Risk Analysis, Vol. 40, No. 2, 2020 DOI: 10.1111/risa.13396

The Psychophysics of Terror Attack Casualty Counts

Matt Baucum and Richard John*

In communicating the risk that terror attacks pose to the public, government agencies and other organizations must understand which characteristics of an attack contribute to the public's perception of its severity. An attack's casualty count is one of the most commonly used metrics of a terror attack's severity, yet it is unclear whether the public responds to information about casualty count when forming affective and cognitive reactions to terror attacks. This study sought to characterize the "psychophysical function" relating terror attack casualty counts to the severity of the affective and cognitive reactions they elicit. We recruited $n = 684$ Mechanical Turk participants to read a realistic vignette depicting either a biological or radiological terror attack, whose death toll ranged from 20 to 50,000, and rated their levels of fear and anger along with the attack's severity. Even when controlling for the perceived plausibility of the scenarios, participants' severity ratings of each attack were logarithmic with respect to casualty count, while ratings of fear and anger did not significantly depend on casualty count. These results were consistent across attack weapon (biological vs. radiological) and time horizon of the casualties (same-day or anticipated to occur over several years). These results complement past work on life loss valuation and highlight a potential bifurcation between the public's affective and cognitive evaluations of terror attacks.

KEY WORDS: Psychophysics; risk perception; subjective value; terrorism

1. INTRODUCTION

Many instances arise where agencies and organizations must effectively communicate information about past terror attacks to the public. The U.S. State Department may raise a foreign country's travel advisory level based on a string of recent terror events, and may need to justify this decision to prospective travelers. A large event venue may enhance its security screening procedures after a terror attack at a similar venue, and may need to explain this rationale to its patrons. Philanthropic organizations may campaign for humanitarian donations to a nation gripped by instability and wish to communicate the severity of the terror attacks that have been carried out.

Department of Psychology, University of Southern California, Los Angeles, CA, USA.

*Address correspondence to Richard John, University of Southern California, 3620 S McClintock Ave., Los Angeles, CA 90089, USA; tel: +1 (213-740-2203); richardj@usc.edu.

In these instances, effective public communication depends on an agency's or organization's ability to understand how the public interprets information about terror attacks. If a string of severe terror attacks leads to a change in transportation security or travel policy, agencies must be able to trust that the attack threat severity will be well-understood by the public. Of course, threat "severity" is a multifaceted construct in the domain of terrorism, but the most prominent and salient component is the number of lives that the attack claims. Intuitively, the perceived severity of a terror attack should monotonically increase with its casualty count, all else being equal—yet the precise nature of this relationship, and whether it obeys this prescriptive monotonicity, is an empirical question that has not been resolved. This study aims to build on previous work on public perceptions of death tolls from terror attacks to identify (1) the degree to which risk perceptions of terror attacks are impacted by the number of lives they

0272-4332/20/0100-0399$22.00/1

claim, (2) which psychological variables are most influenced by casualty counts, and (3) the functional form of these relationships.

1.1. Public Perceptions of Casualty Counts: Past Research

The dominant theoretical perspective relevant to this investigation is that subjective valuations of life loss generally follow a concave function, a phenomenon that has been termed "psychophysical numbing" (Fetherstonhaugh, Slovic, Johnson, & Friedrich, 1997; Friedrich & Dood, 2009; Friedrich et al., 1999; Slovic, 2007). Respondents in these studies placed greater value on saving a fixed number of lives in the context of smaller (as opposed to larger) death tolls, suggesting decreasing marginal value placed on each life as total casualty counts increase (note that we use the term *casualty* to imply deaths, rather than the sum of deaths and nonlethal injuries). This conclusion was also supported by previous work (Summers, Slovic, Hine, & Zuliani, 1994), which found that respondents perceived the magnitude of an armed conflict to be concavely related to the number of lives claimed.

Olivola and Sagara (2009) provide a theoretical basis for the empirical results described above, hypothesizing that the "value" associated with a particular casualty count is determined by its percentile rank among a random sample of remembered casualty counts from past events. Under this model, an individual's value function follows the same form as the cumulative distribution function of casualty counts that the individual has observed over time. The authors demonstrated that the distributions of disaster-related fatalities across multiple countries generally produced concave cumulative distribution functions, which in turn predicted concave value functions over life loss.

