

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/137720>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

© 2020 Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

A Brief History of Risk

Ying Li,^a Thomas Hills,^b Ralph Hertwig^a

^aCenter for Adaptive Rationality, Max Planck Institute for Human Development, 14195 Berlin, Germany; ^bDepartment of Psychology, University of Warwick, University Road, Coventry, CV4 7AL, United Kingdom

We thank Susannah Goss for editing the manuscript.

Corresponding Author: Ying Li
Contact: li@mpib-berlin.mpg.de

20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42

Abstract

Despite increasing life expectancy and high levels of welfare, health care, and public safety in most post-industrial countries, the public discourse often revolves around perceived threats. Terrorism, global pandemics, and environmental catastrophes are just a few of the risks that dominate media coverage. Is this public discourse on risk disconnected from reality? To examine this issue, we analyzed the dynamics of the risk discourse in two natural language text corpora. Specifically, we tracked latent semantic patterns over a period of 150 years to address four questions: First, we examined how the frequency of the word *risk* has changed over historical time. Is the construct of risk playing an ever-increasing role in the public discourse, as the sociological notion of a ‘risk society’ suggests? Second, we investigated how the sentiments for the words co-occurring with *risk* have changed. Are the connotations of *risk* becoming increasingly ominous? Third, how has the meaning of *risk* changed relative to close associates such as *danger* and *hazard*? Is *risk* more subject to semantic change? Finally, we decompose the construct of *risk* into the specific topics with which it has been associated and track those topics over historical time. This brief history of the semantics of risk reveals new and surprising insights—a fourfold increase in frequency, increasingly negative sentiment, a semantic drift towards forecasting and prevention, and a shift away from war toward chronic disease—reflecting the conceptual evolution of risk in the archeological records of public discourse.

Keywords: risk, danger, public discourse, content analysis, topic model, Ngram Corpus

79 Before we turn to our research questions, let us clarify that the term *risk* is often used to mean
80 different things. In the risk management and actuarial literature, for instance, it describes a loss of a
81 certain magnitude (e.g., injury, mortality) weighted by the probability of its occurrence (Rayner &
82 Cantor, 1987; Short, 1984). By this actuarial measure, driving is riskier than flying because it is
83 associated with a greater risk of injury per mile travelled. In the economic discourse, risk commonly
84 refers to the variance in possible (positive or negative) returns. For instance, an investment option with
85 higher return variance is deemed as riskier than an option with lower variance but the same expected
86 mean return (Markowitz, 1952; Pratt 1964). Research in psychology, sociology, and anthropology has
87 consistently demonstrated that these actuarial and economic definitions are too narrow to capture
88 people’s understanding of risk. Lay perceptions are multidimensional, encompassing higher order
89 factors such as *dread* and *equitable exposure* (Bhatia, 2019; Slovic, 1987). *Dread risks*, as opposed to
90 *chronic risks*, are defined by a perceived lack of control and potential large-scale loss of life, making
91 flying a greater perceived risk than driving (e.g., Gaissmaier & Gigerenzer, 2012). Greater dread, in
92 turn, is associated with greater perceived risk and a greater desire for regulation to reduce the risk
93 (Slovic, 1987; Slovic et al., 1985; Sunstein, 2005). All these meanings and others are part of the public
94 discourse and are included in the text corpora that we analyze. In other words, our focus is not on one
95 definition at the expense of another, but rather endorses the rich and inclusive semantic history of *risk*
96 in the natural language.

97 **2. Guiding Research Questions**

98 Our goal in this study was to track change in the public discourse on risk over historical time
99 by addressing four guiding questions. First, we examined how the frequency of the word *risk* has
100 changed over historical time. Word frequency analysis has been used to capture patterns of usage
101 associated with changes in cultural importance (Greenfield, 2013; Twenge et al., 2012; Uz, 2014). Here,
102 it allows us to evaluate the idea that the construct of risk is playing an ever-increasing role in the public
103 discourse, as suggested by the sociological concept of a “risk society” (Beck, 1992) and the
104 anthropological observation that the scope of the word *risk* has broadened over time (Douglas, 1992, p.
105 14). Second, we investigated how the sentiments for the words co-occurring with *risk* have changed.
106 This sentiment analysis allows us to evaluate the hypothesis that risk is becoming a more negative
107 construct, associated with expectations that societies and policy makers should invest ever more in risk
108 reduction and prevention (the precautionary principle; Sunstein, 2005). Third, we asked how the
109 meaning of *risk* has changed by examining change in the semantic relationship between it and other
110 words. The meaning of a word can be reliably inferred from the contexts in which it has been used
111 (Firth, 1957). For example, analysis of the linguistic context of the verb *broadcast* shows that 150 years
112 ago it referred to the spreading of seed, whereas it is now used to mean the spreading of information
113 (Li et al., 2019). We examined the text corpora for indications that *risk* is more subject to semantic
114 change than are close semantic associates such as *danger* and *hazard*. According to social

115 anthropologist Douglas (1992), the concept of *risk* has a strong cultural foundation, but this foundation
116 is not static: Perspectives and social environments change; some dangers are politicized as risks while
117 other worries are backgrounded. If *risk* has become a crucial construct for singling out certain objective
118 dangers and designating them as social concerns, then tracking the use of the term in public discourse
119 can reveal an underlying dynamic mechanism that is constantly responding to the changing
120 sociocultural environment (Douglas, 1992). Fourth, we decomposed the construct of *risk* into the
121 specific topics with which it has been associated and tracked those topics over historical time. Our
122 purpose here was to identify the most prominent risk topics over time and to consider how they have
123 changed in relation to world events.

124 We investigated these questions by analyzing latent semantic patterns in natural language.
125 Tracing the historical meanings of words requires a corpus of texts published over a sufficiently long
126 time period. The Google Books Ngram Corpus (Lin et al., 2012) is one of the few corpora that meet
127 this requirement. Drawing on over 100 sources (e.g., libraries and publishers), it contains over 8 million
128 books published from 1600 to 2008, or 6% of all books ever published. The corpus thus offers a
129 *telescopic view* over a large time period. The corpus has been used to detect large-scale changes in
130 language, which in turn correlate with social and demographic changes (Hills & Adelman, 2015; Hills
131 et al., 2012; Hills et al., 2015; Michel et al., 2011). Any corpus, however, has its limitations. The Google
132 Books Ngram Corpus offers limited contextual information due to a narrow window size (5-grams, or
133 a contiguous sequence of five words); moreover, there has been a surge in the proportion of academic
134 articles in the corpus (Pechenick et al., 2015). We therefore also examined *The New York Times*
135 *Annotated Corpus* (NYT corpus; Sandhaus, 2008) to lend convergent validity to our results. This corpus
136 contains all (1.8 million) articles published in the *New York Times* from 1987 to 2007, and offers a more
137 *microscopic view* on the risks of modern life as reported in the most widely read U.S. newspaper. Let
138 us emphasize that because our analysis draws on English texts only, the present results are limited to
139 English-speaking cultures. Nevertheless, the Google Books Ngram Corpus, in particular, has the
140 advantage of covering a relatively long time period, going beyond short-term analyses of, for instance,
141 media coverage of risk and mortality (see the references in Young et al., 2008).

142 3. Materials and Methods

143 We used word co-occurrence to construct semantic representations of risk in each year of the
144 analysis, such that the meaning of *risk* was approximated by the context in which it was used. The co-
145 occurrence information allowed us to quantify how the sentiment and semantics of *risk* have changed
146 over history. As *risk* may be used in multiple contexts, we used Latent Dirichlet Allocation (LDA,
147 Blei et al., 2003) to identify the historical risk topics. This topic model algorithm detects underlying
148 topics that best explain the structure of the language around *risk*, and allowed us to identify risk topics
149 as they changed over time. In what follows, we describe this procedure in more detail. We begin by

150 briefly describing the Corpora we used. We provide our data on
151 https://osf.io/jctn8/?view_only=988fccae28ca4995b6a3796002a888cc.

152

153 **3.1 Google Books Ngram Corpus**

154 The Google Books Ngram Corpus consists of n -grams: contiguous sequences of n items from
155 a given text (n ranges from 1–5). We used the 5-grams of all English words in our analysis; each data
156 entry therefore displays the number of times a 5-gram appears in the corpus during a specific year. We
157 retrieved all 5-grams starting or ending with the word *risk*. As is standard procedure in many natural
158 language processing tasks, we removed stop words, punctuation, digits, and words containing fewer
159 than three characters before using the WordNet-based NLTK lemmatizer (Bird et al., 2009) to
160 lemmatize each noun to its singular form and each verb to its present tense. Next, we aggregated the
161 corpus by year so that each document contained all 5-grams in a specific year. Aggregating topics by
162 years encourages the topic model to identify the underlying patterns that best explain differences among
163 risk structures over years.

