http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–24
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Predicting Risk Perception: New Insights from Data Science

Sudeep Bhatia[a]

[a] Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104
Contact: bhatiasu@sas.upenn.edu, http://orcid.org/0000-0001-6068-684X (SB)

Received: November 20, 2017
Revised: March 28, 2018
Accepted: May 13, 2018
Published Online in Articles in Advance: ■■ ■■, 2018

https://doi.org/10.1287/mnsc.2018.3121

Abstract. We outline computational techniques for predicting perceptions of risk. Our approach uses the structure of word distribution in natural language data to uncover rich representations for a very large set of naturalistic risk sources. With the application of standard machine learning techniques, we are able to accurately map these representations onto participant risk ratings. Unlike existing methods in risk perception research, our approach does not require any specialized participant data and is capable of generalizing its learned mappings to make quantitative predictions for novel (out-of-sample) risks. Our approach is also able to quantify the strength of association between risk sources and a very large set of words and concepts and, thus, can be used to identify the cognitive and affective factors with the strongest relationship with risk perception and behavior.

History: Accepted by Elke Weber, accounting.
Funding: Financial support is from the National Science Foundation [SES-1626825].

Keywords: risk perception • association • knowledge representation • vector semantics • machine learning • big data

Introduction

The perception of risk is influenced by a number of core psychological variables, and much of the research on risk perception over the past four decades has been dedicated to better understanding the nature of this relationship, and its implications for human behavior and public policy (Fischhoff et al. 1978; Johnson and Tversky 1983, 1984; Slovic et al. 1984; Slovic 1987; Fischhoff 1995; Loewenstein et al. 2001; Slovic and Weber 2002; Slovic et al. 2005; Slovic and Peters 2006; McDowell et al. 2016). Despite these important advances, one challenge remains. Risk perception judgments nearly always involve a naturalistic risk source, for which people have rich and complex knowledge representations. These representations guide both emotional and cognitive assessments of the risk source, and subsequently determine the extent to which the risk source is considered hazardous. Although risk perception research has shed light on how core psychological processes involving emotion and cognition interact with risk perception, it is not currently able to specify exactly *what people know* about different risk sources, and how this knowledge can be used to study their perceptions of risk.

Consider, for example, the risk of terrorism. Most people *know* what terrorism is and have strong associations between terrorism and various objects, individuals, and events. This knowledge is a key determinant of how individuals perceive the risks associated with terrorism and how societies respond to these risks. Current approaches to studying risk perception use various psychometric instruments to uncover people's associations and risk perceptions for risk sources like terrorism (see, e.g., Slovic and Weber, 2002, Jenkin 2006, Sanquist et al. 2008). However, the type of data uncovered through these instruments is fairly limited. For example, a typical survey may require participants to evaluate a risk source like terrorism on various dimensions of interest, such as how much dread it elicits, its catastrophic potential, and its ability to be known and controlled. Although these evaluations would shed light on key aspects of how people think and feel about terrorism, the representations obtained from such methods would still be much sparser than those actually possessed by individuals.

A related limitation involves generalization: Existing techniques need to directly survey people in order to uncover relevant risk representations and associations. Thus they cannot be used to make predictions about novel, out-of-sample, sources of risk, for which survey data does not exist. For example, what we know about risk perceptions for terrorism cannot, by itself, be used to predict the perceived riskiness of a related risk source like cyberespionage (without first querying participants for ratings of cyberespionage on various dimensions of interest). As the list of risk sources is endless, with multiple new risks coming to prominence every year, this greatly restricts the descriptive scope of risk perception research.

One implication of this lack of generalizability is that it is difficult to predict both retrospective and real-time perceptions of risk. People's representations of the world around them are continuously evolving. Existing survey-based techniques, which require extensive participant data, are ill suited to tracking these evolving representations, and thus cannot be used to make immediate assessments of changes to perceived risk in response to notable events (e.g., new types of terrorist attacks). Of course, researchers cannot go back in time to administer surveys, implying that the retrospective study of risk perception is nearly impossible. Thus, for example, we cannot track representations and risk perceptions of terrorism before and after the September 11 World Trade Center bombings, with current techniques.

Fortunately, recent developments in data science have made it possible to uncover high quality knowledge representations for a very large set of real world objects and events without the need for participant survey data. The critical idea underlying this approach is that knowledge—what people know about and associate with different objects and concepts—is reflected in the distribution of words in language. This insight has a long history in psychology and linguistics (Harris 1954, Firth 1957). However, it is only recently that large natural language data sets and the corresponding computational resources for analyzing these data sets have been made available to researchers. With these advancements it is now possible to measure word distributions on a very large scale, and subsequently quantify how people think and feel about hundreds of thousands everyday objects and concepts (Landauer and Dumais 1997, Griffiths et al. 2007, Jones and Mewhort 2007, Dhillon et al. 2011, Mikolov et al. 2013, Pennington et al. 2014).

In this paper, we utilize a popular subclass of word distribution-based knowledge representation techniques, commonly referred to vector space semantic models (also known as word embedding models). As suggested by their name, such models specify knowledge representations for objects and concepts using vectors in a high-dimensional *semantic space*. These vectors are typically obtained by performing a form of dimensionality reduction on word co-occurrence data. The use of word co-occurrence to obtain representations implies that similar concepts, which are often talked about in language in a similar manner, have similar vector representations and are closer to each other in the semantic space.

Vector space semantic models have been shown to predict a wide range of low-level behavioral phenomena involving similarity judgment, categorization, free association, text comprehension, and semantic priming (see Bullinaria and Levy 2007 or Jones et al. 2015 for a review). Most recently, such representations also have been found to capture high-level associative judgment, including factual judgment, probability judgment, forecasting, social judgment, political judgment, and moral judgment for a variety of real-world objects and events (Holtzman et al. 2011; Dehghani et al. 2014; Bhatia 2017a, 2017b; Caliskan et al. 2017; Garten et al. 2017; Bhatia 2018; Bhatia et al. 2018). These successes suggest that the semantic vector approach also can be applied to risk perception research. This would involve using word co-occurrence data to derive high-dimensional vector representations for risk sources, such as terrorism, as well as for other objects and concepts related to risk perception. These representations could be used to predict risk ratings, as well as to better understand features of the vector representations that most strongly map onto high ratings of riskiness.

To test this idea, we obtained participant ratings for a large number of risk sources of varying risk levels, including technologies, activities and occupations, and geopolitical forces. We then used a well-known set of semantic vectors (Mikolov et al. 2013), that have, in prior work, been shown to accurately describe knowledge in associative judgment and consumer decision-making (Bhatia 2017a, 2018), to specify high-dimensional vector representations for each of the risk sources. Six different techniques from machine learning were then applied to map these high-dimensional representations to the risk perceptions of participants, and to examine the out-of-sample predictive power of the semantic vector approach. We also considered how high-dimensional vector representations could be used to uncover the conceptual associates of risk, and how these associates could be analyzed in terms of psychological characteristics such as emotion and concreteness. Finally, we used our vector representations to predict participant ratings of the risk sources on commonly studied risk dimensions (Fischhoff et al. 1978), thereby testing whether our vector representations can accurately characterize the key properties that people associate with real-world sources of risk.

Representation and Perception of Risk

The analysis of risk encompasses many techniques, domains, and academic disciplines (see Fischhoff and Kadvany 2011 for a short introduction). One of the most relevant areas of risk analysis involves risk perception. Risk, at its core, is a subjective construct (Krimsky and Golding 1992, Pidgeon et al. 1992, Slovic 1992), and understanding how people perceive risk is necessary in order to understand how risk interacts with the decisions made by individuals, groups, and organizations (Fischhoff et al. 1978, Slovic 1987, Fischhoff 1995, Slovic and Weber 2002). Although there are a number of different ways to study risk perception, in this paper we present techniques for understanding risk perception that most closely complement the psychometric paradigm (Fischhoff et al. 1978, Slovic et al. 1984, Slovic 1987). In the discussion section we outline techniques for extending our approach to analyze sociocultural determinants of risk perception.

In the psychometric paradigm for studying risk perception, individuals are asked to evaluate the riskiness of various risk sources and are also asked to make judgments about other properties of the risk sources. Commonly studied properties include the outcomes (such as deaths) generated by the risk source, and the probabilities corresponding to these outcomes. However, outcomes and probabilities are typically better predictors of expert risk judgment than lay risk judgment (see, e.g., Slovic and Weber 2002). In this paper, our primary emphasis is on the study of lay judgment. This type of judgment has been shown to depend on other, psychologically richer properties, such as the voluntariness, immediacy, knowledge and certainty, controllability, and novelty, of the risk source, the potential for fatal consequences of the risk source, and the dread elicited by the risk source. Different risk sources are associated with different combinations of these properties, and differences in risk judgment across different risk sources can be attributed to differences in the properties associated with the risk sources. Indeed, these properties are also associated with each other in a structured manner, and factor analysis has shown that these determinants of risk perception can be further condensed into a smaller set of higher-order dimensions (such as "dread risk," which characterizes catastrophic, uncontrollable, fatal, and dread-eliciting risks, and "unknown risk," which characterizes unobservable, new, unknown, and delayed risks) (Fischhoff et al. 1978, Slovic et al. 1984, Slovic 1987; see Slovic and Weber 2002 for a review).

The finding that dread is a core dimension of risk perception illustrates the importance of emotional influences in evaluations of riskiness. Prior work has found that allowing for the effects of emotion on risk perception greatly improves the prediction of risk judgments (compared with using outcomes and probabilities alone) (Holtgrave and Weber 1993). Emotion does have a somewhat complex influence on risk perception, in that feelings of risk both influence and are influenced by nonemotional factors (Johnson and Tversky 1983, Loewenstein et al. 2001). Of course emotions interact with each other as well, so that there is a negative relationship between perceived risk and perceived benefit (Slovic et al. 2002, 2005; Slovic and Peters 2006). Although such relationships are occasionally considered to be harmful or irrational, there are many positive aspects of emotional, and, more generally, association-based processing, in the context of risk (see Slovic et al. 2002 for a discussion).

The role of association in risk perception makes it necessary to understand exactly what people associate with different risk sources. This is a special case of a more general question: what do people *know* about different risk sources, and how can this knowledge be uncovered formally studied? Such questions are intimately related to the study of semantic representation in cognitive psychology, and prior work has already applied insights from cognitive psychology to understand the representation of risk sources. One such application involves multidimensional scaling (Kruskal 1964, Borg et al. 2012), which uses similarity ratings across pairs of risk sources to recover the latent dimensions involved in representing the risk sources. Thus, for example, in addition to asking individuals to evaluate the riskiness of a risk source, this technique also asks them to rate the similarity between different pairs of risk sources. Latent dimensions that best capture the structure of variability across similarity ratings are then used to predict risk evaluations and interpret the psychological underpinnings of risk perception (see Johnson and Tversky 1984 for an early application of this method).

Although multidimensional scaling presents a useful technique for uncovering risk representations, it also involves a number of limitations. Most notably, multidimensional scaling requires explicit participant judgments of similarity. This requirement makes generalizing insights recovered through this technique to new risk sources (for which similarity ratings are unavailable) difficult. Relatedly, such representations cannot be used to specify associations between the risk sources and other objects, concepts, and events not typically considered to be sources of risk (for which similarity ratings may not make sense). Finally, the types of representations uncovered through tasks that require explicit participant ratings are much sparser than the representations actually possessed by individuals. All of these issues also limit the applicability of other psychometric techniques, which use ratings on dimensions like voluntariness or dread, rather than similarity ratings between different sources of risk, to specify risk representations and predict risk perception.

Vector Space Semantic Models

What is necessary then is a way of quantifying representations for risk sources that reflects the richness of human representations of risk, and that can subsequently be used to uncover the associations between different sources of risk and the wide range of objects, concepts and events that play a role in the mental lives of individuals. Ideally, such an approach would not rely on explicit participant ratings of similarity or ratings on various risk-related dimensions, and thus could be applied in an a priori manner to predict out-of-sample risk assessments for a large and diverse range of risk sources.

Fortunately, there have been theoretical and technical advances in data science that have made it possible to do this. These advances rely on large-scale natural language data that has recently become available with the growth of the Internet. By examining the structure of word distribution in this data, it is possible to uncover high-dimensional vector representations for words and phrases, and subsequently specify the association between any two words using the distance between their corresponding vectors (Landauer and Dumais 1997, Griffiths et al. 2007, Jones and Mewhort 2007, Dhillon et al. 2011, Mikolov et al. 2013, Pennington et al. 2014). For example, such an approach may use the structure of word distribution in a large natural language data set to represent a concept like *cat* as a high-dimensional vector $\boldsymbol{c}$. It would also represent thousands of other concepts as vectors in the same semantic space, and using the distance between $\boldsymbol{c}$ and these other vectors, this approach would be able to infer the relatedness or association between cats and various objects, events, people, places, and things.