This notion of a concave disutility function over life loss aligns with other research on decision theory, in which concave utility functions over rewards (such as money) frequently arise (Bernoulli, 1954), especially in affectively rich environments (Hsee & Rottenstreich, 2004). Yet the present study is primarily interested in the *psychophysical* function governing the magnitude of individuals' reactions to events, rather than participants' measured evaluations of their disutility (Olivola & Sagara, 2009) or of the value of preventing the life loss altogether (Fetherstonhaugh et al., 1997; Friedrich & Dood, 2009; Friedrich et al., 1999). Note that, despite the use of the term "psychophysics" in past work on casualty perceptions ("psychophysical numbing" in Fetherstonhaugh et al., 1997; Friedrich & Dood, 2009; Friedrich et al., 1999), these studies used value-based tradeoffs to assess the subjective utility of life-saving interventions. In contrast, true psychophysical functions relate stimuli to the magnitudes of the psychological reactions they elicit, which is not necessarily equivalent to an individual's disutility functions yet still highly relevant to his or her reactions to catastrophic events.

One of the most fundamental principles in psychophysics is that the ability to detect a constant difference between sensory stimuli decreases as their absolute magnitude increases (Fechner, 1966; Weber, 1834); that is, the magnitude of the "just noticeable difference" (the smallest difference an individual can reliably detect) is proportional to the magnitude of the stimulus. Such a principle would suggest that the psychophysical function relating event casualties to individuals' cognitive and emotional reactions follows a concave form, similar to the concave value/disutility functions previously established for life loss. Yet as we discuss in Section 1.2, there are myriad reasons for questioning whether the same degree of concavity that governs value functions and psychophysical functions over physical stimuli should apply to the case of terror attacks.

1.2. Terrorism as a Possible Boundary Condition

Terror attacks, compared to other forms of life loss, possess unique features that may represent boundary conditions for the concavity observed in psychophysical functions and in value functions over life loss in other domains.

Compared to instances of disease or humanitarian crises (on which much of the loss-of-life literature is based), terror attacks are strictly adversarial in nature, and can thus signal to the public that future attacks may be more likely than previously thought (Marshall et al., 2007). This may be why terror attacks are judged as more dangerous than more mundane, higher-likelihood hazards (Sunstein, 2003), and why the public values the prevention of terrorism-related deaths far more than the prevention of natural disaster deaths (Viscusi, 2009).

Such evidence of terrorism's high affective salience may mean that terror attacks instill an even more severe insensitivity to human life than is observed in other domains. Research by Hsee and Rottenstreich (2004) suggests that individuals become less sensitive to differences in reward magnitude as the affective salience of the stimulus increases. It is possible that the fear and unpredictability associated with terror attacks makes distinctions in fatality counts (which, while not a reward stimulus, still carry intrinsic meaning) even less noticeable than in other domains. Empirical work on terrorism risk perception seems to lend preliminary support to this notion. Cui, Rosoff, and John (2016) found that, in a sequence of terror attack vignettes, attack frequency (and thus, overall death toll) did not influence respondents' negative affect, risk perception, or behavioral avoidance tendencies, while a study by Burns and Slovic (2010) found that changing a fictional attack's death toll from 0 to 495 did not influence respondents' risk perceptions of terrorism.

Complete insensitivity to life loss seems especially feasible when dealing with high-casualty terror attacks. The September 11 attacks claimed more than 3,000 lives, while unconventional terror methods (such as nuclear or biological terrorism) could feasibly claim far more. Yet Olivola and Sagara's (2009) model of casualty perceptions was largely based on events claiming less than 1,000 lives, and relied on estimated counts of higher casualty events due to their low prevalence in worldwide news coverage. It is thus unknown whether public perceptions of high-casualty terror attacks would exhibit *any* sensitivity to death tolls—that is, whether an individual would perceive an attack claiming 4,000 lives as any different than one claiming 10,000 lives, if both death tolls exceeded those of any other event he or she had previously encountered.

Keep in mind that we do not interpret these studies to mean that individuals will be insensitive to life loss due to terrorism; we actually hypothesize the opposite (see Section 1.3). We merely argue that the public's sensitivity to casualty count should be positively demonstrated for the specific case of terrorism, rather than merely assumed from past research.