164 **3.2 The New York Times Annotated Corpus**

165 The NYT Corpus contains all articles published in the *New York Times* from 1987 to 2007. We
166 constructed a risk corpus by selecting articles that mentioned the word *risk* or *risks* more than twice.
167 Next, we pre-processed the corpus in the same way as we did the Google Books Ngram data, apart from
168 aggregating articles by year: Each news article was treated as one document.

169 **3.3 Corpus of Historical American English**

170 The Corpus of Historical American English (COHA) is a large structured corpus of historical
171 English. It contains 400 million words of text produced from the 1810s to 2000. COHA is balanced by
172 genre decade by decade, which brings both benefits and concerns. On one hand, it alleviates concerns
173 that insights gained from the corpus are driven by the changing compositions of genres. On the other
174 hand, it may fail to map the reality that public preferences for genres change over history. Although it
175 is difficult to argue whether COHA is a better corpus for analyzing culture change than the Google
176 Books Ngram corpus or vice versa, consistency in the findings from both corpora would lend
177 convergent validity to the results. Therefore, we used COHA to validate some of the historical analysis
178 conducted with the Google Books Ngram Corpus, namely, the analysis of frequency and semantic shift.

179 **3.4 Analysis of Frequency and Contextual Sentiment**

180 Analyses of frequency, contextual sentiment, and semantic drift (Figures 1 and 2) were conducted using
181 the Macroscopic (Li et al., 2019), an interactive linguistic tool that analyzes historical sentiment and
182 semantic change. The Macroscopic was built on the basis of the historical word co-occurrence data made
183 publicly available through the Google Books Ngram Corpus. Frequency was calculated by dividing the
184 count of the selected words by the corpus size to control for the different corpus sizes for each year.
185 Contextual sentiment for the selected words was computed in terms of the averaged valence ratings of

186 co-occurring words during a given year. The valence ratings were retrieved from data collected by
 187 Warriner et al. (2013), which contain valence scores for 13,915 English words, each rated on its
 188 “pleasantness” by around 30 participants. Using contemporary norms to estimate the valence of words
 189 used decades ago is potentially problematic, as all words may have changed their meaning or sentiment
 190 over history. In practice, however, it has been shown that historical sentiment as inferred from averaging
 191 contemporary valence norms of semantic neighbors is similar to the sentiment judged by historical
 192 language experts (Buechel, Hellrich, & Hahn, 2016).

193 3.5 Semantic Shift Analysis

194 The purpose of our semantic drift analysis was to examine how and to what extent the meaning
 195 of *risk* has changed over the past two centuries in relation to related concepts such as *danger*, *fear*, and
 196 *hazard*. It consisted of the following three steps: First, we retrieved the historical word embeddings for
 197 50,000 common English Words trained by Li et al. (2019). Word embeddings provide a vector
 198 representation for each word based on its co-occurring relationship with other words; that is, they
 199 represent the context in which a word has been used. To derive word embeddings, Li et al. (2019) first,
 200 from Google Ngram Corpus, constructed a co-occurrence matrix for 50,000 common English words
 201 that records the number of times any two words were used within the same 5-gram. Next, they computed
 202 the positive pointwise mutual information (PPMI) for each pair of words and then constructed a PPMI
 203 matrix with entries given by:

$$204 \quad PPMI(v_i, v_j) = \max(0, \log(\frac{p(v_i, v_j)}{p(v_i) \times p(v_j)})), \quad (1)$$

205 where v_i, v_j represents a pair of words from the corpus and $p(v)$ corresponds to the empirical
 206 probabilities of those words co-occurring within a sliding window of 5 over the original text. Finally,
 207 Li et al. reduced the dimension of word embeddings to 300 using singular value decomposition (SVD).
 208 This dimensionality reduction acts as a form of regularization and allowed us to compare word
 209 similarities by computing the cosine similarity of word embeddings.

210 Second, drawing on the historical word embeddings trained by Li et al. (2019), we identified
 211 the 8-nearest semantic neighbors for the words *risk*, *danger*, *fear*, and *hazard*. Specifically, we retrieved
 212 word embeddings for each of the four target words and their semantic neighbors in the years 1800 and
 213 2000. For *risk*, we also retrieved the historical embeddings every 20 years between 1800 and 2000. In
 214 order to compare word embeddings from different time periods, we must ensure that the vectors are
 215 aligned to the same coordinate axes. We therefore used Orthogonal Procrustes to align the historical
 216 embeddings (Schönemann, 1966).

217 Third, we visualized semantic shift of words in two-dimensional space. To this end, we used
 218 principal component analysis to reduce the dimensions of word embeddings from 300 to 2. Figure 2
 219 plots the word embeddings retrieved in the second step according to the two orthogonal principal
 220 components (PC1 and PC2). These two principal components represent compressed dimensions that
 221 best explain the variance of the raw data and are therefore not directly interpretable except in relation

222 to relative distance between word embeddings. The background words (semantic neighbors) are always
 223 shown in their “modern” (year 2000) positions. This approximation is necessary since, in reality, all
 224 words are moving. *Risk* and its synonyms are shown in their modern and historical positions. The path
 225 travelled through the semantic space is a proxy for change in historical meaning.

226 Finally, to validate our observations, we quantified semantic change in *risk* and its related
 227 concepts using historical word embeddings trained on COHA (Hamilton et al., 2016) and on the Google
 228 Books Ngram Corpus (Li et al., 2019). For each word, we computed cosine similarity between
 229 embeddings trained on the 1820¹ corpus and on the 2000 corpus.

230 3.6 Topic Modelling

231 We studied historical change in the meaning of the word *risk* by extracting risk topics from the
 232 Google Books Ngram Corpus (Lin et al., 2012) and the NYT corpus (Sandhaus, 2008). The topic model
 233 we used was Latent Dirichlet Allocation (LDA; Blei et al., 2003), a bag-of-words algorithm that
 234 identifies a set of topics that best describe/re-generate the corpus. We took two main steps in analyzing
 235 the data. First, we identified the structure of risk meanings by applying the topic model to the risk corpus.
 236 This step allowed us to understand the key events associated with risk. Next, we applied trend analysis
 237 to understand how the risk topics identified in the first step changed over time.

238 3.7 Interpreting Topics

239 To make sense of the meanings of the risk topics, we used Equation (2) to identify the words
 240 most relevant to each topic. The relevance of term w to topic k given a weight parameter λ was defined
 241 as:

$$242 \quad \gamma(w, k|\lambda) = \lambda \log(P(w|k)) + (1 - \lambda) \log\left(\frac{P(w|k)}{P(w)}\right), \quad (2)$$

243 where $P(w|k)$ is the probability of term w being assigned to topic k and $P(w)$ is the marginal
 244 probability of term w being in the corpus. The first component of the equation, $P(w|k)$, prioritizes terms
 245 with high frequency in a topic. However, it does not consider how unique term w is to topic k , which
 246 can be captured by $\frac{P(w|k)}{P(w)}$, a quantity that Taddy (2012) called *lift*. We set λ to 0.5 to take both
 247 components into consideration; λ determines the weight given to the probability of term w under topic
 248 k relative to its lift.

249 One issue with topic models is that it is not clear which topics capture structures specific to the
 250 risk corpus and which topics capture general features of the source corpus. To find out, we used
 251 Equation (3) to compute the specificity of topic k to the risk corpus:

$$252 \quad \text{Specificity}(k) = \sum_{i=1}^n \left(\frac{\gamma(w_i|k)}{\sum_{i=1}^n \gamma(w_i|k)} * \frac{p(w_i|risk\ corpus)}{p(w_i|general\ corpus)} \right), \quad (3)$$

¹ We chose 1820 instead of 1800 because the frequency of *risk* in COHA in 1810 proved too small to train a stable model.