Representations derived using the above semantic vector methods have been shown to successfully predict behavior in tasks involving similarity judgment, categorization, cued recall, and free association (see Bullinaria and Levy 2007, Jones et al. 2015 for a review). Of course, these representations are also desirable for modelling language use in humans, and, in turn, for facilitating natural language processing in machines (Turney and Pantel 2010). Although a lot of the existing work applying these techniques involves core topics in the study of language, memory, and cognition, recently Bhatia (2017a) has extended this approach to study high-level judgment involving real-world object and events, including probability judgment and forecasting. High-level judgment is often associative and vector space models can be used to formalize the associations at play in this domain. Bhatia (2017a) finds that the associations uncovered through vector space semantic models predict participant response probabilities and probability assignments on a variety of existing and novel judgment tasks. These include tasks in which associative judgment leads to errors, as well as tasks in which associative judgment generates correct responses. Bhatia (2017b) and Caliskan et al. (2017) apply a variant of this technique to study the prejudiced and stereotyped associations at play in social judgment, and Bhatia (2018) uses this technique to study how associations influence the objects that come to mind in memory-based decisions. Holtzman et al. (2011), Dehghani et al. (2014), and Bhatia et al. (2018) have also applied distributional models to study political bias, and Garten et al. (2017) recently have used this approach to study morality-based representations in social networks.

The success of the above work suggests that vector space semantic models can be extended to describe knowledge representations for real-world risk sources. In this paper we attempt to do this with semantic vectors obtained using the continuous bag-of-words (CBOW) and skip-gram techniques of Mikolov et al. (2013). This approach relies on a recurrent neural network that, for the CBOW technique, attempts to predict words using other words in their immediate context, and for the skip-gram technique, attempts to the do the inverse of this. In attempting to predict words and contexts in this manner, this approach gradually learns high-dimensional vector representations for the words in the language data. The vector representations are such that words that often co-occur in the same context have similar vectors. Using both the CBOW and skip-gram techniques allows for high quality vectors which can be trained relatively easily and quickly, permitting the use of very large text corpora. These techniques are also desirable as the vectors they generate capture not just individual semantic characteristics of the various words but also relationships between words, which, in turn, facilitates analogical reasoning and other types of more complex similarity-based inferences.

Mikolov and colleagues at Google Research have released a set of high-quality vectors trained using this method, and we use these pretrained vectors for the purposes of this paper. These vectors were generated by applying the Word2Vec tool (available at https://code.google.com/archive/p/word2vec/) on a corpus of Google News articles with over 100 billion words. They have a vocabulary of 3 million words and phrases, with each word or phrase being defined on 300 dimensions. Bhatia (2017a) found that these vectors performed very well in predicting high-level associative judgment. Bhatia (2018) also used these vectors to predict associative biases in memory-based decision-making. Although there are other sets of vectors that achieved equivalent performance in Bhatia (2017a) (notably the GloVe vectors of Pennington et al. 2014) the Word2Vec vectors are unique in their ability to specify multiword phrases in addition to individual words. Thus, for example, these vectors have representations for risk sources like "nuclear power," which is not the case for the GloVe vectors.

In the studies in this paper, we obtain 300-dimensional vector Word2Vec vector representations for a large number of risk sources and use these 300-dimensional vectors to predict risk assignments and other judgments for the corresponding risk sources. We also examine Word2Vec vector representations for a large number of other concepts that are not risk sources, and use the proximity between the risk source vectors and the concept vectors to better understand the emotional and cognitive substrates of risk representation.

Note that in many ways, vector space semantic models are refinements of multidimensional scaling approaches that have already been applied to the domain of risk perception (Johnson and Tversky 1984). However, instead of using explicit similarity ratings from participants, vector space models use similarities in word distributions in natural language. Indeed, one of the earliest (and perhaps the most prominent) vector space semantic models is latent semantic analysis (Landauer and Dumais 1997). Latent semantic analysis performs singular value decomposition on a word-context co-occurrence matrix, to uncover latent dimensions describing the structure of variability in word distribution in language (which is a simpler version of the skip-gram and CBOW techniques used to derive our Word2Vec representations). Mathematically, this is similar to performing a principal components analysis on a matrix of word-word co-occurrence. Multidimensional scaling, relatedly, involves a principal components analysis on word-word similarity judgments to uncover latent dimensions to describe the structure of variability in word similarity (e.g., Kruskal 1964, Borg et al. 2012).

The close relationship between vectors space semantic models like latent semantic analysis, and more classic multidimensional scaling techniques, implies that this paper can be seen as a continuation of previous work (e.g., Johnson and Tversky 1984) that uses theories of semantic representation to study the underlying cognitive structures involved in risk perception and evaluation. Unlike prior work, however, vector space semantic models present a powerful new set of advantages. Not only are explicit participant data on similarity ratings unneeded but also the use of very large natural language data sets implies that such approaches possess rich and comprehensive representations for hundreds of thousands of words and phrases. Analyzing these representations can shed light on the many associations that people have for risk sources, and by doing so, permit greater complexity, depth, and generalizability in the study of risk perception.

Experimental Methods

Overview

We ran three core studies that elicited risk perception judgments from participants for a variety of risk sources. Our primary goal was to use the Word2Vec vector representations for these risk sources to both predict the risk judgments of participants, and to better understand the cognitive and emotional substrates of these judgments. These three studies also separately elicited participant ratings for the risk sources on nine risk dimensions considered to play a role in risk representation and evaluation (Fischhoff et al. 1978). In our primary analysis, we examined the power of our approach in predicting the risk judgments of participants, and compared this with the power of standard psychometric techniques applied to participant ratings on the nine risk dimensions. In a secondary analysis, we used our vector representations, combined with data from a free association task to predict the participant ratings of the risk sources on the nine dimensions.

The first two studies (studies 1A and 1B) served as a preliminary test of our approach. In these studies, we used a relatively small set of experimenter-generated risk sources. Our second study (study 2) served as the primary test of our approach. It involved a large number of participant-generated risk sources, as well as many ratings per risk source. Study 3 administered the free association task, and the data from this study were used to predict participant ratings of the risk sources on the nine risk dimensions.

Participants and Procedures

All participants were recruited from Prolific Academic, an experimental survey website. Participants were residents of the United States with an approval rating of 90% or higher on Prolific Academic. They were paid at a rate of approximately \$7.50 an hour. The average age across our studies was 31.23 (SD = 11.21) and 43% of the participants were female.

In the risk perception task in studies 1A, 1B, and 2, participants judged the riskiness of the risk sources on a scale of −100 (safe) to +100 (risky). There were a total of 73 participants and 125 different risk sources for study 1A, and 79 participants and 125 different risk sources for study 1B. Each participant rated each risk source. In study 2, our primary study, there were 300 participants and 200 different risk sources. Each participant was given 100 randomly chosen risk sources. The order in which the risk sources were presented in all three studies was randomized, and each risk source rating was elicited on a separate screen.

In the dimensional ratings task in studies 1A, 1B, and 2, participants evaluated each of the risk sources on nine dimensions of risk perception using a seven-point scale. These dimensions were taken from Fischhoff et al. (1978) (see also Slovic et al. 1984, Slovic 1987) and correspond to the voluntariness, immediacy of death, knowledge to the person exposed to the risk, knowledge to science, controllability, novelty, the catastrophic potential of the risk, the potential for fatal consequences, and the amount of dread associated with the risk source. In study 1A, there were 75 participants and 125 risk sources for this task. Each participant rated 15 randomly selected risk sources on all nine risk dimensions, generating an average of nine ratings per dimension per risk source. This was also the case for study 1B. In study 2, there were 301 participants and 200 risk sources for this task. Each participant rated 20 randomly selected risk sources on all nine risk dimensions, generating an average of 30 ratings per risk source. The ratings in all three studies were done item-wise, so that participants rated each risk source on all nine dimensions before proceeding to the subsequent risk source. The order in which the sources were presented was randomized.

Participants in studies 1A, 1B, and 2 were given either the risk perception task or the dimensional rating task. Additionally, studies 1A and 1B were administered simultaneously (with random assignment to either 1A or 1B), whereas study 2 was administered a few months later. Prior to running study 2, we ran a study 2 pretest on 52 participants. The participants in this pretest were each asked to generate 15 everyday sources of risk (five with high risk, five with medium risk, and five with low risk). We used the participant-generated risk sources in this pretest as our stimuli in study 2 (see details below).

Finally, in study 3, 49 participants were asked to list words that they associated with the nine dimensions used in risk dimension ratings task. Specifically, each participant was shown a description of the dimension and asked to list three words that first came to their mind, when thinking of a risk source with that description. These could be any words, including words describing actual sources of risk, words describing emotions and feelings, words corresponding to real world objects, people and places, or words describing abstract concepts. For each dimension we used two descriptions, one corresponding to high values on that dimension and the other corresponding to low values on that dimension. Thus, for example, for the voluntariness dimension we used "a risk source that individuals are exposed to voluntarily" and "a risk source that individuals are exposed to involuntarily" as the two descriptions, and each of these two descriptions served as a separate cue in the free association task. The nine dimensions generated a total of 18 descriptions, which were administered in a random order on separate pages to participants.

We had aimed for 150 total participants each in studies 1A and 1B, 600 total participants in study 2, 50 participants each in the studies 2 pretest and study 3, for a total of 1,000 participants across all our studies. Although final sample sizes diverged slightly from these numbers, we did not exclude any participants or observations. We used large sample sizes to ensure minimal noise in resulting aggregate risk ratings.

Stimuli

The stimuli used in study 1A consisted of a set of 125 technologies of varying risk levels (see Figure 1). This set was experimenter-generated, and was based on the technologies used in Slovic's (1987) experiment. It contained various common technologies, emerging technologies, military technologies, household appliances, energy sources, drugs, and medical procedures. The stimuli used in study 1B consisted of a set of 125 activities of varying risk levels (see Figure 1). This set was also experimenter-generated, and was based on the activities in Slovic's (1987) experiment. It contained various hobbies, sports, and occupations. All items used in studies 1A and 1B are present in the Word2Vec vector vocabulary, implying that we were able to obtain vector representations for the items for our subsequent analysis.

The stimuli used in study 2, our primary study, consisted of a set of 200 risk sources, of varying risk levels. This set was generated in the study 2 pretest, described above. The items obtained in this pretest

Figure 1. Scatterplot of the 432 Unique Risk Sources Used in Studies 1A, 1B, and 2, on the First Two Principal Components of the Matrix of Vector Similarities for the Risk Sources

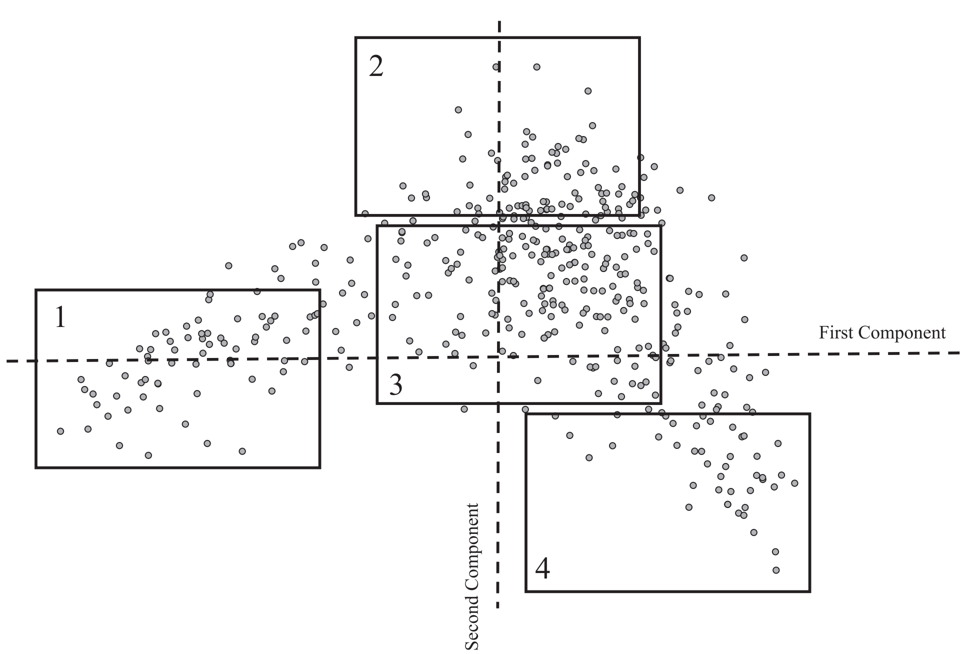

Figure 1. (Continued)

Figure 1. (Continued)

were pruned to select the 200 most frequently listed risk sources which were also present in the Word2Vec vocabulary. Overall, the list of risk sources in study 2 was much more diverse than studies 1A and 1B, as it was participant-generated without any category limitations. It included a range of medical, technological, social, political, climate-based, activity-based, and occupational risks.

Overall, there were a total of 450 risk sources across our three studies. Out of this, 432 risk sources were unique (the remaining 18 risk sources were generated by participants in study 2-pretest, but were also contained in the set of stimuli used for studies 1A and 1B). We had a 300-dimensional vector representation for each of these risk sources, obtained from the Word2Vec database, and we used these

300-dimensional vectors to predict risk perceptions of participants.

