S1. The decision scenario presented in Study 2.

English Version:

Imagine that [*India* / *the U.S.*] is preparing for the outbreak of an unusual disease, which is expected to kill 40 people. There are two alternative programs to combat the disease:

- If Program A is adopted, 20 people will die.
- If Program B is adopted, there is a 50% probability that nobody will die and a 50% probability that 40 people will die.

Which would you choose, Program A or Program B?
Please circle one Program to indicate your choice.

Program A	**Program B**
100% probability of 20 deaths	50% probability of 40 deaths
	50% probability of 0 deaths

Fig. S2. The English-language version of the decision scenario presented in Study 3 (to American and Indian respondents).

Indonesian Version:

Bayangkan Indonesia sedang bersiaga untuk penyakit aneh yang akan menjalar dengan cepat dan diperkirakan akan menyebabkan kematian sejumlah 40 orang. Ada dua program alternatif untuk memberantas penyakit:

- Jika Program A dipilih, 20 orang akan meninggal.
- Jika Program B dipilih, ada 50% kemungkinan tidak ada yang meninggal, dan 50% kemungkinan 40 orang akan meninggal.

Program mana yang akan anda pilih, Program A atau Program B?
Harap lingkari jawaban anda

Program A	**Program B**
100% kemungkinan 20 kematian	50% kemungkinan 40 kematian 50% kemungkinan 0 kematian

Fig. S3. The Indonesian version of the decision scenario presented in Study 3.

Japanese Version:

40人の亡くなる可能性のある特異疾病の急激な蔓延を予想し、日本が対策を立てていると想像してください。この疾病への対抗策として二つのプランがあります。

- プランAが実行された場合、20人が亡くなります。

- プランBが実行された場合、死者を出さない確立が５０％、40人が亡くなる確立が５０％あります。

あなたはどちらのプランを選びますか。
以下のプランのひとつを丸で囲んで選んでください。

プランＡ	プランＢ
１００％の確立の20人の死者	５０％の確立の40人の死者
	５０％の確立の ０人の死者

Fig. S4. The Japanese version of the decision scenario presented in Study 3.

Table S1. GNA search strings (Study 1*B*)

Search no.		String
	Losses (lives lost)	
1		-animal intitle:″X (die OR dead OR died OR deaths OR killed OR fatalities)″
2		-animal intitle:″X (homicides OR murders OR murdered OR massacred)″
3		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (die OR dead)″
4		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (died OR deaths)″
5		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (killed OR fatalities)″
6		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (homicides OR murders)″
7		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (murdered OR massacred)″
	Gains (lives saved)	
1		-animal intitle:″X (saved OR rescued OR survive)″
2		-animal -intitle:″X hundred″ -intitle:″X thousand″ -intitle:″X million″ intitle:″X * (saved OR rescued)″

X represents an integer value. The keywords were adjusted to the singular form for $X = 1$ whenever appropriate (e.g., ″deaths″ was replaced with ″death″).

Table S2. Experimental manipulation used in Study 2

Scenario	How does this event make you feel? Please circle a number for each event, indicating how it makes you feel.
776 people died following an earthquake in Central Asia.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
A week-long heat wave in Mexico led to 9 [*283*] deaths.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
Mudslides in Guyana left 175 [*475*] dead.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
An industrial chemical explosion killed 39 [*426*] people in China.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
A typhoon in the Pacific killed 1,000 people.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
A flash flood in Bangladesh killed 283 [*519*] people.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
2 people were killed in a car accident in Poland.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative
Continuous droughts in Niger were responsible for 94 [*448*] deaths.	1—2—3—4—5—6—7—8—9—10 Neutral Negative Very negative

Numbers outside the brackets represent the death toll magnitudes (i.e., EADTs) that were presented to participants in the concave distribution condition. Italicized numbers inside the brackets represent the death toll magnitudes that were presented to participants in the S-shaped distribution condition. Events without numbers in brackets were those for which the death toll was the same across conditions. In both treatment conditions, the death toll numbers that participants saw (but not the rest of the sentence) were in bold (but not in brackets nor italicized).