253 where $\frac{\gamma(w_i|k)}{\sum_{i=1}^n \gamma(w_i|k)}$ is the normalized relevance of word w to topic k , and $\frac{p(w_i|risk\ corpus)}{p(w_i|general\ corpus)}$ is the ratio
 254 of the frequency of word w in the risk corpus to its frequency in the source corpus. Specificity can range
 255 from 0 to almost infinity. A specificity of 1 means that, on average, the words characterizing the topic
 256 have the same frequency in both the risk corpus and the source corpus, suggesting that the topic reflects
 257 the underlying pattern of the source corpus, not risk. An example of a nonspecific topic is one that
 258 generates the words necessary to construct every document, such as articles and pronouns. The absolute
 259 value of topic specificity is heavily influenced by the data format: NYT articles are more likely than 5-
 260 grams to contain non-risk-specific words (noise) and therefore have smaller values of
 261 $\frac{p(w_i|risk\ corpus)}{p(w_i|general\ corpus)}$. Topic specificity is not comparable across corpora; instead, it should be used
 262 to compare topics from a same corpus.

263 3.8 Tracking Trends in Topics

264 To analyze trends in topics over time, we used the output from the LDA model on the Google Books
 265 Ngram Corpus to calculate the contribution of each topic k in each year by applying Equation (4). For
 266 each document (i.e., all 5-grams in a specific year), the equation controls for document length by
 267 dividing the number of words generated by each topic by the total number of words in the document.
 268 Thus, the yearly topic contribution estimate, $p_d(k)$, is defined as:

$$269 \quad p_d(k) = \frac{|\{w \in d: topic(w) = k\}|}{|d|}, \quad (4)$$

270 where k is a topic and w is a word in a document d . The numerator is the number of words in document
 271 d that are generated by topic k ; the denominator is the total number of words in document d .

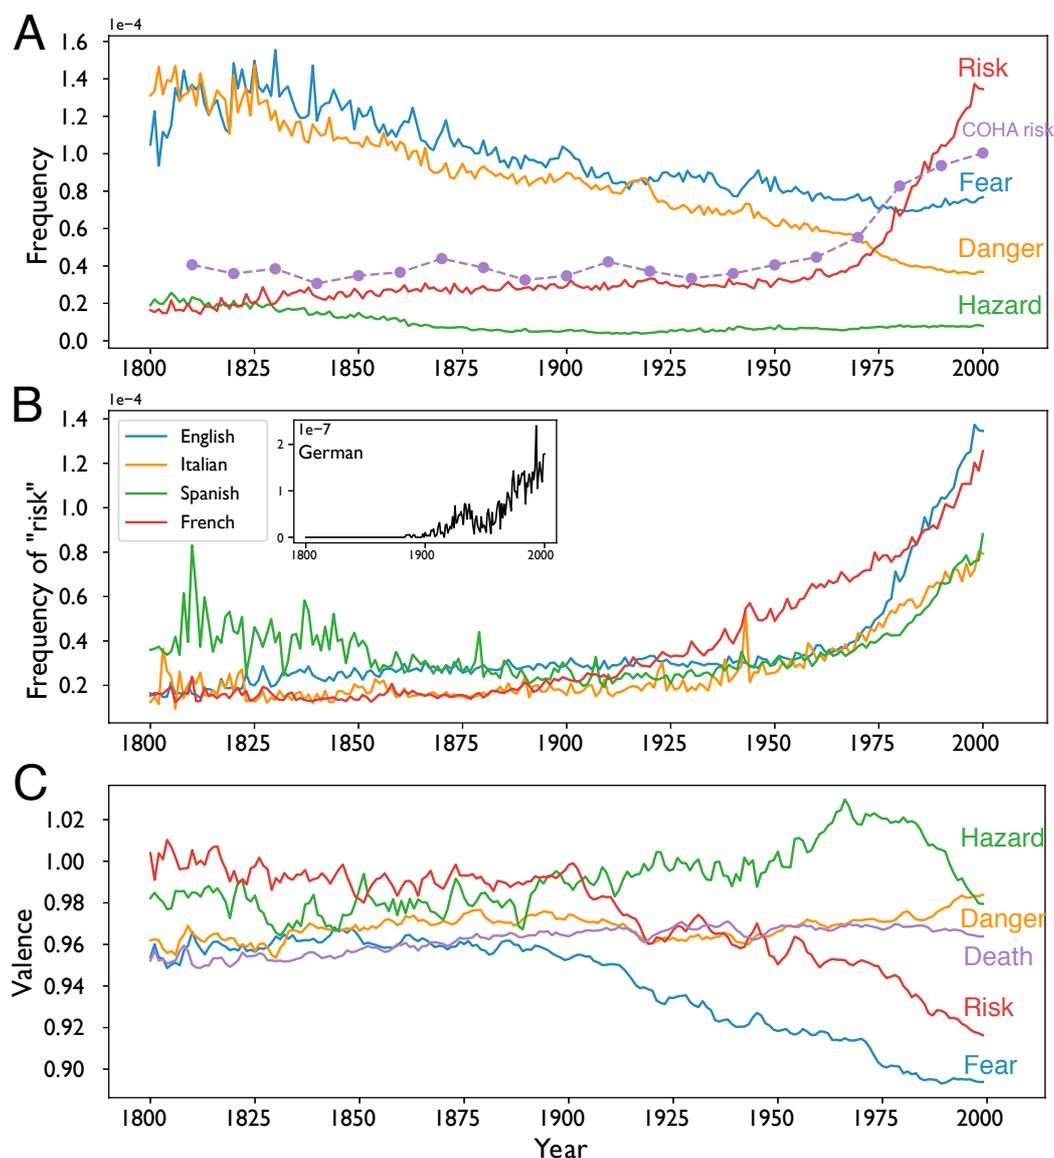
272 4. Results

273 4.1 How Has the Frequency of *Risk* Changed Over Time?

274 We first investigated change in the frequency of the word *risk* over time, starting with the
 275 Google Books Ngram Corpus. As Figure 1A shows, use of the word *risk* has increased dramatically
 276 since about 1970, with an approximately fourfold increase in usage since the 1950s. We checked this
 277 trend in English against other languages and found similar increases in French, German, Italian, and
 278 Spanish (Figure 1B). In addition, we observed a similar proliferation of *risk* in the Corpus of Historical
 279 American English (COHA; Davies 2012). As COHA is balanced by genre and subgenre across
 280 decades,² these findings suggest that *risk* proliferation is not an artifact of increasing numbers of
 281 scientific journals being included in the Google Books Ngram Corpus (Figure 1A). There is, however,
 282 no sign that the public discourse has turned darker in general, as close semantic relatives signifying
 283 undesirable states such as *danger*, *fear*, and *hazard* are not being used more frequently. On the contrary,

² For example, fiction accounts for 48–55% of the total in each decade (1810s–2000s); subgenres such as prose, poetry, and drama are likewise balanced. This balance across genres and subgenres means that researchers can be reasonably certain that patterns in the data do not merely reflect artefacts of a changing genre balance.

284 the use of *danger* and *fear* has declined steadily over the past two centuries, while the use of *hazard* has
 285 remained relatively stable at a low frequency. These results are consistent with the idea that *risk*, more
 286 than other terms, has become a central concept in recent public and political discourse (Beck, 1992;
 287 Bourke, 2005; Douglas, 1992).
 288



289
 290
 291 *Figure 1.* Historical change in the frequency and sentiment for the word *risk* and its close semantic
 292 neighbors in the Google Books Ngram Corpus. (A) Frequency of *risk*, *fear*, *danger*, and *hazard* in the
 293 Google Books Ngram Corpus and frequency of *risk* in the Corpus of Historical American English
 294 (COHA). (B) Frequency of *risk* in five languages—English, Italian, Spanish, French, and German—in
 295 the Google Books Ngram Corpus. German is presented in a separate box because the frequency of *risk*
 296 is much lower in German than in the other languages. (C). Change in the sentiment for words co-
 297 occurring with *risk*, *fear*, *danger*, *hazard*, and *death*. Sentiment was adjusted to mean score of all words,
 298 such that valences > 1 indicate a more positive context than average. The word *death* is included to
 299 provide a sentiment benchmark, as its meaning and sentiment have remained stable over history.

300 4.2 How Have the Sentiments Associated with *Risk* Changed?

301 Next, we examined whether the sentiments³ associated with *risk* have changed over time. For
302 example, is it possible—in line with a more economic interpretation of risk—that the use of the word
303 *risk* is increasingly associated with an appreciation of the large potential rewards that make some risks
304 worth taking (Pleskac & Hertwig, 2014)? This is not the case, as the results presented in Figure 1C
305 show. Computing the frequency-weighted average valence of the words that co-occurred with *risk* over
306 the past 200 years revealed that the sentiment associated with risk has become increasingly negative,
307 showing a roughly monotonic decline from 1800 to 2000. To provide points of comparison, we also
308 analyzed the related concepts of *danger*, *fear*, *hazard* as well as *death* as a benchmark. The sentiment
309 analysis shows that *risk* has undergone a much larger change over time than these inherently undesirable
310 concepts (with the exception of *fear*). In the early 1800s, the sentiment for words co-occurring with *risk*
311 was more positive than that of any of the four comparison words; by the end of 20th century, it was more
312 negative than that of *danger*, *hazard*, or *death* (Figure 1C). In other words, the word *risk* has become
313 not only more prevalent but also more negative in meaning.