Prior to discussing our results, it is useful to briefly explore what the vector representations for the risk sources looked like. Of course, it is difficult to fully visualize the 300-dimensional space within which these representations reside. But we can project these 300-dimensional representations onto a two-dimensional subspace, and use this subspace to understand the relationships between the different risk sources. For this purpose, we first calculated the cosine similarities between the vectors for each pair of risk sources in our stimuli. This metric specifies the proximity between any two vectors x and y by $\text{sim}(x, y) = x \bullet y/(\|x\| \bullet \|y\|)$, and ranges from −1 (for vectors with opposite directions) to 0 (for vectors that are orthogonal) to +1 (for vectors with identical directions). This calculation resulted in a 432×432 matrix with pairwise similarity measures for our 432 unique risk sources. We performed a principal components analysis on this matrix to extract the two largest components of this matrix. Figure 1 presents a scatterplot of the 432 unique risk sources on these two components, with risk source names for four subregions of this scatterplot. This figure reveals certain regularities in the representation of these risk sources. For example, region 1 primarily consists of hobbies and sports activities and region 4 consists primarily of drugs and medical risks. In regions 2 and 3, we frequently see clusters of related risk sources, though these regions are somewhat less coherent. This is likely because the two-dimensional projection omits a lot of relevant information possessed by the 300-dimensional structure (this is also why it is difficult to interpret the two first principal components in terms of intuitive dimensions).

Predicting Risk Perception

Computational Methods

We began our analysis by examining how well our vectors representations predicted the risk judgments of participants, obtained in studies 1A, 1B, and 2. We did this both on the aggregate level (for which we averaged risk perceptions for each risk source to get a single continuous measure of riskiness of the source), as well as on the individual level. In both these cases, we had a rating y_i for a risk source i. We also had, for each source, a 300-dimensional vector representing the source. For source i, we write this vector as $\boldsymbol{x}_i$, where $\boldsymbol{x}_{ij}$ is the value of the risk source on dimension j of the corresponding Word2Vec vector.

Our goal was to predict y_i from the $\boldsymbol{x}_i$. This is a type of high-dimensional regression problem for which there are more independent variables (300) than there are observations (125 in studies 1A and 1B, and 200 in study 2; i.e., one observation for each risk source). This dimensionality problem implies that standard linear regressions are not feasible. Instead, we must use techniques in machine learning that have been shown to work well in these settings. We use six different machine learning techniques. Our first three techniques are support vector regressions, which are able to learn nonlinear relationships between our Word2Vec vectors $\boldsymbol{x}_i$ and our risk ratings y_i. These regressions use a "kernel trick" (which implicitly maps the inputs $\boldsymbol{x}_i$ into high-dimensional feature spaces) to permit nonlinearity. For the purposes of this paper, we considered support vector regressions with radial basis function kernels (SVR-RBF), polynomial kernels (SVR-polynomial), and sigmoidal kernels (SVR-sigmoid). We also considered two other regression-based techniques with linear mappings: Lasso regression and ridge regression. These two perform simple linear least squares, but penalize the coefficients of the inputs $\boldsymbol{x}_i$ based on their size (for this reason, they are also sometimes known as "regularized regressions"). The type of penalty applied by the lasso regression forces a large number of regression coefficients to zero, and thus results in the use of only a small subset of the vector dimensions (this is not necessarily the case for ridge regression). Our final technique was the k-nearest neighbors (KNN) regression. In contrast to the other approaches, KNN uses similarity with previous observations to predict y_i given a new observed $\boldsymbol{x}_i$. Thus, it can be seen as an exemplar-based heuristic technique.

All of our techniques were implemented in the Scikit-Learn Python machine learning library (Pedregosa et al. 2011). There was one key metaparameter in our implementation of the first five techniques. This parameter, $C > 0$, determines the size of the penalty on large errors (it can be seen as a type of regularization parameter). In our estimation, we allowed C to take values in the set $\{10^{-2}, 10^{-1}, 10^{0}, 10^{1}, 10^{2}, 10^{3}, 10^{4}, 10^{5}, 10^{6}, 10^{7}\}$ and evaluated the predictive power of our various techniques for different values of C. We also allowed one free metaparameter for our implementation of KNN. This parameter, $k > 0$, determines the number of nearest neighbors used in the KNN regression, and was allowed to take on values in the set $\{1, 2, \ldots 10\}$. Both the best-fit mappings for these six techniques, as well as the various metaparameters, were inferred through cross validation (which we describe in detail below). We needed to make additional assumptions when applying our techniques. For simplicity we adopted all the default assumptions of the Scikit-Learn library and avoided tweaking or modifying the regression to improve our fits to the data. Of course, it is likely that such modifications or a more fine-grained search through our metaparameters would improve the performance of our approach.

It is also useful to compare the semantic vector approach outlined here to standard methods used in research on risk perception. As discussed above, the most common approach within the psychometric paradigm involves regressing the ratings of the risk sources on nine different dimensions of interest. To apply this approach, we averaged the ratings for each risk source on each dimension, elicited in our studies, to get a single nine-dimensional vector of ratings for each risk source. For source i, we write this vector as $\boldsymbol{x_i}$, where x_{ij} is the average participant rating of the risk source on dimension j. We then used a simple linear regression to predict the perceived riskiness of the risk sources, y_i, from the corresponding x_i.

We measured the predictive power of our fits for both the six machine learning techniques that are part of the semantic vector approach, and the linear regression that is part of the psychometric approach, using the coefficient-of-determination, R^2, in out-of-sample predictions. More specifically, we divided our data set into 10 equal portions. We performed our fits on the risk sources in the first nine portions and then used the fitted models to predict risk assignments for the risk sources in the tenth portion. This cross-validation exercise was repeated 1,000 times, with a random split at each time, to calculate the average R^2 for the out-of-sample predictions. The appendix provides additional details of the methods outlined above, with the help of an example.

Results: Comparing Approaches

Figures 2(a)–2(c) show the out-of-sample R^2 for our best-fitting SVR-RBF, SVR-polynomial, SVR-sigmoid, lasso regression, ridge regression, and KNN regression techniques on the semantic vectors for aggregate risk perception judgments in our three studies. They also plot the out-of-sample R^2 for the psychometric technique, which involves linear regressions on participant ratings on the nine risk dimensions. Cross validation ensures that our modeling avoids overfitting and that performance is evaluated based on model generalizability, rather than flexibility. It is also useful to note that the two sets of approaches rely on different sets of data (preexisting semantic vector representations versus explicit participant-generated ratings on nine key risk dimensions).

For study 1A, the best performing technique for the semantic vector approach, SVR-RBF, was able to achieve a predictive accuracy rate of $R^2 = 0.50$. This is slightly lower than the predictive accuracy rate of $R^2 = 0.54$ for the psychometric approach. There was a somewhat larger discrepancy in the predictive accuracy rates in study 1B. Although the semantic vector approach achieved an accuracy rate of $R^2 = 0.58$ with the SVR-RBF and the ridge regression techniques, the standard psychometric approach was able to predict

Figure 2. Accuracy at Predicting Out-of-Sample Aggregate Risk Ratings Using Either Participant Ratings on Nine Risk Dimensions (Psychometric Approach) or Machine Learning Techniques Applied to Semantic Vectors

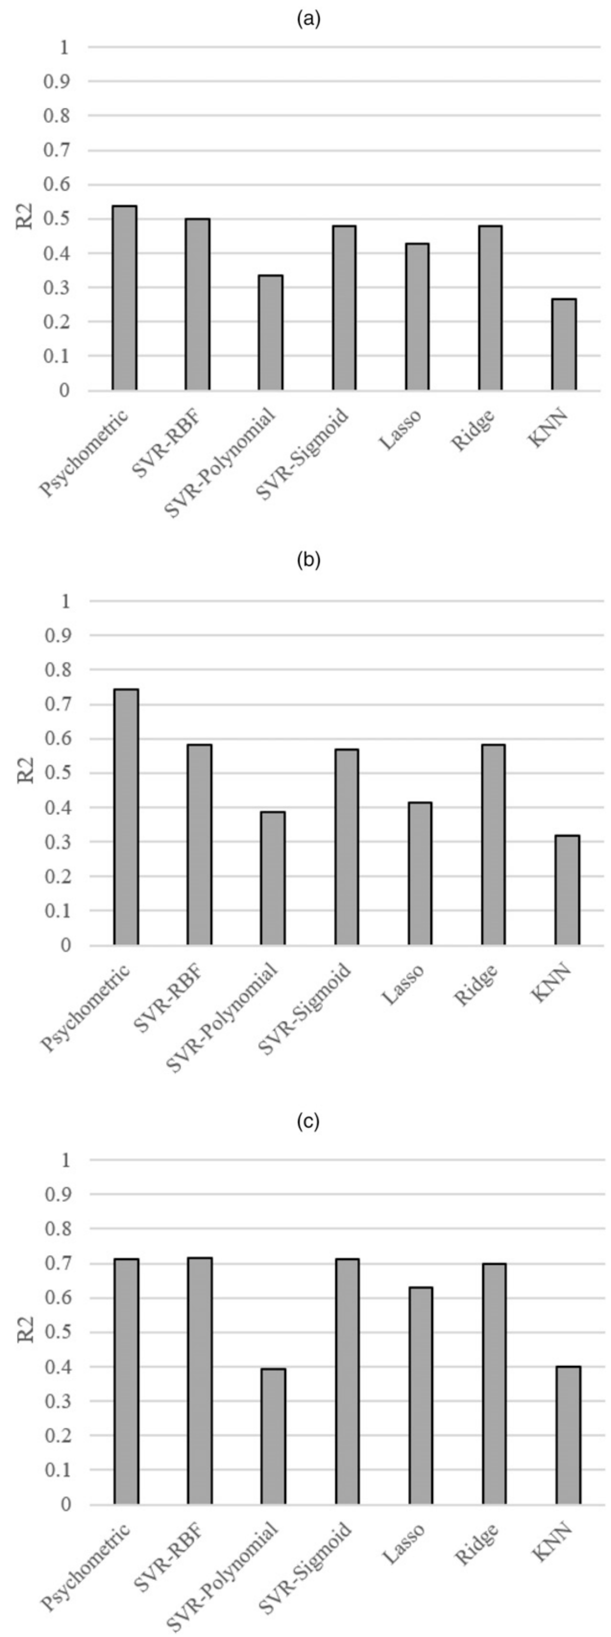

behavior with an $R^2 = 0.74$. Finally, in study 2, our primary study, we found that the semantic vector approach slightly surpassed the psychometric approach, with an accuracy rate of $R^2 = 0.72$ for the SVR-RBF and SVR-sigmoid techniques, compared with an accuracy rate of $R^2 = 0.71$ for the psychometric technique. The predictive power of the semantic vector approach emerges despite the fact that it does not use explicit participant ratings to quantify the values of the risk sources on the relevant risk dimensions.

It is not immediately clear why there was a larger discrepancy between the two approaches in the accuracy rates for study 1B compared with studies 1A and 2. One possibility is that these differences reflect stimuli similarity. As discussed above, study 2 contained a very diverse set of participant-generated stimuli. Study 1A also contained a wide range of different types of technologies (including medicines, household appliances, and energy sources). In contrast, study 1B consisted of mainly hobbies and sports as well a few occupations. Although stimuli similarity is likely to improve the out-of-sample predictive power of both sets of approaches, it is perhaps more beneficial for approaches that rely on participant evaluations, as with the standard psychometric paradigm: Participant ratings in the psychometric tasks are less variable and more internally coherent with similar or comparable risk sources.

In addition to allowing us to compare the semantic vector approach against the psychometric approach, the above tests also permit a comparison between different machine learning techniques for using the semantic vectors. As shown in Figures 2(a)–2(c), the SVR-RBF, SVR-sigmoid, and the ridge regressions are consistently the top techniques, and achieve almost the same accuracy rates in the three studies. In contrast, the SVR-polynomial, lasso regression, KNN regressions achieve much lower accuracy rates.

We obtain very similar results when attempting to predict individual-level risk perception data, with two important caveats. First, the average accuracy rates for both the semantic vector approach and the psychometric approach are smaller than the analogous rates for predicting aggregate data. Second, the discrepancy between the semantic vector and the psychometric approach for study 1B is diminished. Overall, for all three studies, we find statistically indistinguishable average accuracy rates for the psychometric approach compared with the best performing machine learning technique applied to the semantic vectors ($p > 0.05$ for all three studies when comparing individual-level accuracy rates with a t-test). Again, out of the six machine learning techniques, SVR-RBF, SVR-sigmoid, and the ridge regression, are consistently the top techniques for the three studies. These results are illustrated in Figures 3(a)–3(c).