1.3. Present Study

Our primary goal is to investigate whether individuals' perceptions of terror attacks are sensitive to the number of lives the attack claims, given terrorism's unique affective features and the null effect of casualty count seen in past research on terrorism. Support for this notion would not only bolster the literature on public reactions to life loss, but would better inform how government agencies communicate to the public about terror attack severity.

To robustly test this research question, our study involves the following:

- *Realistic terror attack vignettes*: We measure respondents' reactions to vivid, hypothetical terror attack scenarios, a common technique in research on terrorism risk perception (Burns & Slovic, 2010; Cui et al., 2016; Rosoff, John, & Prager, 2012; Rosoff, Siko, John, & Burns, 2013). Relying on carefully crafted scenarios, rather than on real events, allows us to randomly assign respondents to different casualty counts while holding all other attack features constant.
- *Unconventional terror attack methods*: Our scenarios describe terror attacks carried out with CBRN (chemical, biological, radiological, nuclear) methods. Although most terror attacks use "conventional" tactics such as explosives or firearms, "unconventional" CBRN attacks have a far greater potential for life loss. This allows us to test a greater range of casualty counts than vignettes involving firearms or explosives (which typically do not claim more than 100–200 lives).
- *Variation in casualty time horizon*: To test the robustness of any notable effects of casualty count, we administered scenarios that involved either immediate or delayed fatalities. In our scenarios, immediate fatalities occurred on the day of the attack, whereas delayed fatalities were described as *anticipated* over a 20-year time horizon, due to the health effects of the chemical/biological substance used in the attack (see Section 2). Public communication about terror attack casualties may involve fatalities that have already occurred, or may involve fatalities that are anticipated or counterfactual (e.g., an attack would have killed 30 people if it had not been prevented); thus, we sought to test whether actual versus hypothetical casualties differentially impacted individuals' reactions to an attack.
- *Variation of attack method*: We also sought to test our results' robustness against different mechanisms through which an attack might cause fatalities. Biological weapons represent a unique case within the realm of unconventional terrorism since they can be transmitted between individuals, unlike other attack methods whose victims must be present at the initial attack. We thus randomly assigned respondents' attack scenarios to describe a biological or chemical

attack. This allowed us to ensure that all respondents' scenarios involved hazardous substances, while testing for our results' robustness against the specific fatality mechanism (transmittable vs. nontransmittable agent).

- *Multiple dependent variables*: Individuals' reactions to hazards are determined by both cognitive and emotional factors (Loewenstein, Weber, Hsee, & Welch, 2001), and it is possible that cognitive and emotional appraisals of a terror attack will differentially depend on the attack's death toll. We thus measure respondents' affective reactions (self-reported fear and anger) and cognitive perceptions (perceived severity) for each attack vignette.

In line with past research on perceptions of life loss, we hypothesize that respondents' self-reported emotional and cognitive reactions will exhibit (concave) sensitivity to casualty count across all attack vignettes, regardless of attack method (chemical vs. biological) or casualty time horizon (immediate vs. delayed).

Table I. Demographic Characteristics by Scenario Condition

	Immediate Time Horizon		Delayed Time Horizon	
	Chemical ($n = 168$)	Biological ($n = 173$)	Chemical ($n = 182$)	Biological ($n = 161$)
Male	54.8%	51.4%	52.7%	46.6%
Female	44.0%	48.0%	46.7%	53.4%
Hispanic/Latino	9.5%	8.1%	8.8%	8.7%
Not Hispanic/Latino	90.5%	91.9%	91.2%	91.3%
Black	4.8%	6.9%	6.6%	9.3%
Caucasian	82.7%	75.1%	78.6%	78.9%
East Asian	5.4%	5.8%	7.7%	3.7%
South Asian	0.0%	4.0%	0.5%	1.9%
Multiracial	4.2%	5.8%	2.7%	3.1%
Other/no answer	2.9%	2.4%	3.9%	3.1%
Median age (years)	33	33	34	33

2. METHODS

2.1. Respondents

A total of 802 respondents were recruited from Amazon Mechanical Turk and paid $0.50 each to complete a brief survey on their reactions to a terror scenario. Of the original sample, we removed 108 (13.5%) who failed an attention check question (a four-option item asking which kind of weapon was used in the attack scenario) and 10 who failed to enter their city location, for a final sample size of $n = 684$. Table I shows the demographic breakdown of the sample for each combination of scenario type (chemical vs. biological) and time horizon (immediate vs. delayed). Note that racial categories with less than 1% representation in the sample are reported as "other."