314

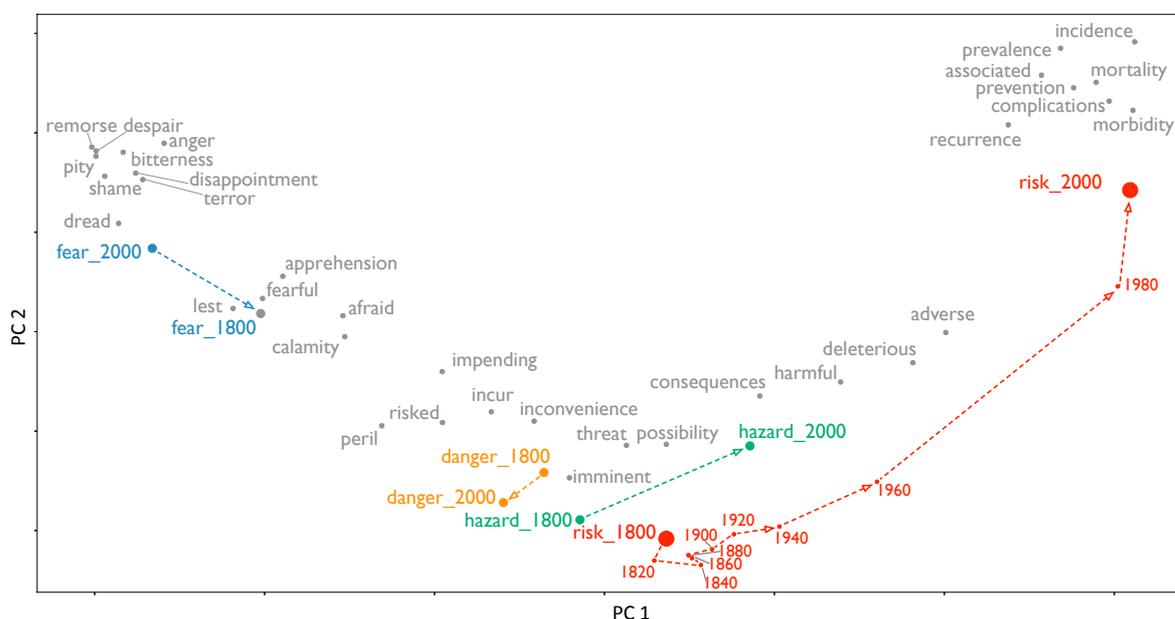
315 4.3 How Have the Semantic Relationships of Risk Changed?

316 The increasing negativity of *risk*'s sentiment, relative to the stability of the sentiment for *danger*,
317 *fear*, and *hazard*, might be driven by the changing contexts in which these words have been used. In
318 this section, we therefore turn to an analysis of semantic drift, which likewise suggests that the *risk* has
319 experienced more semantic change over historical time than have its close semantic associates.
320 Specifically, Figure 2 visualizes the semantic associates of *risk*, *danger*, *fear*, and *hazard* in two-
321 dimensional space relative to their k most similar words in 1800 and 2000 ($k = 9$ for each word). A
322 larger distance between two words suggests less similarity in the contexts in which they appeared. The
323 pattern is clear: *risk*, *danger*, and *hazard* started as close semantic neighbors in 1800 and moved apart
324 over time. By the year 2000, the underlying semantics of *risk* had grown more similar to those of
325 *prevalence* and *prevention*, terms associated with the quantification, reduction, and avoidance of risk.
326 *Danger* and *hazard*, in contrast, remained in the semantic area defined by words such as *harm*, *threat*,
327 *adverse*, and *peril*. This finding suggests that the word *risk* has moved from merely representing the
328 presence of threats to also being associated with the scientific examination, quantification, and
329 prevention of threats.

330 It is possible that this pattern is a result of an increase in the number of academic (especially
331 medical) articles in the Google Books Ngram Corpus (Pechenick et al., 2015). Therefore, we again used
332 COHA, a smaller yet genre-balanced corpus, to validate our findings. We analyzed the semantic shift

³ Because we inferred historical sentiment by averaging the valence of contextual neighbors, what we measured is sentiment of the context associated with *risk*, not directly sentiment of the word *risk*. However, the two are conceptually related: Because the meaning of a word can be learnt from the linguistic companions it keeps (Firth, 1957), words used in negative contexts are likely to carry negative connotations.

333 of *risk* using historical word embeddings trained on COHA (Hamilton et al., 2016) and compared the
 334 results with results derived from embeddings trained on the Google Books Ngram Corpus (Li et al.,
 335 2019). For each word, we quantified semantic similarity over history by computing the cosine similarity
 336 of embeddings trained on the 1820 corpus and the 2000 corpus. Cosine similarity scores range from 0
 337 to 1, with larger scores indicating greater semantic similarity. Comparison of results from the two
 338 corpora confirmed that the semantics of *risk* was much less stable than the semantics of *danger*, *fear*,
 339 and *hazard* (Table 1). In addition, we searched for the nearest semantic neighbors for *risk* in COHA in
 340 1820 and 2000. Again, we found that *risk* acquired associations with medical concepts over time: its
 341 top-5 nearest semantic neighbors changed from *loss*, *expense*, *danger*, *trouble*, and *run* in 1820 to
 342 *disease*, *diabetes*, *cancer*, *rate*, and *factors* in 2000.



343
 344 *Figure 2.* Semantic drift of *risk*, *danger*, *fear*, and *hazard* from 1800 to 2000 in the Google Books
 345 Ngram Corpus. The target words (in color) are shown in relation to their near associates⁴ (in gray) in
 346 the years 1800 and 2000. The meaning of *Risk* is shown at 11 historical points from 1800 to 2000 with
 347 a 20-year interval. PCA was performed to reduce the dimension of word embeddings from 300 to 2 so
 348 that words can be visualized in two-dimensional space. The axes represent the two principal
 349 components. A larger distance between two words indicates lower semantic similarity. The words *risk*,
 350 *danger*, and *hazard* started as near neighbors in 1800 but moved apart over time.

351

⁴ The semantic neighbors of *risk* identified from Google Ngram Corpus (as shown in the top right corner of this figure) are different from the words people report in the Small World of Word project (SWOW; De Deyne, Navarro, Perfors, Brysbaert & Storms, 2019) where they were required to “enter the first word that comes to mind” when reading the word *risk* (the most popular responses in SWOW are *danger*, *game*, *reward*, *gamble*, *chance*, *money*, to name a few). We argue that language corpora may have greater ecological validity because it reflects real situations where people find *risk* most appropriate to label an object. In contrast, behaviors in a free association task are likely to be influenced by non-semantic cues such as word frequency, phonological overlap, etc. However, we acknowledge that the need to be informative in language production may cause language corpora to underrepresent certain information, such as obvious fact like banana is yellow.

352
353
354
355
356
357
358
359

Table 1*Semantic Similarity Between 1820 and 2000*

	Risk	Danger	Fear	Hazard
Google Books Ngram	0.36	0.61	0.58	0.56
COHA	0.42	0.81	0.80	0.54

360 *Note.* For each word, semantic similarity was quantified by calculating the cosine similarity of word
361 embeddings between 1820 and 2000. The embeddings were normalized such that the similarity scores
362 range from 0 to 1, with 1 and 0 representing maximum and minimum similarity, respectively.
363

364 4.4 How Have Risk Topics Changed Over Time?

365 The semantic drift analysis shows how *risk* has diverged from its semantic neighbors over the
366 last two centuries, but it cannot provide detailed insights into the topical dimensionality of risk in this
367 period. As noted by Blais and Weber (2006), risk is a multidimensional concept encompassing
368 numerous topics. We therefore applied LDA to investigate the topics that have driven the proliferation
369 of *risk* in the public discourse and its increasingly negative sentiment. We inferred topic meanings by
370 inspecting their most relevant words (see Equation 2 in the Methods section), as summarized for each
371 topic in Table 2. Applying the topic model to the Google Books Ngram Corpus identified six risk
372 categories: **war** (topic 1, 2, 3), **nuclear** (topic 4), **health** (topic 5, 6, 7, 8, 9), **HIV/AIDS** (topic 10, 11),
373 **risk society** (topic 12), **economy** (topic 13, 14), and a non-specific topic on risk analysis (topic 15).