Figure 3. Accuracy at Predicting Out-of-Sample Individual-Level Risk Ratings Using Either Participant Ratings on Nine Risk Dimensions (Psychometric Approach) or Machine Learning Techniques Applied to Semantic Vectors

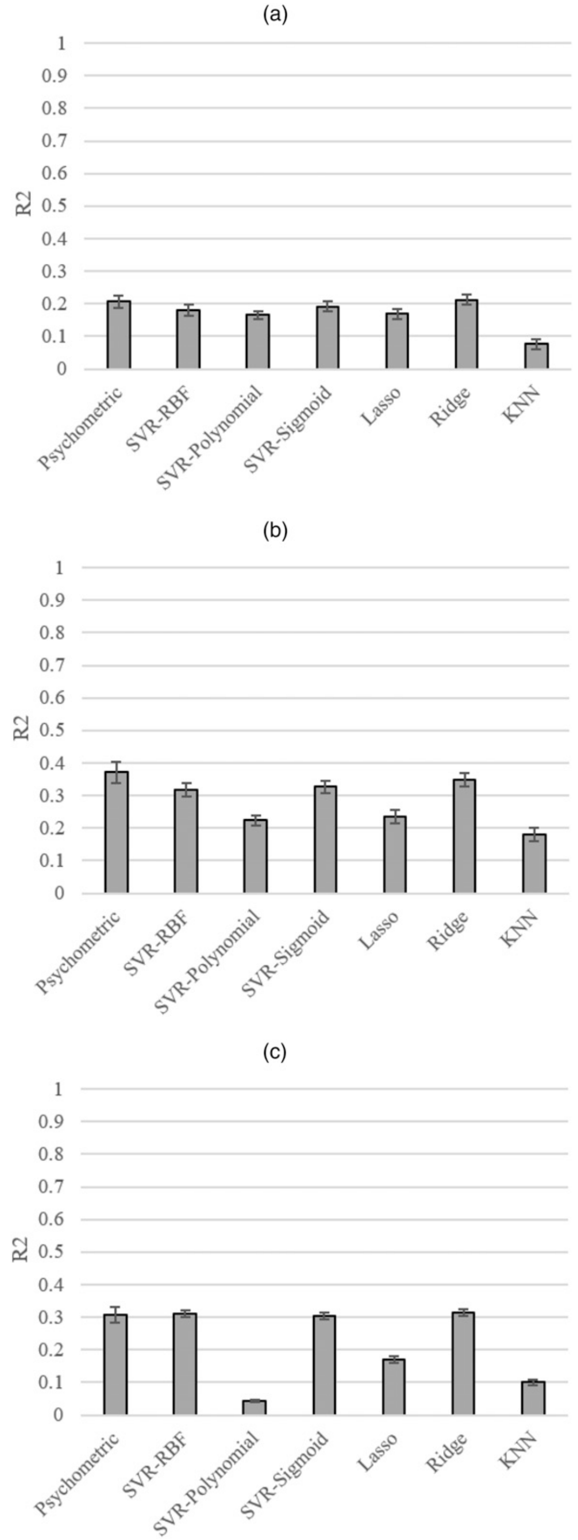

Note. Error bars represent ±1 SE for individual-level accuracy rates.

Results: Combining Approaches

The above analysis attempted to directly compare the proposed semantic vector approach against the standard psychometrical approach, in order to establish the relative predictive power of the semantic vector approach. However, it is likely that the two approaches are in fact complementary and that the highest accuracy rates can be attained by combining the 300-dimensional vector representations with the nine-dimensional participant risk ratings. Such an analysis would attempt to predict the aggregate or individual-level participant risk rating y_i for each risk source i with a 309-dimensional vector $\boldsymbol{x_i}$. Here, x_{ij} would be the value of the risk source on the jth semantic vector dimension for $j \leq 300$. For $301 \leq j \leq 309$, j would capture the rating of the risk source on 300-jth dimension of the risk dimension rating task.

We implemented the combined analysis using the six machine learning techniques and cross-validation procedure described above. The results of this analysis are plotted in Figures 4(a)–4(c) for predictions on aggregate data, and Figures 5(a)–5(c) for predictions on individual-level data, alongside the highest out-of-sample R^2 values from just the semantic vector approach and just the psychometric approach. Here, we see that the 309-dimensional vectors, comprising both the semantic vector representations and participant dimension ratings, greatly exceeded the predictive power of the two sets of independent variables alone. Overall, we were able to achieve out-of-sample accuracy rates of $R^2 = 0.74$, $R^2 = 0.83$, and $R^2 = 0.86$ for studies 1A, 1B, and 2 on the aggregate data, and average out-of-sample accuracy rates of $R^2 = 0.37$, $R^2 = 0.50$, and $R^2 = 0.52$ for studies 1A, 1B, and 2 on the individual data. Note that in the latter case, the accuracy rates obtained using the combined semantic vector and psychometric ratings data are significantly higher than the accuracy rates obtained using either the semantic vectors or the psychometric ratings alone ($p < 0.01$ for all three studies using a t-test on individual-level accuracy rates). These high accuracy rates suggest that the information contained in the semantic vectors and the information elicited through the dimension ratings task are complementary and that both play a somewhat independent role in predicting the risk evaluations of participants. In a later section we will examine the relationship between these two types of representations in more detail.

Figure 4. Accuracy at Predicting Out-of-Sample Aggregate Risk Ratings Using Machine Learning Techniques Applied to a Combination of Semantic Vector Representations and Participant Ratings on Nine Risk Dimensions

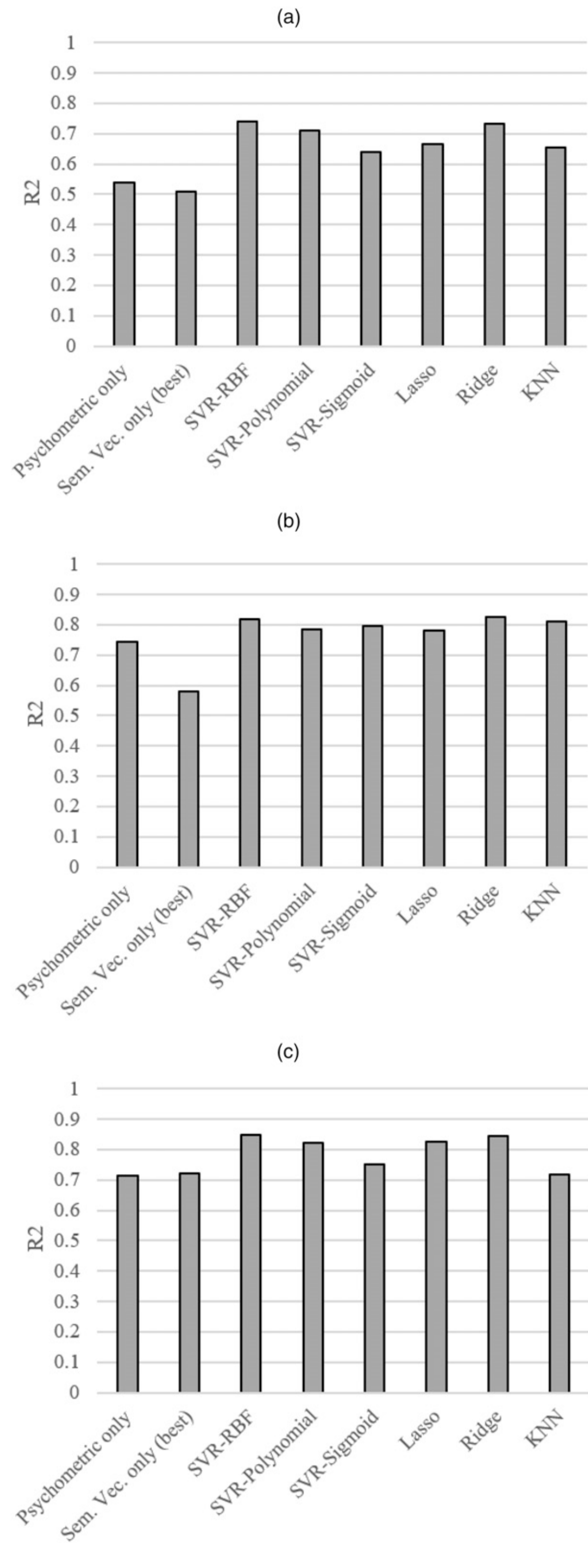

Note. For comparison, "Psychometric only" and "Sem. Vec. only (best)" indicate accuracy rates using only the psychometric approach and using only the best-performing semantic vector approach (taken from Figure 2).

Risk Associations

Computational Methods

The analysis in the previous section showed that the 300-dimensional vector representations for the risk sources can be mapped accurately onto the risk judgments of participants. In this section we wish to better

Figures 5. Accuracy at Predicting Out-of-Sample Individual-Level Risk Ratings Using Machine Learning Techniques Applied to a Combination of Semantic Vector Representations and Participant Ratings on Nine Risk Dimensions

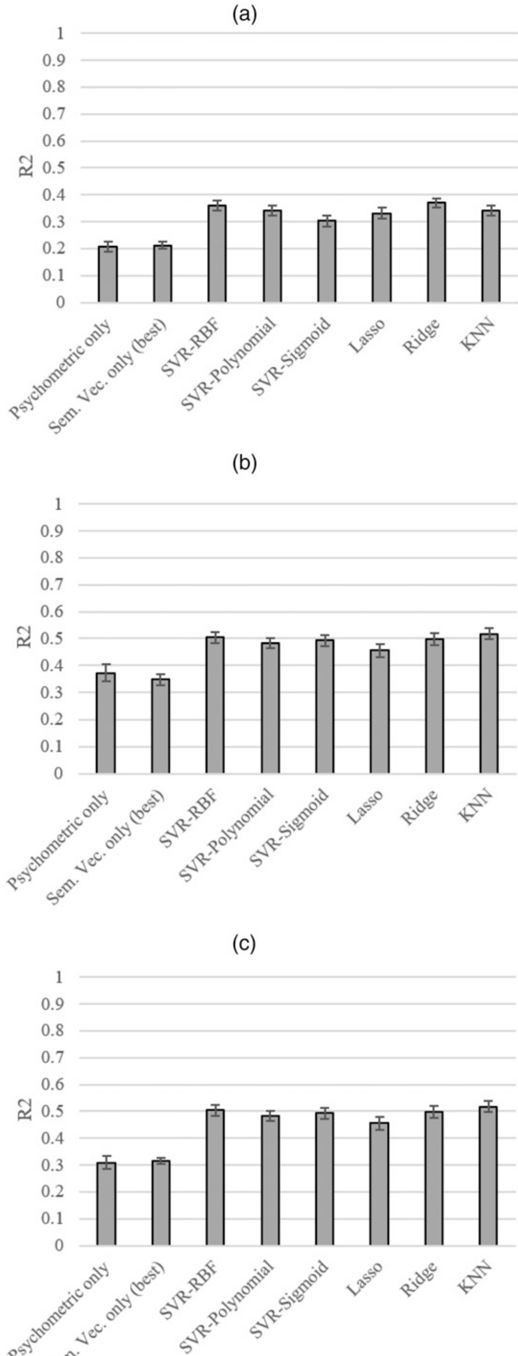

Notes. For comparison, "Psychometric only" and "Sem. Vec. only (best)" indicate accuracy rates using only the psychometric approach and using only the best-performing semantic vector approach (taken from Figure 3). Error bars represent ±1 SE for individual-level accuracy rates.

understand the conceptual underpinnings of this mapping. What are the features of the risk representations that best predict participant risk judgments?

One way to address this question would be to examine the relative weights assigned to the 300 vector dimensions by our various machine learning techniques, so as to understand the dimensions with the greatest influence on risk perception. However, the dimensions themselves do not have an intuitive interpretation, and such an analysis is unlikely to yield intuitive results. A better solution to this issue thus involves analyzing many of the other words and phrases that are present in the Word2Vec vocabulary. Each of these words is described with a 300-dimensional vector, and we can use vector distance to determine the associations between these other words and our various risk sources. Subsequently, we can map the words themselves onto our participants' risk ratings, to determine the words with the strongest association with risk.

For the analysis in this section we specified the association between words and risk sources using the cosine similarity metric introduced above. For a risk source i with vector s_i and a word j with vector w_j, this metric gives us $\text{sim}(s_i, w_j) = s_i \bullet w_j / (\|s_i\| \bullet \|w_j\|)$, which ranges from −1 to +1. Now recall that we have risk ratings for each risk source, and thus for each word we can calculate the correlation between its associations with the risk sources and the risk ratings of the risk sources. More specifically for each word j, we can measure the correlation between $\text{sim}(s_i, w_j)$ and y_i, using all the risk sources i in the given study. This risk association value captures how strongly the word is associated with risk sources high in risk relative to risk sources low in risk. Words with the highest risk associations are words that are disproportionately associated with risk sources that are rated as being extremely risky. The appendix provides additional details of our approach, with the help of an example.

