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).