374 Each topic represents a probability distribution over all words. In order to validate our
375 interpretation of risk topics from the Google Books Ngram Corpus, we selected a collection of words
376 (see the left column of Figure 3A) that characterize each of the risk categories identified above and
377 examined how those words were distributed over topics (see the left panel of Figure 3A). Instead of
378 selecting words from Table 2, we chose a different set of associates that we felt exemplified our
379 interpretation of the topics, based on events occurring at the time the topics peaked. For example, under
380 the war category, we selected words that reflect the major participants in 20th century wars (e.g., *Soviet*,
381 *American, Japan, Germany*) as well as war-related words (e.g., *battle, invasion, army*). For the cancer
382 category, we included names of the most common cancers. If our interpretation was correct, topics
383 grouped under the same category should be more likely to generate corresponding words but not others.
384 This is indeed what we found. For example, Figure 3A shows that topics 1, 2 and 3 in the Google Books
385 Ngram corpus (identified as *war* topics in Table 2) were associated with the set of words we selected
386 under the *war* category. This pattern, visualized as probability loadings on the diagonal of the word-
387 topic probability heat map in Figure 3A, lends further support to our interpretation of topic meanings
388 presented in Table 2.

389

390

391 **Table 2**392 *Most Relevant Words for Each Risk Topic, Ordered by Relevance as Defined in Equation 2*

Index	Google Books Ngram Corpus	Index	NYT Corpus
1	Life, imminent, battle, resolve	1	Military, war, Iraq, troop
2	Life, war, bureau, loss	2	China, Japan, country, foreign
3	War, uncertainty, loss, prepare		
4	Nuclear, carcinogenic, patient, infant	3	Environmental, plant, energy, gas
5	Heart, coronary, injury, bear	4	Cancer, woman, study, breast
6	Breast, cancer, osteoporosis, fetus	5	Drug, patient, doctor, hospital
7	Stroke, cancer, disease, capital		
8	Prostate, cancer, event, Alzheimer		
9	Management, diabetes, cardiovascular, overweight		
10	AIDS, nation, HIV, immunodeficiency	6	AIDS, virus, infect, vaccine
11	HIV, deficit, assess, volume		
12	Management, value, assessment, society	7	Child, school, parent, student
13	Confrontation, return, equilibrium, preference	8	Fund, stock, investor, market
14	Rate, free, interest, return		
15	Behavio[u]r, group, death, population		
		9	Food, fat, eat, diet
		10	Insurance, bank, loan, insurer
		11	Law, court, abortion, tobacco
		12	Airline, flight, shuttle, space
		13	Company, business, executive, industry
		14	Investigation, Enron, prison, police
		15	Think, people, way, thing
		16	Republican, Clinton, Bush, Democrat
		17	Game, player, sport, team
		18	Day, car, hour, walk
		19	City, build, York, new
		20	Film, art, movie, theater

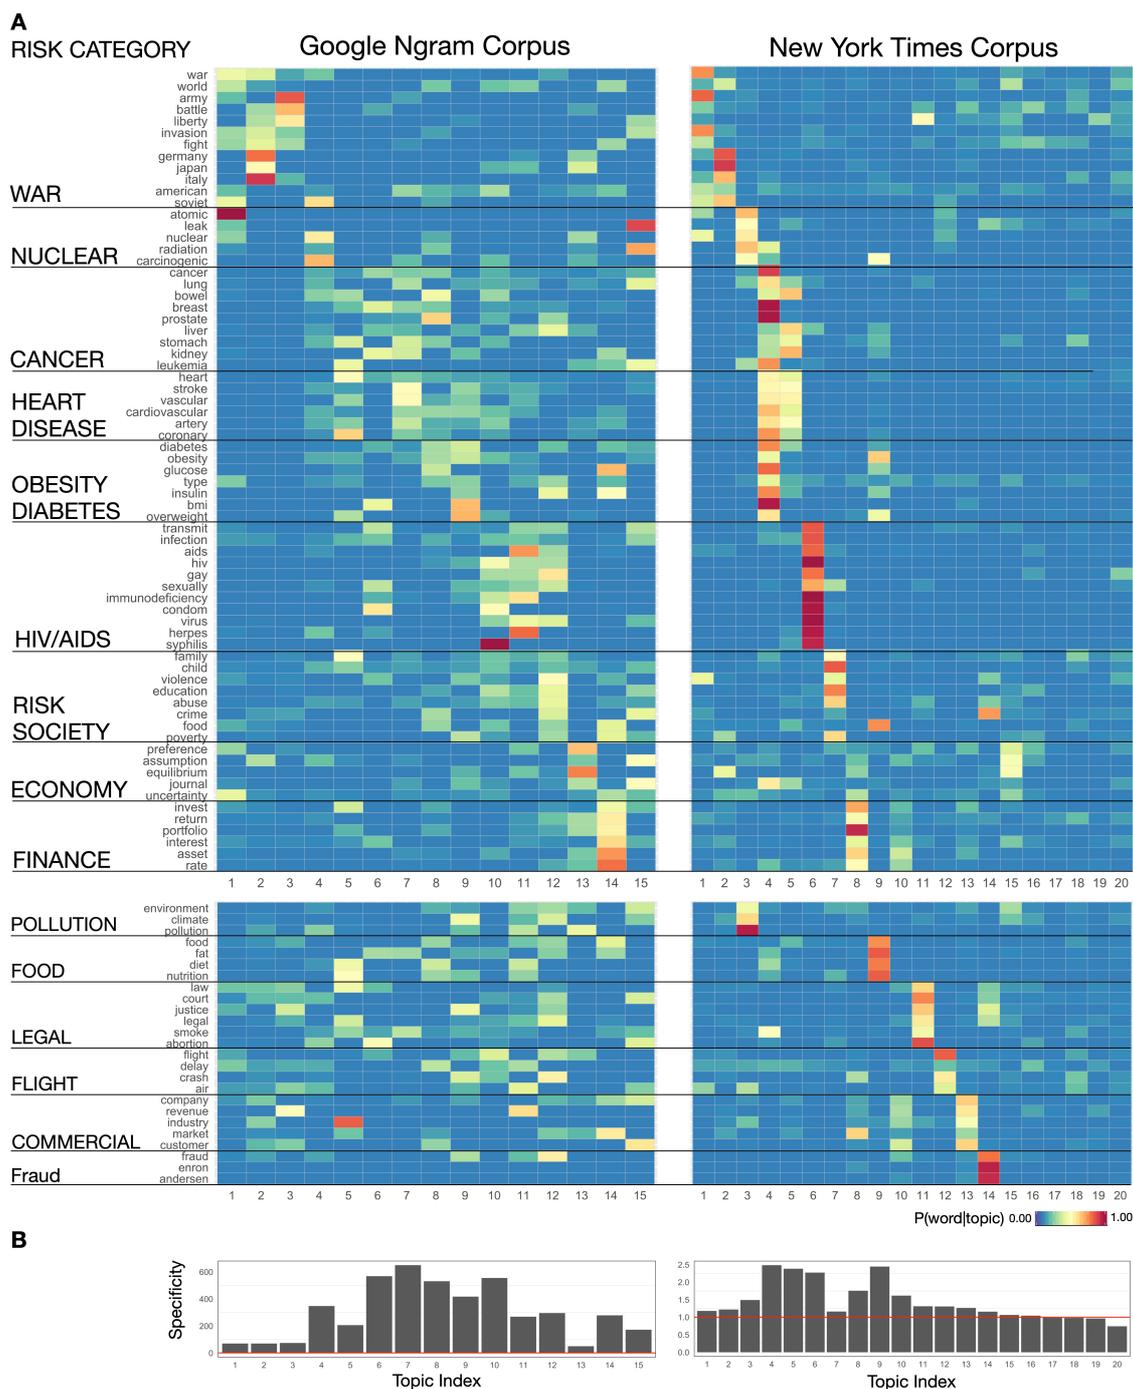
393 *Note:* Topics 15–20 of the NYT corpus are shown in gray to indicate that these topics are not specific
394 to articles that contain the word *risk* and its inflections. Topic specificity is defined in equation 3.
395

396 How replicable is this category structure? To find out, we also analyzed the NYT Corpus.
397 Applying the same procedure to the NYT Corpus confirmed all risk categories inferred for the Google
398 Books Ngram Corpus (visualized as probability loadings on the diagonal of the right panel of Figure
399 3A). We can therefore conclude that the meanings of risk derived in our analysis of the Google Books
400 Ngram dataset are not corpus-specific results associated with a non-representative sample, but reflect
401 general trends in the topicality of risk over both relatively long and short time scales.