Results: Common Words

As a first test we considered a list of 5,000 words with the greatest frequency in American English, obtained from the corpus of contemporary American English (http://corpus.byu.edu/coca/). All, except for 18, of these 5,000 words are present in our Word2Vec vocabulary, and we were able to obtain vector representations for these words and subsequently calculate the risk associations of the words, using the techniques specified above. Figures 6(a)–6(c) show word clouds of words with the strongest risk associations for each of the three studies. These word clouds each take the 50 words with the strongest risk associations, exclude all

Figure 6. Word Cloud of Words with the Strongest Risk Associations for Each of the Three Studies

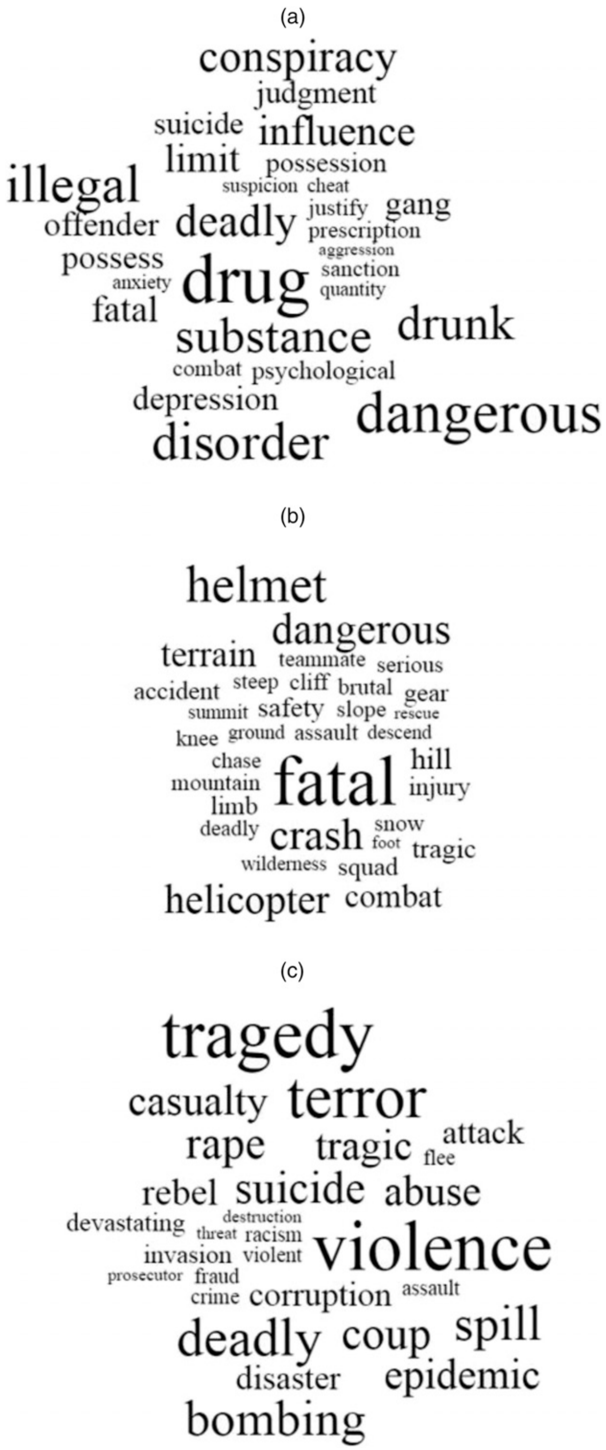

words that refer to the risk sources used in the three studies, and pool together identical words with different syntactic roles (e.g., "abuse" as a noun and "abuse" as a verb) resulting in around 30 words in each cloud. The font size of the words maps linearly to the computed risk association, so that larger words in the word clouds are more associated with risk.

These figures show numerous distinct conceptual associates of risk. Some of the associates occurred in multiple studies (e.g., "fatal," "dangerous," and "tragic"). These likely reflect common features of risk across different risk domains. Indeed, there were substantial similarities in the risk associations of words across our three studies. For example, the risk associations for the 5,000 words in study 1A and the risk associations for these words in study 1B were correlated with a strength of $\rho = 0.28$ ($p < 0.001$). The analogous correlations between studies 1A and 2 and between studies 1B and 2 were $\rho = 0.50$ ($p < 0.001$) and $\rho = 0.28$ ($p < 0.001$). Thus, on aggregate, our approach finds that a similar set of words tend to be associated with risk across the three data sets.

That said, many of the associates in Figures 6(a)–6(c) are also unique. This likely reflects the differences in underlying stimuli. Study 1A had unique associates like "drug," "disorder," and "substance," corresponding to the relatively large number of medical risks (drugs, medical procedures) present in the stimuli set. Study 1B, which involved various risky sports and occupations had unique associates like "crash," "combat," and "helmet." Study 2, involved participant-generated risk sources, and had unique associates like "terror," "violence," and "bombing" (perhaps reflecting the contemporary focus on terrorist attacks and related forms of geopolitical conflict).

Results: Emotion Words

Although the word clouds in Figures 6(a)–6(c) provide us with an intuitive understanding of the close associates of risk, a more rigorous analysis would use the above techniques to quantitatively specify the risk association for a number of different psychological variables of interest. Such an analysis would also compare various psychological variables against each other, to determine the variables with the greatest relationship with the perception of risk. Ideally, such an analysis should also reveal consistent effects for the psychological variables considered, across our three studies.

We first attempted this analysis for a list of six emotions commonly studied in psychology: anger, disgust, happiness, fear, sadness, and surprise. Emotional processing has been strongly implicated in the perception of risk and the proposed analysis can be used to examine which emotions have the strongest associations with risk. For this purpose, we used a database of English language words with binary ratings (1 if the word is related to the emotion in consideration; 0 otherwise) on the six emotions (Mohammad and Turney 2013). This database was generated by a large-scale crowdsourced study on Amazon Mechanical Turk, and consists of over 14,000 words. For each of these words we used the techniques specified at the start of this section to generate a measure of risk association. We then tested how well this measure predicted the emotion rating of the word. This was done with logistic regressions, applied individually to each emotion. These regressions set the emotion of the word as the dependent variable (1 if the word is linked to the corresponding emotion in the Mohammad and Turney emotion ratings data set, and 0 otherwise) and the risk association of the word as the independent variable. The coefficients obtained from these regressions capture the relationship between the emotions and the risk associations of the words. Large positive coefficients indicate emotions that are most likely to be elicited by the conceptual associates of risk, and thus emotions that are most likely to be present when individuals give high risk evaluations. Large negative coefficients indicate emotions that are least likely to be elicited by the conceptual associates of risk, and thus emotions that are least likely present when individuals give high risk evaluations. The appendix provides additional details of our approach, with the help of an example (see also Garten et al. 2017, Pennebaker et al. 2001 for a discussion of related techniques).

The results of these six regressions for the three studies are displayed in Figures 7(a)–7(c) (and details of regression outputs are provided in Table 1). These figures present bar plots with regression coefficients for each of the regressions. These three figures reveal strong regularities across our studies. As expected we obtained a positive regression coefficient fornegative emotions (anger, disgust, fear and sadness) and a negative regression coefficient for positive emotions (happiness). This corresponds to prior work showing that negative affect leads to increased perceptions of risk relative to positive affect (e.g., Johnson and Tversky 1983). Additionally, we obtained systematic differences across the emotions: Across all studies, fear was the emotion with the strongest positive relationship with risk judgment. Although anger was also significantly associated with risk for all three studies, the strength of its effect was much milder. Other negative emotions like disgust and sadness did not have significant effects in all our studies. This again corresponds to prior results showing that fear is the emotion with the strongest effect on the perception of risk (Lerner and Keltner 2001, Loewenstein et al. 2001, Lerner et al. 2003, Slovic and Peters 2006).

Figures 7. Coefficients for Logistic Regressions of Emotions on Risk Associations

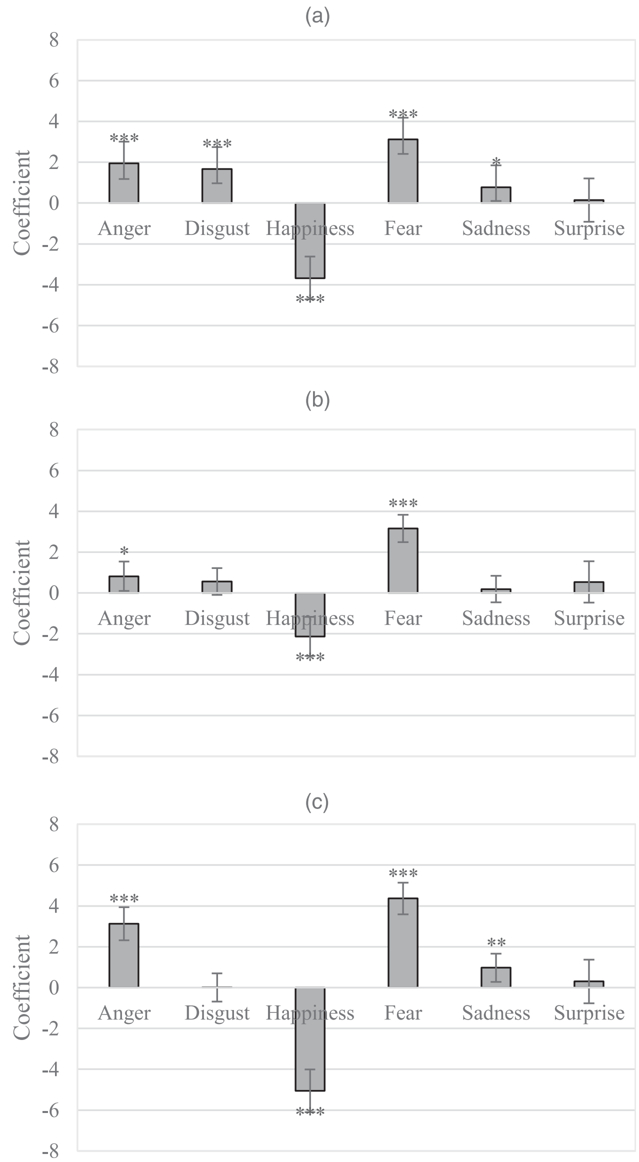

Notes. Positive coefficients indicate emotions that are most likely to be elicited by the conceptual associates of risk. Error bars capture 95% confidence intervals for regression coefficients.
$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

Results: Concreteness

To illustrate the broad applicability of the approach outlined in this paper, we applied it to another psychological variable: concreteness/abstraction. Although risk perception is not often studied in terms of concreteness, some prior work does suggest that risk sources that vary in terms of their concreteness could be associated with different levels of riskiness. For example, temporal distance, an important determinant of whether objects and events are construed concretely or abstractly, has been shown to influence risk perception (Chandran and Menon 2004, Trope and Liberman 2010). Particularly, risks communicated as being temporally proximate are imagined more concretely and are considered more probable. Relatedly, vividness has been shown to play a key role in the strength of emotional responses to stimuli, subsequently influencing risk perception: risks that are vivid elicit stronger feelings and (if the feelings are negative) are considered to be more likely (Loewenstein et al. 2001).

Table 1. Details for Regressions of Emotion Ratings on Word Risk Associations for the Three Studies

	Coef.	SD	*z*	*p*	95% CI-L	95% CI-H
Study 1A						
Anger	1.94	0.39	4.96	0.00	1.18	2.71
Disgust	1.67	0.36	4.67	0.00	0.97	2.37
Happiness	−3.68	0.52	−7.04	0.00	−4.71	−2.66
Fear	3.12	0.37	8.53	0.00	2.40	3.83
Sadness	0.78	0.35	2.24	0.03	0.10	1.46
Surprise	0.15	0.54	0.27	0.79	−0.92	1.21
Study 1B						
Anger	0.82	0.37	2.24	0.03	0.10	1.53
Disgust	0.56	0.34	1.67	0.10	−0.10	1.22
Happiness	−2.13	0.49	−4.36	0.00	−3.09	−1.17
Fear	3.15	0.34	9.20	0.00	2.48	3.83
Sadness	0.19	0.33	0.57	0.57	−0.46	0.83
Surprise	0.54	0.52	1.04	0.30	−0.47	1.55
Study 2						
Anger	3.12	0.41	7.55	0.00	2.31	3.94
Disgust	0.01	0.35	0.02	0.98	−0.69	0.70
Happiness	−5.06	0.53	−9.48	0.00	−6.11	−4.01
Fear	4.36	0.39	11.06	0.00	3.59	5.14
Sadness	0.97	0.35	2.77	0.01	0.29	1.66
Surprise	0.30	0.55	0.55	0.58	−0.77	1.37

Notes. Here, positive coefficients indicate emotions that have a positive relationship with risk, and negative coefficients indicate emotions that have a negative relationship with risk. Coef., coefficient; CI, confidence interval (L = Lower, H = Higher).

The relationship between concreteness and risk perception can also go in the opposite direction. One of the two core dimensions underlying risk perception involves the knowledge of the risk, and risk sources that are unknown to the individual exposed to the risk or unknown to science are considered to be riskier (Slovic 1987). As concrete risk sources are easier to imagine and are thus more "known" it is possible that concreteness actually has a negative effect on risk perception.

We tested the relationship between risk perception and concreteness using a data set of concreteness word ratings compiled by Brysbaert et al. (2014). This data set has participant ratings for 40,000 English words on a scale of 1 (abstract) to 5 (concrete). Of these 40,000 words, 34,121 are present in our Word2Vec vocabulary. For each of these words it is possible to compute the risk association with the risk assignments of participants in studies 1A, 1B, and 2, and subsequently regresses the concreteness of the words on their risk association measures.

Table 2 displays the coefficients obtained from these linear regressions for studies 1A, 1B, and 2. These coefficients specify the relationship between the concreteness of a word and its association with risk for the three studies. Although risk had a highly significant relationship with concreteness, the direction of the relationship varied across the three studies. In study 1A and 2, there was a significant negative coefficient, showing that more abstract words were associated with higher risks. In study 1B, however, we observed a significant positive coefficient, showing that more concrete words were associated with higher risks. This is likely due to the stimuli in these three studies. Study 1A used primarily technological risks. Abstract technologies are typically less known and could be subsequently considered to be riskier. This is also the case for study 2, which had both a large number of technologies and various geopolitical risk factors. In contrast, the stimuli set for study 1B comprises hobbies, sports, and occupations. Physical activities are more concrete than are intellectual activities. Physical activities are also riskier, likely accounting for the positive relationship observed for this study.