2.2. Procedure

After affirming their consent to participate, respondents read a brief, hypothetical terror attack scenario. Before proceeding to the scenario, respondents indicated the name of the large city (with a population greater than 100,000) that they lived nearest, or a large city they were familiar with or had visited if they did not wish to provide information specific to their location. This value was then piped into the fictional terrorism vignette as the attack location; for example, if a participant entered "Dallas" as his or her nearest large city, the scenario would indicate "there has been a major terror attack in downtown Dallas." Each scenario spanned three screens, each of which allowed respondents to advance only after 10 seconds (which was deemed adequate time to read the page).

Scenarios were constructed as breaking news stories, describing a group of terrorists that drove through the downtown area of each respondent's chosen city and released an unknown gaseous substance that forced many people to seek medical treatment. Scenarios described either a chemical attack, involving a gaseous "blister agent" that damaged the lungs, or a biological attack, involving the bacteria *Yersinia pestis*, which also targeted the lungs. Furthermore, respondents were randomly assigned to one of two possible time frames for casualties from the attack, which were either immediate (i.e., had already occurred) or delayed (i.e., expected to occur in the future), resulting in four distinct scenarios. The immediate chemical and immediate biological scenarios stated that a specific number of casualties had already occurred due to the chemical/biological weapon. The delayed chemical scenario indicated that casualties are anticipated to occur over the next 20 years, as the result of the chemical weapon's delayed yet harmful effects on the respiratory system of exposed individuals. The delayed biological scenario suggested that the infection that was released would continue to spread to new victims, and was anticipated to claim further casualties over the next 20 years because of this. Finally, within each scenario, respondents were also randomly assigned to read that the attack had killed (or was expected to kill) either 20, 60, 200, 600, 1,800, 5,400, 16,400, or 50,000 people. These casualties counts were designed to ensure a sufficiently large range of values, and to ensure a relatively constant multiple between consecutive values; each casualty count value is roughly triple the next lowest value.

After reading the scenario, respondents completed a series of self-report items and demographic questions. Median completion time was six minutes, with an interquartile range of four minutes.

2.3. Materials

Respondents' emotional reactions to the vignettes were operationalized as their self-reported levels of fear and anger after reading their assigned scenario (see Lerner & Keltner, 2001). Two items were used to measure each emotion:

(1) Think about the fact that these kinds of attacks are possible, and that there are people who might try to commit an attack like the one just described. How [fearful/angry] does this make you feel? (0–10 scale, ranging from "not at all fearful" to "extremely fearful").
(2) I feel [scared/angry] about the threat of terrorism after reading about this possible terrorist attack (seven-point scale, "strongly disagree" to "strongly agree").

Respondents' cognitive reactions to the attack were operationalized as how severe they perceived the attack to be, as measured by the following items:

(1) On a scale of 1–10, how extensive do you think the consequences of this attack were? That is, how extensive was the overall damage that the attack caused? (0–10 scale, "contained/not at all extensive" to "incredibly widespread/extensive").
(2) On a scale of 0–10, how severe do you think this attack was? (0–10 scale, "not severe at all" to "extremely severe").

Finally, to control for differences in the perceived realism of each scenario and casualty condition, respondents rated the plausibility of the attack scenario on a nine-point scale that ranged from "extremely believable" to "extremely unbelievable." Note that the mean believability score was 2.82 out of 9 (between "moderately believable" and "very believable").

For each of the two-item scales for fear, anger, and severity, Cronbach's alpha values suggested sufficient internal consistency ($\alpha_{\text{fear}} = 0.79$, $\alpha_{\text{anger}} = 0.76$, $\alpha_{\text{severity}} = 0.90$), and each item was more highly correlated with its other same-construct item than with the other-construct items. Thus, each item's score was converted to a 0–10 scale, with the fear, anger, and severity items being averaged together to create composite scores for each construct. Perceived severity scale scores correlated at $r = 0.34$ with both the fear and anger scale scores, and the latter two being correlated $r = 0.62$.