402 In order to ensure that the topics were risk-specific and did not just reflect the background
403 features of the corpus, we next computed *topic specificity* (see Equation 3 in the Methods section) to
404 quantify the relative correspondence of each topic with the risk corpus as compared with the entire

405 corpus (see Figure 3B). A topic specificity score around or below 1 means that the topic has a
406 distribution of words similar to that seen in the entire corpus; the topic therefore represents the general
407 features of the entire corpus. For the Google Books Ngram Corpus, we found the topic specificity of all
408 risk topics to be above 1 (ranging from 50 to 650), suggesting that all topics were risk-relevant. In
409 contrast, the specificity of NYT topics ranged from 0.7 to 2.5, with six topics being irrelevant to risk
410 (the specificity scores of topics 15–20 were close to or less than 1). This notable difference in the topic
411 specificity of the two corpora may be attributable to differences in data format: Recall that the Google
412 Books Ngram data contain words that co-occurred with *risk* within a narrow window size, whereas the
413 NYT data contain entire articles that mention the word *risk*. As such, NYT articles are more likely than
414 Google Books Ngrams to contain words not specific to *risk*.

415 Nevertheless, both corpora rendered a similar set of high-specificity topics: nuclear, heart
416 disease, cancer, diabetes, and HIV/AIDS. War-related topics had low specificity in the NYT Corpus.
417 This result is not surprising because, as we show in the following analysis, war topics have gradually
418 disassociated from *risk* since World War II, and the NYT Corpus only dates back to 1987. Beyond the
419 risk topics identified for the Google Books Ngrams, we found only one additional topic in the NYT
420 Corpus with specificity clearly above 1 (topic 9, featuring words such as *food*, *fat*, *eat*, and *diet*), and
421 four additional NYT topics slightly above 1 (topics 11–14, which we interpreted as legal, flight,
422 commercial, and fraud, respectively). Correspondingly, the key words associated with topics 11–14
423 showed low co-occurrence with *risk* in the Google Books Ngram Corpus throughout history. This
424 comparison suggests that, overall, both corpora converged on a similar set of important risk categories.
425



426
 427 *Figure 3.* Visual quantification of risk topics. (A) Heatmap of the probability that word w was generated
 428 by topic k in models derived from the Google Books Ngram Corpus (left) and the NYT Corpus (right).
 429 Words on the y-axis were selected by referring to the list of most relevant words for each topic
 430 (relevance defined by Equation 2) and they were grouped by categories. (B) Topic specificity (as
 431 defined by Equation 3). The red horizontal line indicates topic specificity equal to 1. Topics with
 432 specificity above this reference line can be considered risk-specific and therefore capture one or more
 433 aspects of the meaning of risk. Topics with specificity below 1 can be considered generic words that
 434 are not informative with respect to risk meanings.

435

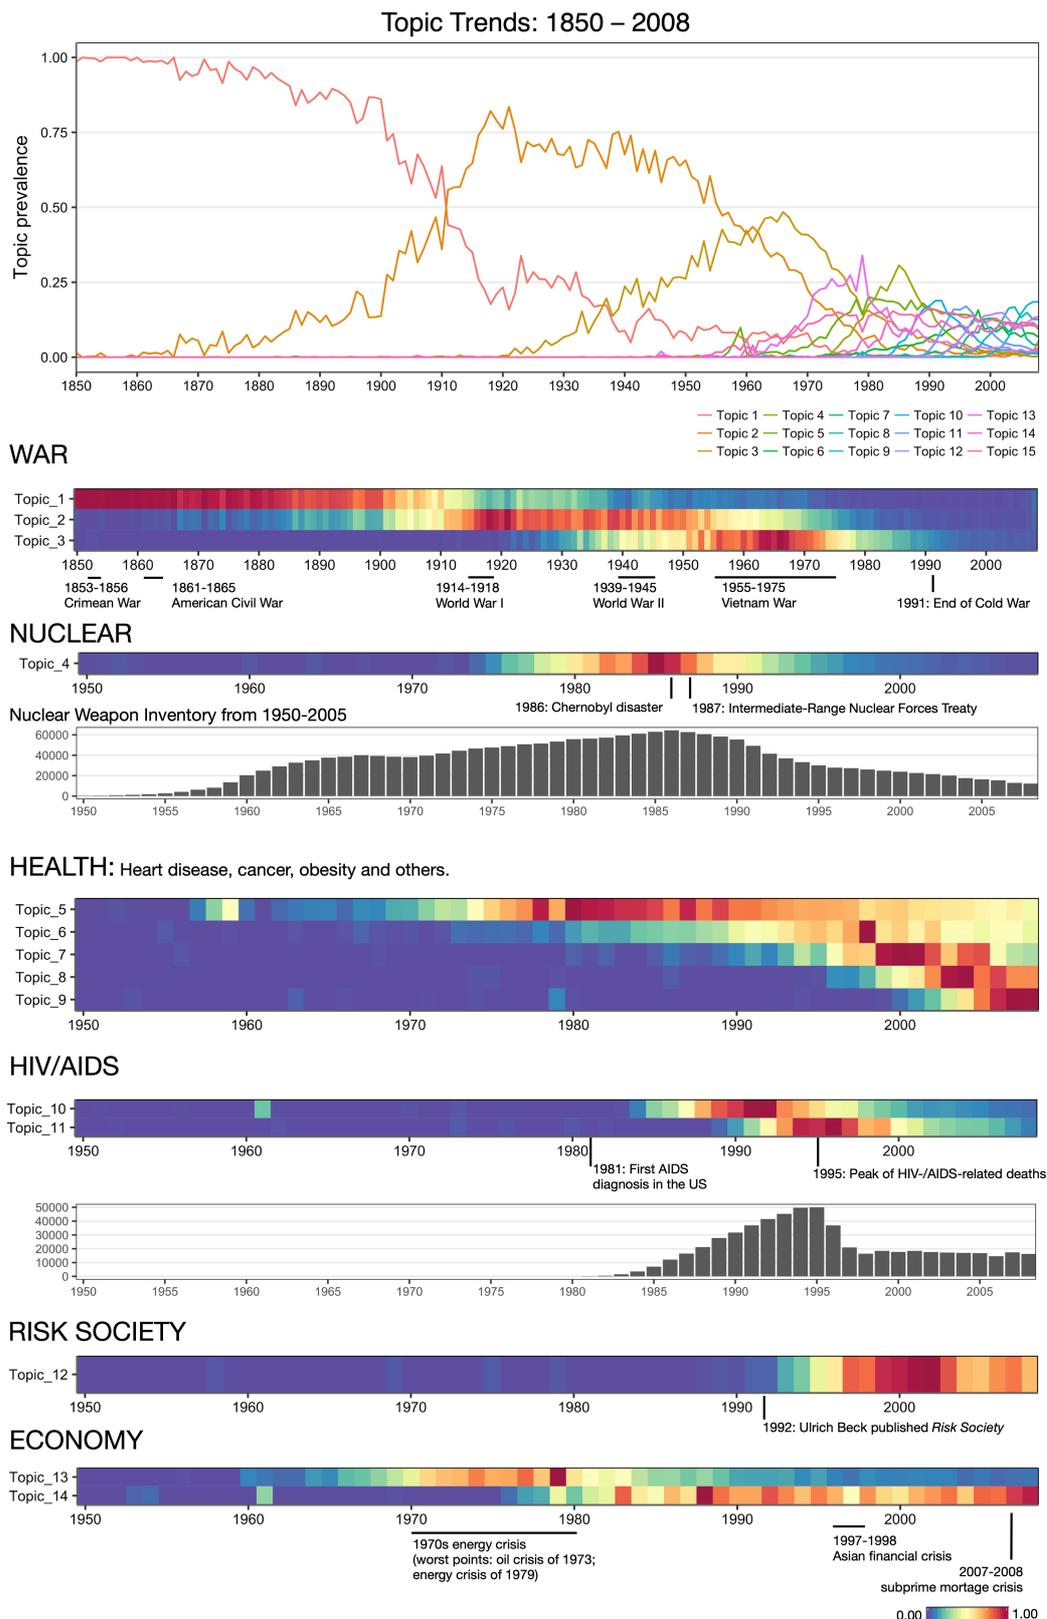
436 4.5 How Are Changes in Risk Categories Associated with Other Events and Developments?