Predicting Ratings of Risk Dimensions

Computational Methods

The analysis reported in the previous section used data sets of word ratings on emotional and concreteness variables to examine the conceptual associates of risk. We can also apply a similar technique to test the degree to which the nine key risk dimensions (Fischhoff et al. 1978) were associated with risk judgment, as assessed by the vector space semantic models. Specifically, our three studies elicited actual participant ratings of the risk sources on the nine dimensions, and we can predict these ratings with our uncovered associations. This exercise would provide a strong test of whether conceptual associates of risk obtained through the proposed approach correspond to actual human associations studied in risk perception research.

Now unlike emotions and concreteness there is no existing data set of word ratings for the nine risk dimensions. Thus, in study 3, we used a free association task to elicit sets of words that are closely associated with the nine dimensions, and then used the calculated vector similarity between our risk sources and these words to predict the dimensional ratings of participants in studies 1A, 1B, and 2. As outlined in the experimental methods section, the free association task in study 3 showed participants 18 descriptions, and for each description, participants listed words that first came to their mind. Two descriptions for each risk dimension correspond to high and low ratings on the dimension (specifically, voluntary versus involuntary, immediate versus delayed, known vs

Table 2. Details for Regressions of Concreteness Ratings on Word Risk Associations for the Three Studies

	Coef.	SD	*z*	*p*	95% CI-L	95% CI-H
Study 1A	−1.26	0.03	−43.02	0.00	−1.32	−1.20
Study 1B	0.78	0.03	27.99	0.00	0.72	0.83
Study 2	−1.27	0.03	−42.82	0.00	−1.33	−1.21

Notes. Here, positive coefficients indicate a positive relationship between word concreteness and risk. Coef., coefficient; CI, confidence interval (L = Lower, H = Higher).

unknown to individual exposed to the risk, known versus unknown to science, controllable versus uncontrollable, new versus old, chronic versus catastrophic, risk that individuals can reason about calmly versus one for which they experience dread, and nonfatal versus fatal). With three listed words per participant per description, we obtained a set of 2,700 (nonunique) word associates for the nine risk dimensions. Out of these, 2,646 words were in the Word2Vec vocabulary. For each of these words we knew whether the word was positively associated with the dimension (e.g., listed for the "involuntary" description), negatively associated with the dimension (e.g., listed for the "voluntary" description), or unassociated with the dimension (e.g., listed for some other description). We subsequently calculated the association of the risk sources with each dimension using the relative cosine similarities of the semantic vectors for the risk sources with the free associates. Thus, for a given risk source *i* and a given dimension *j*, we calculated the average of the cosine similarities of the semantic vector for the risk source with the semantic vectors for the words that were positively associated with the dimension. We also calculated the average of the cosine similarities of the semantic vector for the risk source with the semantic vectors for the words that were negatively associated with the dimension. The difference between these two gave us a measure of the relative strength of association of the risk source with the risk dimension. We repeated this for all risk sources *i* and all dimensions *j* to obtain predicted associations between the risk sources and the dimensions.

Results

How well do our predicted associations describe the actual associations of participants? In studies 1A, 1B, and 2, we obtained participant ratings of the risk sources on the nine dimensions. Like in the earlier analysis in this paper, these ratings were averaged to generate a single rating for each risk source on each dimension. We then regressed our participant ratings on the predicted associations for each of the nine dimensions.

Figures 8(a)–8(c) plot the coefficients from each of these regressions for our three studies (and Table 3 provides detailed outputs of these regressions). As can be seen in these figures, our predicted associations were positively related to participant ratings for all dimensions and all studies. Out of these, all except for unknown to science in studies 1A and 1B, and unknown to the individual exposed to the risk, in study 1B, reached statistical significance. Overall, these results indicate that for the vast majority of dimensions in our three studies (and for all dimensions in our primary study, study 2) we are able to quantitatively

Figure 8. Coefficients for Regressions of Participant Ratings on the Predicted Associations for Each of the Nine Risk Dimensions

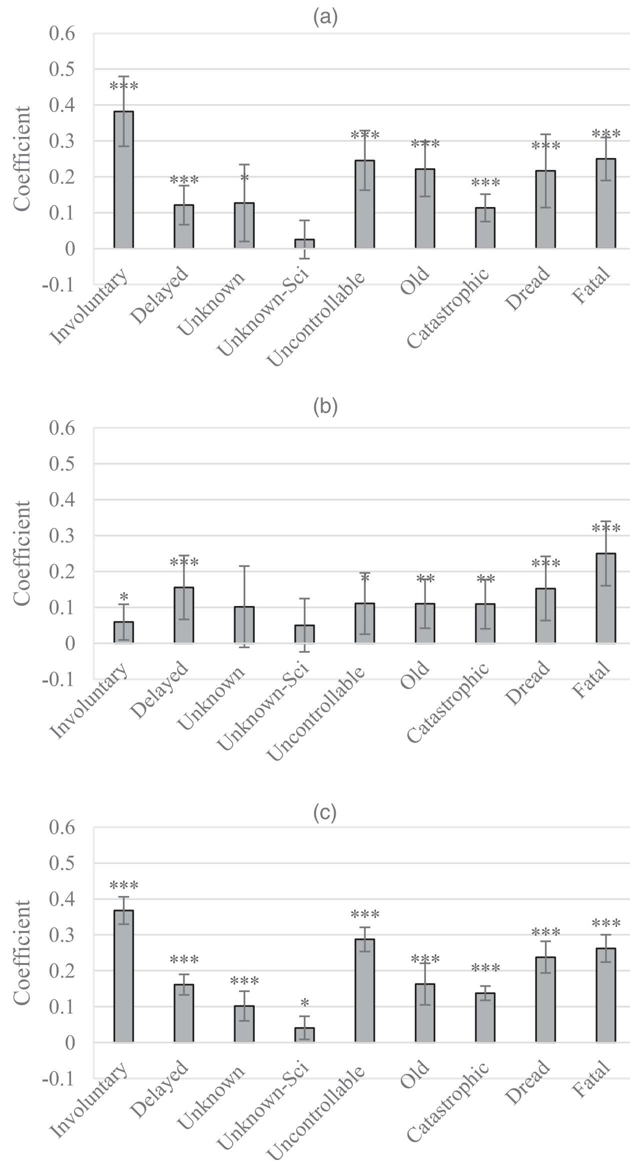

Notes. Error bars capture 95% confidence intervals.
$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

predict the ratings of the risk sources on the dimensions, using our semantic vectors and the free associations of participants with the dimensions.

General Discussion

Predicting Risk Perception

People's knowledge about the world around them determines how they evaluate its objects and events. These evaluations form the basis of decisions made by individuals and societies. To predict (and subsequently understand and influence) people's decisions, we therefore need to uncover and quantify people's knowledge representations.

Table 3. Details for Regressions of Participant Ratings on the Predicted Associations for Each of the Nine Dimensions for the Three Studies

	Coef.	SD	*z*	*p*	95% CI-L	95% CI-H
Study 1A						
Involuntary	0.38	0.05	7.79	0.00	0.29	0.48
Delayed	0.12	0.03	4.42	0.00	0.07	0.18
Unknown	0.13	0.05	2.34	0.02	0.02	0.23
Unknown-sci	0.03	0.03	0.95	0.34	−0.03	0.08
Uncontrollable	0.25	0.04	5.85	0.00	0.16	0.33
Old	0.22	0.04	5.74	0.00	0.15	0.30
Catastrophic	0.11	0.02	5.96	0.00	0.08	0.15
Dread	0.22	0.05	4.20	0.00	0.11	0.32
Fatal	0.25	0.03	7.93	0.00	0.19	0.32
Study 1B						
Involuntary	0.06	0.03	2.37	0.02	0.01	0.11
Delayed	0.16	0.04	3.47	0.00	0.07	0.24
Unknown	0.10	0.06	1.78	0.08	−0.01	0.21
Unknown-sci	0.05	0.04	1.34	0.18	−0.02	0.12
Uncontrollable	0.11	0.04	2.57	0.01	0.03	0.20
Old	0.11	0.03	3.19	0.00	0.04	0.18
Catastrophic	0.11	0.03	3.15	0.00	0.04	0.18
Dread	0.15	0.05	3.38	0.00	0.06	0.24
Fatal	0.25	0.05	5.40	0.00	0.16	0.34
Study 2						
Involuntary	0.35	0.02	18.91	0.00	0.33	0.41
Delayed	0.16	0.01	11.08	0.00	0.13	0.19
Unknown	0.10	0.02	4.89	0.00	0.06	0.14
Unknown-sci	0.04	0.02	2.50	0.01	0.01	0.07
Uncontrollable	0.29	0.02	16.74	0.00	0.25	0.32
Old	0.16	0.03	5.53	0.00	0.10	0.22
Catastrophic	0.14	0.01	13.65	0.00	0.12	0.16
Dread	0.25	0.02	10.66	0.00	0.19	0.28
Fatal	0.26	0.02	13.60	0.00	0.22	0.30

Notes. Here, positive coefficients indicate a positive relationship between predicted associations and participant ratings. Coef., coefficient; CI, confidence interval (L = Lower, H = Higher). Unknown-sci, unknown to science.

This paper has attempted this in the context of risk perception (Fischhoff et al. 1978, Slovic 1987, Fischhoff 1995, Slovic and Weber 2002, Slovic and Peters 2006). It has exploited recent advances in data science involving the automated recovery of knowledge representations for real-world object and events, from natural language data (Dhillon et al. 2011, Griffiths et al. 2007, Jones and Mewhort 2007, Landauer and Dumais 1997, Mikolov et al. 2013, Pennington et al. 2014). Such knowledge representations are often specified as high-dimensional vectors, and these vectors have been shown to predict responses in wide array cognitive, linguistic, and high-level judgment tasks (for reviews, see Bullinaria and Levy 2007 or Jones et al. 2015). Across three studies, involving over 400 experimenter and participant-generated sources of risk, we have found that high-dimensional vector representations of risk sources (specifically, those generated by the Word2Vec methods of Mikolov et al. 2013) also predict the risk levels assigned to those risk sources by participants, both on the aggregate and the individual level. On the aggregate level, our accuracy rates are comparable to those of existing psychometric techniques (Fischhoff et al. 1978, Slovic et al. 1984, Slovic 1987) in studies 1A and 2, which involve a diverse array of technological and geopolitical risk sources. In study 1B, which involves risky activities and hobbies, the relative predictive power of our approach is somewhat lower. On the individual-level, these differences diminish, with both the psychometric and the semantic vector approaches attaining equivalent accuracy rates. Future work should attempt to better understand the relative merits of the proposed semantic vector approach and provide a more rigorous characterization of the settings in which it is likely to exceed or fall short of the accuracy of existing psychometric methods in risk perception research.

Regardless of the relative accuracy rates of the semantic vector and the psychometric approaches, the highest accuracy in all three of our studies, both on the aggregate and on the individual level, is achieved by combining the two approaches. For example, in study 2, our primary study, using both semantic vectors and participant risk ratings allows us to predict over 85% of the variability in our aggregate out-of-sample data. This indicates that vector space knowledge representations complement the types of representations obtained through psychometric tasks, and the two techniques can be used together to best predict lay assessments of risk.

The results of our tests also shed light on the specific machine learning techniques that could be used for predicting risk ratings using semantic vectors as inputs. We have found that the techniques that consistently achieve the highest accuracy rates are support vector regressions with the radial basis function and sigmoidal kernels and ridge regressions. However, the reason for the superiority of these techniques is unclear, and more work is necessary to better understand the linguistic, behavioral, and statistical properties our results. This work could involve an analysis of the relationship between predictive accuracy for each technique and the risk domain, as well as a test of the nonlinearities involved in mapping semantic vector representations onto participant judgments. This analysis would also benefit from a learning curve test, which systematically varies the proportion of the data used to fit the models. Such a test may reveal that different methods excel with different amounts of data, which would help guide future applications of the semantic vector approach.

Associates of Risk

Although the semantic vector approach does not require explicit participant similarity ratings between risk sources, it can nonetheless be seen as an extension of existing techniques like multidimensional scaling. As with multidimensional scaling (Johnson and Tvesrky 1984; see also Kruskal 1964, Borg et al. 2012) our approach attempts to use latent dimensions of concept representation to predict the risk perceptions of individuals. Additionally, these dimensions are uncovered based on similarity measures (applied to word co-occurrence statistics in natural language). Indeed, similarity between the vector representations for the risk sources reflects the actual similarity of the risk sources, as indicated in Figure 1.

One key difference between the semantic vector approach and multidimensional scaling involves the fact that similarity can be computed not only between the vectors for two risk sources but also between the vectors of risk sources and the vectors for other words and concepts. Thus it is possible to uncover the set of words that are most associated with a given risk source, and subsequently uncover the words and concepts with the overall strongest association with risk. We have attempted to do this in all three of our studies. Our analysis finds that the words and concepts with the strongest risk associations are often those that have been previously implicated in risk perception (e.g., "fatal," "dangerous," and "tragic"). Many of these associates are shared across the three studies, reflecting domain general characteristics of risk. Yet other associates are unique to the studies, corresponding to the idiosyncratic features of the risk sources in consideration. Overall, however, the risk association of the words across our three studies is highly correlated, indicating that there are some common underlying associations at play in risk, regardless of the specific risk sources in consideration.