Table II. Mean (*SD*) of Fear, Anger, and Severity Scores by (Expected) Casualty Count

(Expected) Casualties	Fear	Anger	Severity
20	6.79 (2.39)	7.47 (2.21)	6.31 (2.25)
60	7.11 (2.48)	7.89 (2.17)	6.79 (2.10)
200	7.16 (2.04)	6.76 (2.16)	6.91 (2.34)
600	6.91 (2.61)	7.55 (2.33)	7.58 (2.06)
1,800	6.37 (2.86)	7.17 (2.75)	7.62 (1.86)
5,400	6.76 (2.31)	7.63 (1.93)	7.93 (1.61)
16,400	6.83 (2.48)	7.40 (2.29)	8.15 (1.82)
50,000	6.77 (2.59)	7.66 (2.31)	8.44 (1.71)

3. RESULTS

3.1. Effect of Casualty Count

Table II and Fig. 1 present respondents' fear, anger, and severity scores across all assigned casualty counts, collapsed across scenario. Note that scores are plotted against (natural) log-transformed casualty counts, rather than raw casualty counts. Given that each assigned casualty count is roughly three times the next lowest value, the natural log-transformation ensures almost equal spacing between them. That is, for any two casualty counts x and $3x$, the distance between their log transformations is constant and does not depend on x, as given by $ln(3x) - ln(x) = [ln(3) + ln(x)] - ln(x) = ln(3)$. Note that perceptions of severity, but not fear or anger scores, demonstrate an approximately linear relationship with log-transformed casualty counts,

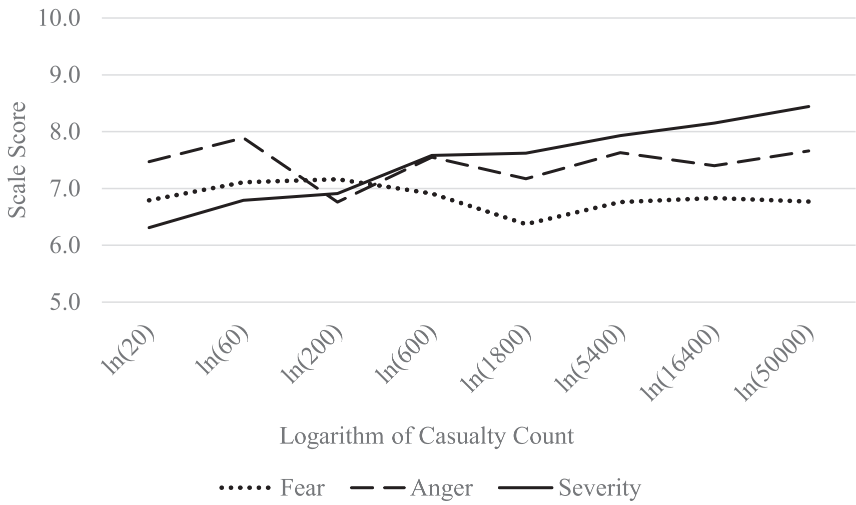

Fig. 1. Mean fear, anger, and severity scale scores by ln-casualty count.
Note: Casualty counts plotted on natural log-scale; a linear relationship between self-report scores and log-transformed casualty counts suggests a logarithmic relationship between self-report scores and raw casualty counts.

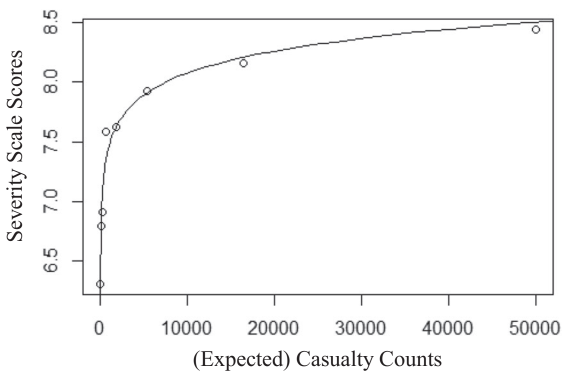

Fig. 2. Severity score group means as a function of casualty count, along with the best-fitting logarithmic curve.

suggesting a logarithmic relationship between raw casualty counts and severity scores.

Linear regression models testing the effects of log-transformed casualty count on self-report scores confirmed this pattern of results. Casualty counts exhibited no logarithmic effect on self-reported fear ($b = -0.035$, $p = 0.34$) or anger ($b = -0.022$, $p = 0.52$), but did exhibit a significant logarithmic effect on perceptions of attack severity ($b = 0.26$, $p < 0.001$), suggesting that self-reported severity perceptions increased by 0.29 scale points (roughly 0.15 *SD*) for each tripling of the casualty count. Perceived scenario plausibility also significantly predicted perceived severity ($b = 0.18$, $p < 0.001$), but did not attenuate the effect of casualty count when added into the model ($b = 0.27$, $p < 0.001$).