437 One advantage of Google Books Ngram Corpus is that it allows us to investigate change in the
438 sources of risk across a period of over 150 years and to speculate on how those changes relate to other
439 historical events and developments. Specifically, we performed a trend analysis on the topic model
440 derived from the Google Books Ngram Corpus over the years 1850 to 2008. As Figure 4 shows, the
441 structure of the Google Books Ngram risk topics underwent major changes over this period. The three
442 war-related topics emerge early in the distribution: Topic 1 (*life, imminent, battle, resolve*) dominated
443 the risk structure in the second half of the 19th century, which witnessed several major wars (e.g.,
444 Crimean War, American Civil War). Topic 2 (*life, war, bureau, loss*) emerged and reached its peak
445 during World Wars I and II. Topic 3 (*war, uncertainty, loss, prepare*) reached its peak during the
446 Vietnam War. Topic 4 (*nuclear, carcinogenic, patient, infant*) peaked around 1985, capturing the risks
447 associated with the proliferation of nuclear weapons during the Cold War (see the histogram in Figure
448 4) and the growing use of nuclear power in the 1970s and 1980s.

449 Chronic diseases such as heart disease and cancer are now the leading global risks for mortality
450 (World Health Organization, 2009). Topics reflecting this development (topics 5–9) started to emerge
451 from the 1970s and remain the most prominent risk topics. Due to the large proportion of shared words
452 associated with the different health conditions, topics 5, 6, 7, and 8 show considerable overlap, that is,
453 they share words that describe cancer, heart and coronary issues, and other severe diseases. Topic 9,
454 associated with obesity and diabetes, emerged after 2000. The data for topics 10 and 11 show that
455 concerns over AIDS and HIV emerged within 2 years of the first AIDS diagnosis in the US in 1981 and
456 soon reached a peak around 1995, when the reported annual mortality from HIV/AIDS peaked in the
457 United States (Centers for Disease Control and Prevention, 2010). Potentially reflecting the dramatic
458 medical advances in treatments for HIV and the ensuing drop in mortality rates, this risk topic decreased
459 in prominence after 2000 (see the histogram of AIDS-related deaths in the US in Figure 4).

460 Finally, topic 12 (*management, value, assessment, society*) is about management of various
461 social risks. It seems to relate to Beck's conceptualization of the *risk society*, being associated with
462 words such as *Ulrich, Beck, and modernity*. Topics 13 and 14 relate to the economy, and emerged from
463 the 1970s: topic 13 features words like *preference, assumption, equilibrium, and journal*, whereas topic
464 14 features words such as *return, portfolio, and interest*. Lastly, topic 15 (*behavior, group, death,*
465 *population*) seems to be concerned with general risk analysis, without reference to any specific risk
466 event.

467



468
 469 *Figure 4.* Trend analysis on risk topics derived from the Google Books Ngram Corpus. Topics are
 470 grouped into six categories: war, nuclear, health, HIV/AIDS, risk society, and economy. Relevant
 471 historical events are labeled to indicate how changes in the meanings of risk were associated with
 472 historical events and developments. Top panel: historical trends of 15 risk topics (computed using
 473 Equation 4). Bottom panel: normalized topic trend for each individual topic. Topic 15 is not included
 474 as it does not refer to a specific risk topic.

475 **5. Discussion**

476 Risk is a complex multidimensional construct. It takes a variety of forms in public discourse
477 and has, accordingly, been investigated in various ways. Each approach focuses on some aspects of the
478 discourse at the expense of others. One common approach has been to analyze media coverage of risk
479 as a leading source of information for the general public and experts alike (see, e.g., Combs & Slovic,
480 1979, and various references in Young et al., 2008). Our approach consisted in a large-scale analysis of
481 historical text corpora. Such corpora are attractive because they collate a vast array of perspectives on
482 an extensive historical time window: in the case of the Google Book Ngrams Corpus, over 8 million
483 books and 150 years. What did we learn about the risk-related discourse in English-speaking countries?

484 First, we found—consistent with Beck’s (1992) diagnosis of post-industrialist Western
485 societies as risk societies facing a wide variety of unique and human-made risks and with Giddens’s
486 (1990) idea that society is increasingly preoccupied with the future and its safety—that the word *risk*
487 has become much more prevalent (Figure 1A). There is evidence of an approximately fourfold increase
488 in its usage since the 1950s. Beck also stressed that risks in the post-modern world are increasingly
489 unknowable and unpredictable due to scientific and technological innovations having unanticipated
490 consequences. It is possible that this process has contributed to our second major observation, namely,
491 that the sentiments associated with risk have become much more negative, starting around 1900 and
492 confirming Pinker’s (2011) observation that humans have become increasingly preoccupied with the
493 negative aspects of risk. Interestingly, the same does not apply to its close semantic relatives (Figure
494 1C). What is also puzzling is that this change in sentiments is happening at a time when the semantics
495 of risk have become increasingly associated with notions of quantification, reduction, and prevention—
496 findings that also challenge the idea that the increase in negative sentiments has been caused by the
497 unknowability of risks. In addition, we found that the risk categories to some extent reflect real-world
498 changes in the prevalence and magnitude of the respective risks (see Figure 4 and our analyses of
499 nuclear proliferation and AIDS-related deaths). Finally, we also found a shift from macro-risks, such
500 as war and battle, to more individual-specific, chronic risks such as disease (Holzmann & Jørgenson,
501 2000) as well as shift toward more variability in risk topics. The strong focus on modern diseases
502 suggests that the public discourse is generally oriented toward the most prevalent causes of death and
503 harm. This is noteworthy, as several authors have argued that people tend to be afraid of the wrong
504 things (see Glassner, 2018; Renn, 2014; Schröder, 2018).

505 Many of these patterns observed are remarkable in part because they are monotonic: the notable
506 increase in the frequency and negativity of the risk construct, and the increase in number of topics it
507 encompasses. These changes are perhaps related to one another. One potential underlying mechanism
508 is the social amplification of risk (Jagiello & Hills, 2018; Kaspersen et al., 1988; Moussaïd et al., 2015):
509 as information is transferred from one individual to another, people tend to share the more negative
510 aspects of a risk at the expense of potential gains. In Jagiello and Hills (2018), an individual exposed to

511 a balanced argument on nuclear power shared that information with another individual. As information
512 was communicated from one individual to the next, the focus shifted increasingly to the downsides of
513 nuclear power and away from its benefits. This pattern is consistent with the substantial evidence that
514 negative information has more influence on decision making than positive information (Baumeister et
515 al., 2001; Ito et al., 1998; Rozin & Royzman, 2001). A second, related factor is that this effect may be
516 further amplified by increasing communication over the period of our analysis. As Herbert Simon (1971)
517 noted, “a wealth of information creates a poverty of attention” (pp. 40–41). With the unprecedented
518 amounts of information now available, all other things being equal, the absolute amount of negative
519 information has increased. In this environment, information that is better at being received, remembered,
520 and reproduced has a selective advantage (Hills, 2019). This mechanism may apply particularly to
521 information on prominent risks, which may self-reinforce more rapidly via intensified social
522 communication (Jagiello & Hills, 2018).

523 What can be concluded from our results about the state of the public discourse on risk? First
524 and foremost, our analysis can offer only a glimpse of this complex and multi-dimensional construct.
525 Yet, we found results that were both disconcerting and reassuring. Primarily, the increasing prevalence
526 of the word *risk* is an indicator of its growing significance, which is in itself a double-edged sword.
527 Classifying something as a potential risk is likely to burden it with negative sentiments. Yet, branding
528 something a risk also appears to imply the chance of changing our fortune in relation to it. Importantly,
529 the text corpus analyses suggest that risk categories track real threats over the 20th and 21st century,
530 shifting from violent death to chronic disease and major risks for morbidity and mortality in the modern
531 day. In this sense, the risk discourse reflects changes in threats as well as changes in the potential to
532 mitigate them.