It is useful to note our vectors do not possess explicit information for the outcomes and probabilities (e.g., probability distribution over number of deaths) for the risk sources. They only have association-based representations, which are themselves based on the structure of word co-occurrence in natural language. The fact that such representations are able to provide such a good account of human judgment indicates that lay risk perception is largely associative. Many scholars have already highlighted the associative nature of risk perception (Loewenstein et al. 2001, Slovic et al. 2002, Slovic and Peters 2006). Additionally, in recent work, Bhatia (2017a) has shown how semantic vectors are able to quantitatively predict high-level associative judgments (see also Holtzman et al. 2011; Dehghani et al. 2014; Bhatia 2017a, b; Caliskan et al. 2017; Garten et al. 2017; Bhatia 2018; Bhatia et al. 2018 for applications to social and political judgment and decision-making). The findings of this paper further emphasize the key role of associations in many high-level judgment tasks and demonstrate the desirability of vector space representations for modeling behavior in these tasks.

Uncovering the conceptual associates of risk also allows us to study the relationship between risk perception and various psychological variables of interest. We have attempted to do so with three important sets of variables. The first of these involves six emotions: anger, disgust, happiness, fear, sadness, and surprise. Mohammad and Turney (2013) have compiled measures of the emotionality of a very large list of words, in terms of these six emotions, and it is possible to correlate the emotional rating of a word with the strength of association of that word with risk. Across all three of our studies, we found that the emotions with the strongest relationships with risk are fear and happiness. Specifically, the more closely associated a word is with a risk source, the more likely it is that the word is fear related and the less likely it is that the word is happiness related. We also found a positive relationship between the risk association of words and the anger-relatedness of the words, but this relationship was not as robust. Overall, these results corroborate existing findings on the relationship between risk perception and emotions (Johnson and Tversky 1983; Holtgrave and Weber 1993; Lerner and Keltner 2001; Loewenstein et al. 2001; Lerner et al. 2003; Slovic et al. 2002, 2005; Slovic and Peters 2006).

The second set of psychological variables we considered involved concreteness and abstraction. As with our analysis of emotions, we had concreteness ratings for a large number of words (compiled in a data set by Brysbaert et al. 2014) and were able to correlate the risk association of the word with its concreteness rating. Unlike in our previous analysis, we found different effects for our three studies. In studies 1A and 2, concreteness was negatively associated with risk, so that the words with the highest risk association were the words that were the most abstract. These studies involved technological and geopolitical risks, and it is likely that the abstraction of the risk source correlates with the uncertainty of the risk source, thereby increasing perceptions of risk in this domain. In contrast, in study 1B, concreteness was positively associated with risk, so that the words with the highest risk associations were the words that were the most concrete. This study involved activity-based risk, and concrete activates are also more physical and thus more dangerous than abstract activities, potentially explaining the observed relationships. These results suggest that there are some important domain-level differences in the conceptual associates of risk. They also make novel predictions regarding the effect of concreteness and abstraction on risk, which can be tested in future experimental work.

The final set of variables we examined were the nine key risk dimensions themselves (Fischhoff et al. 1978, Slovic et al. 1984, Slovic 1987). For this purpose, we elicited free associations from participants for each of these dimensions. Using the words generated in this free association task, we were able to categorize each risk source as being positively or negatively associated (or unassociated) with the risk dimension. We compared participant evaluations of the risk sources on these dimensions (obtained in studies 1A, 1B, and 2) with our predicted associations and found a positive relationship for nearly all risk dimensions in all studies.

Overall, these results suggest that vector space representations for risk sources can be used not only to predict the risk assignments of participants but also to understand their conceptual associates and psychological underpinnings. Of course, there are limitations to using this method. For example, it is impossible to disentangle causality. All we can say is that words associated with risk sources are also associated with other psychological variables (such as fear). A more nuanced understanding of the role of these psychological variables in risk perception thus necessitates experimental research. Nonetheless, these results highlight the richness of vector space representations and show novel ways for using these representations to uncover the mental structures involved in the perception of risk for real-world objects.

Language of Risk

A key assumption of the approach proposed in this paper is that the language of risk matters. Prior experimental work has repeatedly shown that perceptions of risk depend on the words used to describe the risk source (see, e.g., Fischhoff 1995 for a review). What is new in this paper, however, is the way in which language influences risk representations and subsequent risk perceptions. Risk sources that co-occur with a certain set of words in everyday language are also likely to have vector representations that are more proximate to those words in the resultant semantic space. As the point a risk source occupies in the semantic space determines its eventual mapping onto risk ratings, changing how a risk source is described in natural language has a direct impact on whether we predict that it will be judged as being risky. For example, if positron emission tomography (PET) is described as "nuclear medicine" in our underlying language data set (as it was in the early days of the technology) then the resultant vector representation for PET would be closer to other applications of nuclear technology (like bombs). If these other applications are rated to be highly risky, then our approach would generalize this to predict that PET also would be considered to be risky.

The relationship between risk perception and language, implicit in our approach, can also contribute to development of better surveys for eliciting risk judgments. The words used in these surveys critically influence participant responses, and the approach proposed in this paper provides a formal technique for modeling this relationship. Indeed, with the help of the approach in the current paper it also may be possible to add diagnostic survey questions to existing psychometric tools, for example, by identifying words that most likely to relate to risk-relevant representations. Examining the feasibility of this application is a useful topic for future work.

Computational Analysis of Risk Perception

Perhaps the greatest benefit of the proposed approach, compared with existing psychometric techniques, is that it does not require additional participant data. The high predictive accuracy rates are achieved without previously elicited participant dimensional ratings or similarity judgments. Rather, they are a product of the computational analysis of large amounts of natural language data. This offers the semantic vector approach a number of unique advantages for modeling and predicting risk perception. Firstly, our approach is able to make out-of-sample predictions: it is possible to train our models on a list of risk sources and use these trained models to estimate risk perceptions for a range of other novel risk sources (without having additional data on these novel sources). Thus, we can identify whether a hypothetical new technology would be perceived as being more or less risky than the one currently in place. Indeed, using the techniques outlined in this paper, it is possible to generate a risk map for thousands of potential sources of risk, including not only traditional risk sources but also novel hazards not traditionally viewed as risks (e.g., social media) as well as for prosaic items that may, from time to time, be considered to be risky (e.g., beef). Such a risk map would be difficult to compile if participants had to explicitly rate each (real or hypothetical) risk source on various risk dimensions. Numerous researchers have argued for the necessity of out-of-sample predictive power in modeling judgment and decision-making (Dawes et al. 1989, Gigerenzer and Brighton 2009; Erev et al. 2017, Plonsky et al. 2017; see also Yarkoni and Westfall 2017 for a general discussion), and this paper can be seen as addressing this issue in the domain of risk perception.

The automated nature of our analysis also allows us to examine the conceptual associates of risk on a much larger scale than is feasible with human data alone. Indeed, in our tests, we have calculated the risk association of tens of thousands of words. Not only does this permit a more comprehensive understanding of the psychological underpinnings of risk but it also allows us to test psychological hypotheses (e.g., hypotheses pertaining to the relationship between fear and risk or between concreteness and risk) in a novel manner. Such tests can be used to both evaluate current theories of risk perception, and to generate novel predictions for subsequent laboratory-based research. Additionally, a comprehensive understanding of the conceptual associates of different sources of risk can inform risk communication strategies and help policy makers identify and remedy misconceptions regarding the riskiness of various technologies, products, or activities.

Our approach is of particular value in settings in which participant data are difficult to obtain. One such setting involves retrospective evaluations of risk. Often it is useful for scholars and policy makers to determine how the perception of a particular risk source has evolved over time, or changed in response to a given event. It is difficult to obtain retrospective evaluation of the riskiness for a given risk source from human participants. However, with appropriate natural language data (e.g., newspaper articles extending many years into the past) such an analysis using our computational techniques is fairly straightforward. Indeed, such an analysis not only provides estimates of the perceived riskiness of a risk source over time but also can be used to calculate changes in the psychological structure of this risk source, including changes in its close associates over time (see, e.g., Iliev et al. 2016, Garg et al. 2018 for examples of such an approach applied to other domains in psychology).

The proposed approach also can be used to predict real-time changes in risk perception and representation. With the use of language data obtained from online news media and social media, it is possible to update learned vector representations (and thus update predicted risk perceptions) based on current events. Although it remains to be seen whether such an application can provide actionable practical insights for policy makers, there is no doubt that a technique for measuring risk perception in real-time holds many potential benefits for researchers (for a similar use of social media in other domains, see Asur and Huberman 2010, Choi and Varian 2012, Curme et al. 2014). We hope to contribute to such an application in the near future.

The power of the proposed approach extends beyond just testing for temporal differences. It also can be used to model changes in risk representation and perception as a function of various social and cultural factors (Peters and Slovic 1996, Finucane et al. 2000, Bickerstaff 2004). Specifically, the models used in our analysis can be trained on representative natural language data sets for different cultural groups, and thus measure cultural differences in the representation of risk sources, and subsequently cultural differences in risk perception. These differences can then be compared against observed findings on the role of culture in the perception of risk, and can similarly be used to derive new testable predictions regarding culture and risk representation (see, e.g., Chen 2013, Noguchi et al. 2014 for related techniques for studying cultural differences in decision-making). More generally, with appropriate natural language data sets, our approach allows for the computational analysis of individual-level risk representations: If we are able to obtain language data from the information sources that a given individual is exposed to, we can build a model that describes that individual's knowledge about the world; a model that can also quantitatively predict that person's perceptions of risk (and corresponding risk associates). Recent computational and societal developments have made such data sets readily available, with important applications for the study of management (George et al. 2016), health behavior (Hawn 2009), consumer choice (Humphreys and Wang 2017), cognition (Griffiths 2015, Jones 2017), and other domains in the behavioral sciences (Harlow and Oswald 2016, Kosinski and Behrend 2017). We look forward to research that exploits these new and exciting data sources, to further enhance our ability to understand and predict risk perception.

Acknowledgments
The author thanks Minkwang Jang for expert research assistance, and Alex Davis and Ellen Peters for valuable feedback.

Appendix
In this appendix, we outline our computational approach with the help of an example. Recall that our primary data involves participant risk ratings for a set of risk sources. For each of these risk sources we also have a 300-dimensional vector representation. A simplified version of this data set, involving five risk sources, each with a hypothetical 10-dimensional representation and a hypothetical participant rating is shown in Table A.1. In Table A.2 we show the pairwise cosine similarities of these risk sources. This table shows risk sources such as (1) cycling and mountain biking and (2) mountain biking and hiking are quite similar (with high cosine similarity values). Hiking and cycling are neither similar nor dissimilar, and basketball is dissimilar to all other risk sources (with negative cosine similarity values). A principal components analysis on this table can be used to recover the first two latent dimensions characterizing the space of risk sources, like in Figure 1.

As part of our primary analysis we applied various machine learning techniques to predict participant risk ratings from our semantic vector representations. These regressions used 9/10th of the data for model training and the remaining 1/10th of the data for calculating model predictive power (formalized in terms of the R^2 of the predictions on the test data) and repeated this 1,000 times, with a random split at each time. With our hypothetical example we can similarly divide our data set into two parts,

Table A.1. Risk Sources with Hypothetical 10-Dimensional Vector Representations and Average Participant Risk Ratings

Risk source	Vector representation	Risk rating
Cycling	[2,–3,0,4,–1,–2,–1,1,–2,–1]	12
Mountain biking	[2, –1,0,4, –1,0, –1,1,0, –1]	44
Rock climbing	[1,2,0,4,–1,1,–1,1,1,–1]	48
Hiking	[0,2,1,3,2,1,1,1,1,0]	4
Basketball	[–1,1,0,–4,0,–1,1,1,–1,1]	–22

Table A.2. Pairwise Cosine Similarities Between Risk Sources in Table A.1

	Cycling	Mountain biking	Rock climbing	Hiking	Basketball
Cycling	1.00	0.84	0.36	0.00	–0.59
Mountain biking	0.84	1.00	0.77	0.34	–0.83
Rock climbing	0.36	0.77	1.00	0.66	–0.72
Hiking	0.00	0.34	0.66	1.00	–0.44
Basketball	–0.59	–0.83	–0.72	–0.44	1.00

such as {cycling, rock climbing, basketball} for model training and {hiking, mountain biking} for model testing. We can use the vector representations and participant risk ratings for the three risk sources in the first part to fit one of our six machine learning techniques and then use mapping learned by these techniques, applied to the vector representations for the remaining two risk sources, to predict participant risk ratings for these two sources. A measure of model predictive power can be obtained by comparing the R^2 of our predictions relative to actual participant ratings for these two risk sources. This can be repeated multiple times with different data splits, and the resultant R^2 values can be averaged to obtain single measures of the predictive power of the six machine learning techniques applied to the vector representations. These measures can be compared against each other and against similar measures for other existing techniques (e.g., those involving linear regressions on participant ratings on nine core risk dimensions).

Now in the main text, we also used our vector representations to examine the conceptual associates of risk. For this purpose, we obtained a large set of words, with each word possessing a 300-dimensional vector representation. Using cosine similarity, we then calculated the association between each word and each risk source. In Table A.3, we present a list of words with hypothetical 10-dimensional vector representations. Like in the paper, we calculated the association between these words and the risk sources using cosine similarity. These associations are shown in Table A.4. As can be seen here, helmet is strongly associated with cycling and mountain biking, weakly associated with rock climbing and hiking, and negatively associated with basketball. Similar associations exist for the other words.