Of course, this regression forces a logarithmic functional form on the casualty–severity relationship, an assumption that should be checked. Fig. 2 plots severity scale scores as a function of casualty count, along with the best-fitting logarithmic function (i.e., the log-transformation of the best-fit regression line from the previous analysis). The group means cluster tightly around the best-fit curve and suggest that the data are indeed well-modeled by a logarithmic function. Note that the concavity of the severity–casualty function is evident even at 5,400 casualties, suggesting that this effect is not a mere artifact of the highest casualty condition (50,000).

3.2. Robustness Across Scenario Variables

Attack method (chemical vs. biological) and casualty time horizon (immediate vs. 20 years) both influenced respondents' mean ratings of attack severity, but did not appreciably alter the relationship between casualty count and any of the dependent variables. Specifically, adding the scenario variables as dichotomous contrasts to the linear regression analyses did suggest that respondents perceived the chemical attack as more severe than the biological attack ($d = 0.15$, $p = 0.005$), and the immediate casualty attack more severe than the delayed casualty attack ($d = 0.39$, $p < 0.001$). These main effects were qualified by a significant interaction between weapon type and time horizon ($p = 0.009$). Delayed casualties (compared to immediate) substantially reduced average severity ratings for the biological weapon scenarios ($M_{\text{bio.delayed}} = 6.72$, $M_{\text{bio.immediate}} = 7.83$, $d = 0.96$) but had only a small effect for the chemical weapon scenarios ($M_{\text{chem.delayed}} = 7.36$, $M_{\text{chem.immediate}} = 7.88$, $d = 0.26$). There were no significant scenario effects on levels of fear or anger ($ds < 0.05$).

Notably, attack method and casualty time horizon did not significantly moderate the effects of log-transformed casualty counts on fear, anger, or perceived severity ($ps > 0.05$). Furthermore, 95% confidence intervals for the effect of log-transformed casualty counts on severity ratings exhibited considerable overlap between the four scenarios. The confidence intervals for the two delayed casualty scenarios ([0.116, 0.360] and [0.187, 0.447]) overlapped considerably with those of the immediate casualty scenarios ([0.146, 0.351] and [0.191, 0.395]), and the same was true for the biological weapon scenarios ([0.146, 0.351] and [0.187, 0.447]) versus the chemical weapon scenario ([0.116, 0.360] and [0.191, 0.395]).

Fig. 3 shows the effect of casualty count on anger, fear, and severity ratings across all four scenarios. The plots show considerable overlap in the scenario-specific mean scores for anger and fear, suggesting the null effect of time horizon and weapon type on these variables. Conversely, the parallel trend lines for severity ratings demonstrate how scenario type impacted overall perceptions of severity, but not the strength or form of their dependence on casualty counts. That is, severity ratings were approximately linear in log-transformed casualty counts across all four scenario conditions, with no significant variations in the slope of these relationships.

4. DISCUSSION

Consistent with past research on public perceptions of life loss, this experiment demonstrated that respondents' perceptions of a terror attack's severity were sensitive to the attack casualty count, though such sensitivity diminished as casualty counts increased. This logarithmic functional form of the casualty–severity relationship aligns with past research on sensory perception and decision theory, given its prevalence in psychophysical and utility functions, and builds on the notion that sensitivity to a constant change in casualty count should decrease as the total casualty count increases.

Interestingly, this effect was observed only for respondents' perceptions of attack severity (which we deemed a "cognitive" perceptual measure), but not self-reported anger or fear. This may explain why some past studies (Burns & Slovic, 2010; Cui et al., 2016) did not find main effects of casualty counts on respondents' perceptions, given that these studies focused on affective and risk perception measures rather than perceived severity. It also suggests an interesting bifurcation in how respondents made sense of the attack scenario they were presented—their cognitive assessment of the attack's impact seemed to take death toll into account, but self-reported emotions did not. The same was true of the scenario manipulations—there was a particularly strong reduction in perceived severity in the delayed biological scenario, but no accompanying effect on fear or anger ratings. Of course, this may partly be a limitation of our self-report measures of fear and anger, but it may also suggest a fruitful avenue for future research on public terror attack perceptions. Psychologists have long known that cognitive and emotional appraisals of disaster events depend on different sources of information (Loewenstein et al., 2001), though this study raises the question of why casualty count, arguably the most salient and dreaded effect of a terror attack, did not factor into respondents' self-reported emotional responses.