533

534

535

536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571

Reference

- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370.
- Beck, U. (1992). *Risk society: Towards a new modernity*. Sage.
- Bhatia, S. (2019) Predicting risk perception: New insights from data science. *Management Science*. Advance online publication. <https://doi.org/10.1287/mnsc.2018.3121>
- Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python: Analyzing text with the Natural Language Toolkit*. O'Reilly Media.
- Blais, A. R., & Weber, E. U. (2006). A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgment and Decision Making*, 1(1), 33–47.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003) Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Bourke, J. (2005). *Fear: A cultural history*. Virago.
- Buechel, S., Hellrich, J., & Hahn, U. (2016). Feelings from the past: Adapting affective lexicons for historical emotion analysis. In E. Hinrichs, M. Hinrichs, & T. Trippel (Eds.). *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH)* (pp. 54–61). The COLING 2016 Organizing Committee.
- Centers for Disease Control and Prevention. (CDC). (2010). *HIV/AIDS surveillance report*. CDC.
- Combs, B., & Slovic, P. (1979). Newspaper coverage of causes of death. *Journalism Quarterly*, 56(4), 837–849.
- Davies, M., Hegedüs, I., & Fodor, A. (2012, November). The 400 million word corpus of historical American English (1810–2009). In *English Historical Linguistics 2010: Selected Papers from the Sixteenth International Conference on English Historical Linguistics (ICEHL 16)*, Pécs, 23-27 August 2010 (Vol. 325, p. 231). John Benjamins Publishing.
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior research methods*, 51(3), 987-1006.
- Douglas, M. (1992). *Risk and blame: Essays in cultural theory*. Routledge.
- Firth, J. R. (1957). *Papers in linguistics 1934–1951*. Oxford University Press.
- Gaissmaier, W., & Gigerenzer, G. (2012). 9/11, Act II: A fine-grained analysis of regional variations in traffic fatalities in the aftermath of the terrorist attacks. *Psychological Science*, 23(12), 1449–1454.
- Giddens, A. (2013). *The consequences of modernity*. John Wiley & Sons.
- Glassner, B. (2018). *The culture of fear: Why Americans are afraid of the wrong things*. Basic Books.
- Greenfield, P. M. (2013). The changing psychology of culture from 1800 through 2000. *Psychological Science*, 24(9):1722–1731.

- 572 Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016). *Diachronic word embeddings reveal statistical*
573 *laws of semantic change*. arXiv. <https://arxiv.org/abs/1605.09096>
- 574 Harari, Y. N. (2016). *Homo Deus: A brief history of tomorrow*. Random House.
- 575 Hills, T. T. (2019). The dark side of information proliferation. *Perspectives on Psychological*
576 *Science, 14*(3), 323–330.
- 577 Hills, T. T., & Adelman, J. S. (2015). Recent evolution of learnability in American English from 1800
578 to 2000. *Cognition, 143*, 87–92.
- 579 Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory.
580 *Psychological Review, 119*(2):431–440.
- 581 Hills, T. T., Proto, E., Sgroi, D., & Seresinhe, C. I. (2019). Historical analysis of national subjective
582 wellbeing using millions of digitized books. *Nature human behaviour, 1*-5.
- 583 Holzmann, R., & Jørgenson, S. (2000). Social risk management: A new conceptual framework for
584 social protection, and beyond. *International Tax and Public Finance, 8*(4), 529–556.
- 585 Hornik, K., & Grün, B. (2011). Topic models: An R package for fitting topic models. *Journal of*
586 *Statistical Software, 40*(13), 1-30. [not in text]
- 587 Ito, T. A., Larsen, J. T., Smith, N. K., Cacioppo, J. T. (1998). Negative information weighs more
588 heavily on the brain: The negativity bias in evaluative categorizations. *Journal of Personality*
589 *and Social Psychology, 75*(4), 887–900.
- 590 Jagiello, R. D., & Hills, T. T. (2018). Bad news has wings: Dread risk mediates social amplification in
591 risk communication. *Risk Analysis, 38*(10), 2193–2207.
- 592 Kasperson, R. E., Renn, O., Slovic, P., Brown, H., Emel, J., Goble, R., Kasperson, J. X., & Ratick, S.
593 (1988). The social amplification of risk: A conceptual framework. *Risk Analysis, 8*(2), 177–
594 187.
- 595 Li, Y., Engelthaler, T., Siew, C. S., & Hills, T. T. (2019). The MacroScope: A tool for examining the
596 historical structure of language. *Behavior Research Methods, 51*, 1864–1877.
597 <https://doi.org/10.3758/s13428-018-1177-6>
- 598 Lin, Y., Michel, J. B., Aiden, E. L., Orwant, J., Brockman, W., & Petrov, S. (2012). Syntactic
599 annotations for the Google Books Ngram Corpus. *Proceedings of the ACL 2012 System*
600 *Demonstrations* (pp. 169–174). Association for Computational Linguistics.
- 601 Markowitz, H. (1952). Portfolio selection. *The Journal of Finance, 7*(1):77–91.
- 602 Michel, J. B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., Hoiberg, D., Clancy,
603 D., Norvig, P., Orwant, J., & Pinker, S. (2011). Quantitative analysis of culture using millions
604 of digitized books. *Science, 331*(6014):176–182.
- 605 Moussaïd, M., Brighton, H., & Gaissmaier, W. (2015). The amplification of risk in experimental
606 diffusion chains. *Proceedings of the National Academy of Science, 112*(18), 5631–5636.
- 607 Oeppen, J., & Vaupel, J. W. (2002). Broken limits to life expectancy. *Science, 296*(5570):1029–1031.

- 608 Pechenick, E. A., Danforth, C. M., & Dodds, P. S. (2015). Characterizing the Google Books corpus:
609 Strong limits to inferences of socio-cultural and linguistic evolution. *PLOS*
610 *ONE*, 10(10):e0137041.
- 611 Pinker, S. (2011). *The better angels of nature: The decline of violence in history and its causes*.
612 Penguin.
- 613 Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the
614 environment. *Journal of Experimental Psychology*, 143(5), 2000–2019.
- 615 Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1–2), 122–136.
- 616 Rayner, S., & Cantor, R. (1987) How fair is safe enough? The cultural approach to societal
617 technology choice. *Risk Analysis*, 7(1):3–9.
- 618 Renn, O. (2014). *Das Risikoparadox: Warum wir uns vor dem Falschen fürchten* [The risk paradox:
619 Why we are afraid of the wrong things]. Fischer Taschenbuchverlag.
- 620 Rothman, L. (2016, January 6). *Why Americans are more afraid than they used to be*. TIME.
621 <https://time.com/4158007/american-fear-history/>
- 622 Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion.
623 *Personality and Social Psychology Review*, 5(4), 296–320.
- 624 Sandhaus, E. (2008). *The New York Times Annotated Corpus*. Linguistic Data Consortium.
- 625 Short, J. F., Jr. (1984). The social fabric at risk: Toward the social transformation of risk analysis.
626 *American Sociological Review*, 49(6):711–725.
- 627 Schönemann, P. H. (1966). A generalized solution of the Orthogonal Procrustes problem.
628 *Psychometrika*, 31(1), 1–10.
- 629 Schröder, M. (2018). *Warum es uns noch nie so gut ging und wir trotzdem ständig von Krisen reden*
630 [Why we've never had it so good and we still constantly talk about crises]. Benevento.
- 631 Simon, H. (1971). Designing organisations for an information-rich world. In Greenberger, M. (Ed.).
632 *Computers, communications, and the public interest* (pp. 40–41). The Johns Hopkins Press.
- 633 Slovic, P. (1987). Perception of risk. *Science*, 236(4799):280–285.
- 634 Slovic, P., Fischhoff, B., & Lichtenstein, S. (1985) Characterizing perceived risk. In Kates, R. W.,
635 Hohenemser, C., & Kasperson, J. X. (Eds.). *Perilous progress: Managing the hazards of*
636 *technology* (pp. 91–125). Westview.
- 637 Stearns, P. N. (2012). *American fear: The causes and consequences of high anxiety*. Routledge.
- 638 Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*. Cambridge University
639 Press.
- 640 Taddy, M. (2012). On estimation and selection for topic models. In T. Sichelman & K. Atkinson,
641 *Proceedings of the Fifteenth International Conference on Artificial Intelligence and*
642 *Statistics* (pp. 1184–1193). Association for Computing Machinery.

- 643 Twenge, J. M., Campbell, W. K., & Gentile, B. (2012). Increases in individualistic words and phrases
644 in American books, 1960–2008. *PLOS ONE*, 7(7):e40181.
- 645 Uz, I. (2014). Individualism and first person pronoun use in written texts across languages. *Journal of*
646 *Cross-Cultural Psychology*, 45(10):1671–1678.
- 647 Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance
648 for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.
- 649 World Health Organization. (2009). *Global health risks: Mortality and burden of disease attributable*
650 *to selected major risks*. World Health Organization.
- 651 Young, M. E., Norman, G. R., & Humphreys, K. R. (2008). Medicine in the popular press: The
652 influence of the media on perceptions of disease. *PLoS One*, 3(10), e3552.