In the main text, we determined the words that were most associated with risk based on the structure of associations between the words and the risk sources. This was done by correlating the associations with the risk sources and the risk ratings of the risk sources for each of the words. When applied to our hypothetical risk sources and words, this involves correlating the associations in Table A.4 with the participant ratings in Table A.1. Thus, for example, for helmet, we correlated {0.97, 0.95, 0.56, 0.16, –0.72} with {12, 44, 48, 4, –22} to obtain a risk association value of 0.78. A similar technique can be applied to the remaining four words. Such risk associations were used to generate Figures 6(a)–6(c) in the main text.

In the paper we also mapped various psychological variables on to risk associations in order to calculate the relationships between these variables and the risk ratings of participants. We can perform a similar analysis with our hypothetical vector representations. Here, in Table A.3, we have hypothetical ratings of the five words in terms of fear (similar to the Mohammad and Turney ratings used in the main text). Crash and fall are rated as being fearful whereas the remaining words are not. To calculate the relationship between fear and risk, we can perform a logistic regression with the fear ratings for the five words as the dependent variable and the risk associations of the five words as the dependent variable. The magnitude of the resultant coefficient would specify the association between fear and risk. Such regressions were used in the main text for five other emotions, as well as for word concreteness/abstraction to generate Tables 1 and 2, and Figures 7(a)–7(c). The main text also used a similar technique, combined with data from a free association task, to predict participant ratings of the risk sources on nine risk dimensions. The results of this analysis are provided in Table 3 and Figures 8(a)–8(c).

Table A.3. Words with Hypothetical 10-Dimensional Vector Representations and Fear Ratings

Word	Vector representation	Fear rating
Helmet	[2,–2,0,4,–1,–1,–1,1,–1,–1]	0
Mountain	[0,1,0,3,2,1,1,1,1,0]	0
Ball	[0,1,0,–4,0,–1,1,1,–1,1]	0
Crash	[2,-1,0,4,–1,–1,–1,1,–1,–1]	1
Fall	[1,2,1,4,–1,1,–1,1,1,0]	1

Table A.4. Cosine Similarities Between Risk Sources in Table A.1 and Words in Table A.3

	Helmet	Mountain	Ball	Crash	Fall
Cycling	0.97	0.11	–0.53	0.93	0.33
Mountain biking	0.95	0.42	–0.77	0.96	0.73
Rock climbing	0.56	0.64	–0.70	0.67	0.96
Hiking	0.16	0.95	–0.45	0.25	0.70
Basketball	–0.72	–0.54	0.98	–0.72	–0.68

References

Asur S, Huberman BA (2010) Predicting the future with social media. *Proc. 2010 IEEE/WIC/ACM Internat. Conf. Web Intelligence Intelligent Agent Tech.*, vol. 1 (IEEE, Piscataway, NJ), 492–499.

Bhatia S (2017a) Associative judgment and vector space semantics. *Psych. Rev.* 124(1):1–20.

Bhatia S (2017b) The semantic representation of prejudice and stereotypes. *Cognition* 164:46–60

Bhatia S (2018) Semantic processes in preferential choice. *J. Experiment. Psych. Learn.* Forthcoming.

Bhatia S, Goodwin G, Walasek L (2018) Trait associations for Hillary Clinton and Donald Trump in news media: A computational analysis. *Soc. Psych. Personality Sci.* 9(2):123–130.

Bickerstaff K (2004) Risk perception research: Socio-cultural perspectives on the public experience of air pollution. *Environ. Internat.* 30(6):827–840.

Borg I, Groenen PJ, Mair P (2012) *Applied Multidimensional Scaling* (Springer, New York).

Brysbaert M, Warriner AB, Kuperman V (2014) Concreteness ratings for 40 thousand generally known English word lemmas. *Behav. Res. Methods* 46(3):904–911.

Bullinaria JA, Levy JP (2007) Extracting semantic representations from word co-occurrence statistics: A computational study. *Behav. Res. Methods* 39:510–526.

Caliskan A, Bryson JJ, Narayanan A (2017) Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334):183–186.

Chandran S, Menon G (2004) When a day means more than a year: Effects of temporal framing on judgments of health risk. *J. Consumer Res.* 31(2):375–389.

Chen MK (2013) The effect of language on economic behavior: Evidence from savings rates, health behaviors, and retirement assets. *Amer. Econom. Rev.* 103(2):690–731.

Choi H, Varian H (2012) Predicting the present with Google Trends. *Econom. Record* 88(1):2–9.

Curme C, Preis T, Stanley HE, Moat HS (2014) Quantifying the semantics of search behavior before stock market moves. *Proc. Natl. Acad. Sci. USA* 111(32):11600–11605.

Dawes RM, Faust D, Meehl PE (1989) Clinical vs. actuarial judgment. *Science* 243(4899):1668–1674.

Dehghani M, Sagae K, Sachdeva S, Gratch J (2014) Analyzing political rhetoric in conservative and liberal weblogs related to the construction of the "Ground Zero Mosque." *J. Inform. Tech. Polit.* 11(1):1–14.

Dhillon P, Foster DP, Ungar LH (2011) Multi-view learning of word embeddings via CCA. Shawe-Taylor J, Zemel RS, Bartlett PL, Pereira F, Weinberger KQ, eds. *Advances in Neural Information Processing Systems 24* (Curran Associates, Hook, NY), 199–207.

Erev I, Ert E, Plonsky O, Cohen D, Cohen O (2017) From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psych. Rev.* 124(4): 369–409.

Finucane ML, Slovic P, Mertz CK, Flynn J, Satterfield TA (2000) Gender, race, and perceived risk: The 'white male'effect. *Health Risk Soc.* 2(2):159–172.

Firth JR (1957) *Papers in Linguistics* (Oxford University Press, Oxford, UK).

Fischhoff B (1995) Risk perception and communication unplugged: twenty years of process. *Risk Anal.* 15(2):137–145.

Fischhoff B, Kadvany J (2011) *Risk: A Very Short Introduction* (Oxford UP, Oxford, UK).

Fischhoff B, Slovic P, Lichtenstein S, Read S, Combs B (1978) How safe is safe enough? A psychometric study of attitudes toward technological risks and benefits. *Policy Sci.* 9(2):127–152.

Garg N, Schiebinger L, Jurafsky D, Zou J (2018) Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proc. Natl. Acad. Sci. USA* 115(16):E3635–E3644.

Garten J, Hoover J, Johnson K, Boghrati R, Iskiwitch C, Dehghani M (2017) Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behav. Res. Methods* 50(1):344–361

George G, Osinga EC, Lavie D, Scott BA (2016) Big data and data science methods for management research. *Acad. Management J.* 59(5):1493–1507.

Gigerenzer G, Brighton H (2009) Homo heuristicus: Why biased minds make better inferences. *Topics Cognitive Sci.* 1(1):107–143.

Griffiths TL (2015) Manifesto for a new (computational) cognitive revolution. *Cognition* 135:21–23.

Griffiths TL, Steyvers M, Tenenbaum JB (2007) Topics in semantic representation. *Psych. Rev.* 114(2):211.

Harlow LL, Oswald FL (2016) Big data in psychology: Introduction to the special issue. *Psych. Methods* 21(4):447–457.

Harris ZS (1954) Distributional structure. *Word* 2:146–162.

Hawn C (2009) Take two aspirin and tweet me in the morning: How Twitter, Facebook, and other social media are reshaping healthcare. *Health Affairs* 28(2):361–368.

Holtzman NS, Schott JP, Jones MN, Balota DA, Yarkoni T (2011) Exploring media bias with semantic analysis tools: Validation of the contrast analysis of semantic similarity (CASS). *Behav. Res. Methods* 43(1):193–200.

Holtgrave DR, Weber EU (1993) Dimensions of risk perception for financial and health risks. *Risk Anal.* 13(5):553–558.

Humphreys A, Wang RJ-H (2017) Automated text analysis for consumer research. *J. Consumer Res.* 6(1):1274–1306.

Iliev R, Hoover J, Dehghani M, Axelrod R (2016) Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proc. Natl. Acad. Sci. USA* 113(49): E7871–E7879.

Jenkin CM (2006) Risk perception and terrorism: Applying the psychometric paradigm. *Homeland Security Affairs J.* 2:1–14.

Johnson EJ, Tversky A (1983) Affect, generalization, and the perception of risk. *J. Personality Soc. Psych.* 45(1):20.

Johnson EJ, Tversky A (1984) Representations of perceptions of risks. *J. Experiment. Psych. General* 113(1):55.

Jones MN, Mewhort DJ (2007) Representing word meaning and order information in a composite holographic lexicon. *Psych. Rev.* 114(1):1.

Jones MN, Willits JA, Dennis S (2015) Models of semantic memory. Busemeyer JR, Townsend JT, eds. *Oxford Handbook of Mathematical and Computational Psychology* (Oxford University Press, Oxford, UK), 232–254.

Jones MN, ed. (2017) *Big Data in Cognitive Science, Frontiers of Cognitive Psychology* (Psychology Press, Taylor & Francis Group, Abingdon, UK).

Krimsky S, Golding D (1992) *Social Theories of Risk* (Praeger, Santa Barbara, CA).

Kruskal JB (1964) Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* 29(1):1–27.

Kosinski M, Behrend T (2017) Big data in the behavioral sciences. *Curr. Opin. Behav. Sci.* 18:4–6.

Landauer TK, Dumais ST (1997) A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psych. Rev.* 104(2):211–240.

Lerner JS, Keltner D (2001) Fear, anger, and risk. *J. Personality Soc. Psych.* 81(1):146.

Lerner JS, Gonzalez RM, Small DA, Fischhoff B (2003) Effects of fear and anger on perceived risks of terrorism a national field experiment. *Psych. Sci.* 14(2):144–150.

Loewenstein GF, Weber EU, Hsee CK, Welch N (2001) Risk as feelings. *Psych. Bull.* 127(2):267.

McDowell M, Rebitschek FG, Gigerenzer G, Wegwarth O (2016) A simple tool for communicating the benefits and harms of health interventions: A guide for creating a fact box. *Medical Decision Making Policy Practice* 1(1):1–10.

Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J (2013) Distributed representations of words and phrases and their compositionality. Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, eds. *Adv. Neural Inform. Processing Systems (NIPS 26)*, 3111–3119.

Mohammad SM, Turney PD (2013) Crowdsourcing a word–emotion association lexicon. *Comput. Intell.* 29(3):436–465.

Noguchi T, Stewart N, Olivola CY, Moat HS, Preis T (2014) Characterizing the time-perspective of nations with search engine query data. *PLoS One* 9(4):e95209.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Prettenhofer P, Weiss R, Dubourg V, et al. (2011) Scikit-learn: Machine learning in Python. *J. Machine Learn. Res.* 12(Oct): 2825–2830.

Pennebaker JW, Francis ME, Booth RJ (2001) *Linguistic Inquiry and Word Count (LIWC 2001)* (Lawrence Erlbaum Associates, Mahwah, NJ).

Pennington J, Socher R, Manning CD (2014) Glove: Global vectors for word representation. *Proc. 2014 Conf. Empirical Methods Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Stroudsburg, PA), 1532–1543.

Peters E, Slovic P (1996) The role of affect and worldviews as orienting dispositions in the perception and acceptance of nuclear power. *J. Appl. Soc. Psych.* 26(16):1427–1453.

Pidgeon N, Hood C, Jones D, Turner B, Gibson R (1992) Risk perception. Risk: Analysis, Perception and Management: Report of a Royal Society Study Group, The Royal Society, London, 89–134.

Plonsky O, Erev I, Hazan T, Tennenholtz M (2017) Psychological forest: Predicting human behavior. *Proc. 31st AAAO Conf. Artificial Intelligence (AAAI-17)* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA), 656–662.

Sanquist TF, Mahy H, Morris F (2008) An exploratory risk perception study of attitudes toward homeland security systems. *Risk Anal.* 28(4):1125–1133.

Slovic P (1987) Perception of risk. *Science* 236(4799):280.

Slovic P (1992) Perception of risk: Reflections on the psychometric paradigm. Krimsky S, Golding D, eds. *Social Theories of Risk* (Praeger, New York), 117–152.

Slovic P, Peters E (2006) Risk perception and affect. *Curr. Dir. Psych. Sci.* 15(6):322–325.

Slovic P, Weber EU (2002) Perception of risk posed by extreme events. Working paper, University of Oregon, Eugene.

Slovic P, Fischhoff B, Lichtenstein S (1984) Behavioral decision theory perspectives on risk and safety. *Acta Psych.* 56(1):183–203.

Slovic P, Finucane M, Peters E, MacGregor DG (2002) Rational actors or rational fools: Implications of the affect heuristic for behavioral economics. *J. Sociol. Econom.* 31(4):329–342.

Slovic P, Peters E, Finucane ML, MacGregor DG (2005) Affect, risk, and decision making. *Health Psych.* 24(4S):S35.

Trope Y, Liberman N (2010) Construal-level theory of psychological distance. *Psych. Rev.* 117(2):440.

Turney PD, Pantel P (2010) From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res.* 37(1):141–188.

Yarkoni T, Westfall J (2017) Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspect. Psych. Sci.* 12(6):1100–1122.