From a theoretical standpoint, these results make a strong contribution to the literature on loss-of-life perceptions due to the study's experimental, between-subject design. Rather than infer a concave casualty–perception relationship based on risk preferences (Olivola & Sagara, 2009) or relative comparisons (Fetherstonhaugh et al., 1997), this experiment explicitly generated a "psychophysical curve" for casualty count severity perceptions using a between-subject design in which respondents were exposed to only a single casualty count. Remarkably, a strong logarithmic relationship between casualty count and severity perceptions clearly emerged. This result strongly bolsters findings from previous research that people react to relative, though not absolute, changes in life loss—a defining feature of the logarithmic relationship shown here. Note that similar results might have emerged by exposing participants to multiple casualty counts; however, doing so might have induced participants to respond based on rational comparisons between casualty conditions (e.g., "surely 5,000 deaths are worse than 200") rather than on their intrinsic reactions to a single scenario.

From a practical standpoint, this study suggests that casualty count can effectively convey information about the severity of a terror attack. Agencies or organizations needing to communicate to the public about past (or possible) terror attacks can likely trust that the attack's death toll will be generally perceived as a proxy for its severity across a wide range of attack types. Of course, such agencies should keep in mind the diminishing marginal sensitivity of this relationship. The perceived severity of a fixed number of lost or saved lives (e.g., a security procedure

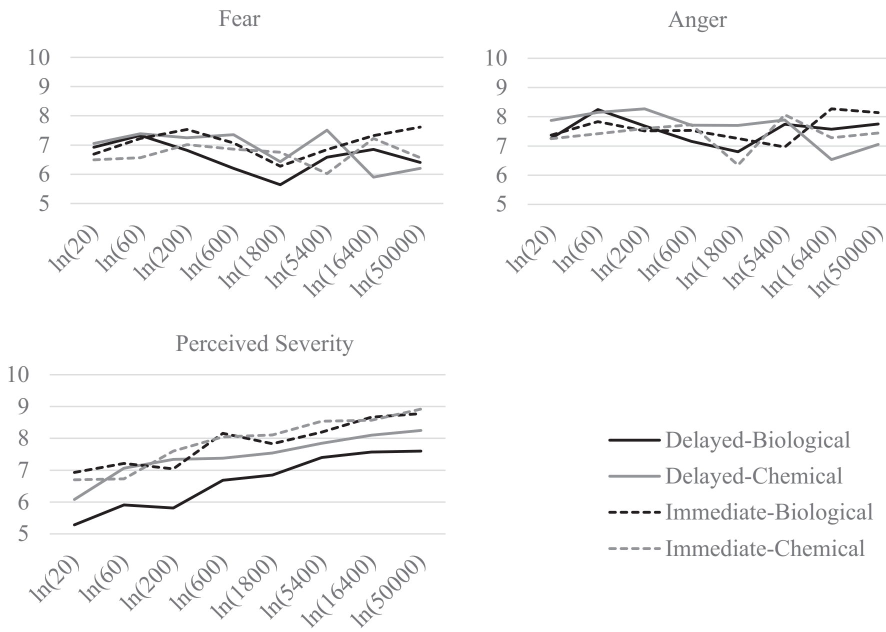

Fig. 3. Mean severity, fear, and anger scale scores by log-transformed casualty count and by scenario.

that can save 50 lives) will likely decrease as the total number of lives in question increases—a previously established pattern that has now been specifically demonstrated in the context of terror attacks.

Although not central to our theoretical concerns, it is interesting to note the effects of casualty time horizon and attack method on respondents' perceptions of attack severity. Specifically, severity ratings were particularly low among those who viewed the delayed casualty biological attack. In this scenario, most of the attack's eventual victims were not exposed to the weapon during the initial attack, possibly causing its lethality to seem less immediate/tangible than the immediate scenarios or the delayed chemical scenario (in which the victims were exposed to the lethal agent during the attack, and simply took time for its effects to be realized). However, future studies should better investigate the public perceptions of unconventional terror attack methods.

Among this study's primary limitations are its reliance on self-reported emotion and on hypothetical terror attack vignettes. Although the use of vignettes was necessary to control for all attack features besides casualty count, future studies may attempt similar analyses using real-world occurrences. Future work should also supplement these findings with other measures of reaction intensity, such as psychophysiological measures; such methods (e.g., Hsee & Rottenstreich, 2004) could aid in corroborating (or discrediting) our study's lack of a casualty–emotion link. Future work should also investigate whether "perceived attack severity" has any meaningful attitudinal or behavioral correlates in the aftermath of a terror attack. Still, we argue that this study lends strong theoretical support to the study of life loss perceptions, while extending previously demonstrated effects to the nuanced and complicated domain of terrorism.

ACKNOWLEDGMENTS

This research was supported by funding from the U.S. Department of Homeland Security (DHS), PTE Federal Award No. ASU-209466 (prime 2017-ST-061-QA001-01), and from the U.S. Defense Advanced Research Projects Agency (DARPA), Award No. HR00111820004.

REFERENCES

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, *22*(1), 23–36.

Burns, W. J., & Slovic, P. (2010). Predicting and modeling public response to a terrorist strike. In P. Slovic (Ed.), *The feeling of risk: New perspectives on risk perception* (pp. 285–306). London: Earthscan Publishing.

Cui, J., Rosoff, H., & John, R. S. (2016). Cumulative response to sequences of terror attacks varying in frequency and trajectory. *Risk Analysis*, *36*(12), 2272–2284.

Fechner, G. (1966). *Elements of psychophysics* (Vol. 1). New York: Holt, Rinehart and Winston.

Fetherstonhaugh, D., Slovic, P., Johnson, S., & Friedrich, J. (1997). Insensitivity to the value of human life: A study of psychophysical numbing. *Journal of Risk and Uncertainty*, *14*(3), 283–300.

Friedrich, J., Barnes, P., Chapin, K., Dawson, I., Garst, V., & Kerr, D. (1999). Psychophysical numbing: When lives are valued less as the lives at risk increase. *Journal of Consumer Psychology*, *8*(3), 277–299.

Friedrich, J., & Dood, T. L. (2009). How many casualties are too many? Proportional reasoning in the valuation of military and civilian lives 1. *Journal of Applied Social Psychology*, *39*(11), 2541–2569.

Hsee, C. K., & Rottenstreich, Y. (2004). Music, pandas, and muggers: On the affective psychology of value. *Journal of Experimental Psychology: General*, *133*(1), 23–30.

Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of Personality and Social Psychology*, *81*(1), 146–159.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, *127*(2), 267–286.

Marshall, R. D., Bryant, R. A., Amsel, L., Suh, E. J., Cook, J. M., & Neria, Y. (2007). The psychology of ongoing threat: Relative risk appraisal, the September 11 attacks, and terrorism-related fears. *American Psychologist*, *62*(4), 304–316.

Olivola, C. Y., & Sagara, N. (2009). Distributions of observed death tolls govern sensitivity to human fatalities. *Proceedings of the National Academy of Sciences*, *106*(52), 22151–22156.

Rosoff, H., John, R. S., & Prager, F. (2012). Flu, risks, and videotape: Escalating fear and avoidance. *Risk Analysis*, *32*(4), 729–743.

Rosoff, H., Siko, R., John, R., & Burns, W. J. (2013). Should I stay or should I go? An experimental study of health and economic government policies following a severe biological agent release. *Environment, Systems, and Decisions*, *33*(1), 121–137.

Slovic, P. (2007). "If I look at the mass I will never act": Psychic numbing and genocide. *Judgment and Decision Making*, *2*(2), 79–95.

Summers, C., Slovic, P., Hine, D., & Zuliani, D. (1994). "Psychophysical numbing": An empirical basis for perceptions of collective violence. Unpublished manuscript.

Sunstein, C. R. (2003). Terrorism and probability neglect. *Journal of Risk and Uncertainty*, *26*(2–3), 121–136.

Viscusi, W. K. (2009). Valuing risks of death from terrorism and natural disasters. *Journal of Risk and Uncertainty*, *38*(3), 191–213.

Weber, E. H. (1834). *De Pulsu, Resorptione, Auditu et Tactu: Annotationes Anatomical et Physiological*. Leipzig: Koehler.