informs DECISION ANALYSIS SOCIETY

DECISION ANALYSIS TODAY

Vol. 38, No. 1, July 2019

The newsletter of the INFORMS Decision Analysis Society

Inside:

President's Letter

Karen Jenni

Fellow DAS members,

Summer greetings! In this "President's Letter" I want to preview some coming attractions, invite your input on how to strengthen to core of DAS (or, really, on any aspect of our society that captures your imagination and energy), and close with some kudos and thanks.

By the time you receive this the **Advances in Decision Analysis Conference** will be completed. This is the third Advances conference DAS has organized—this year we were hosted by Bocconi University in Milan, and organized by a fabulous committee led by Emanuele Borgonovo (our own VP). We had nearly 150 attendees and an exciting program, and I promise a full report through INFORMS Connect and in future newsletters.

Coming attractions include:

- INFORMS keeps an up-to-date calendar of various INFORMS-sponsored conference on their website. One that I think might be of more interest to DAS members is INFORMS Healthcare 2019 - Transforming Health with Data, Mind and Hand, July 27-29, 2019, in Cambridge, MA.
- And of course, October 20-23rd will see the **INFORMS Annual Meeting** in Seattle, WA. Christian Wernz and Allison Reilly are this year's DAS cluster chairs. Thanks to them and to all of you for your enthusiasm for the conference. The DA tracks are shaping up nicely, and we are excited to see how our "flash talk" session initiative works out.

This newsletter shows additional upcoming conferences on page 4.

Request for input and invitation to engage. At the annual meeting in Phoenix, I outlined three areas of focus for the Society: strengthening our core, building connections, and raising awareness. In my series of letters I want to expand on these one at a time, and request your input. First up is strengthening our core.

The annual conference highlights the things that we do as a society and our continuing need to support and refresh these activities: our annual awards, this newsletter, the annual INFORMS conference, the Advances conference, membership and mentorship opportunities, joint webinars with SDP, and more. As a result of a member survey last year, we re-chartered a special subcommittee to focus on increasing the vibrancy of the Society and creating opportunities for all of our members to engage and get more benefit from their membership. Matthias Siefert, Jay Simon, and Mehmet Ayvaci will continue their efforts aimed at student and new members, while Eva Regnier, Lea Delaris, and Sinan Tas will spearhead efforts aimed at the full membership. Please reach out to me or to any of these folks if you have ideas, needs, or want to help.

Delayed thanks and kudos. My thanks to Jason Merrick for his enthusiastic stewardship of our Society for the past several years. We passed the leadership baton during the annual conference in Phoenix, so thanks and kudos to "outgoing" Council members Saurabh Bansal and Emanuele Borgonovo, to all continuing Council members, and all our many amazing volunteers. DAS would not be DAS without all of your hard work! Special thanks to Saurabh and to Christian Wernz for organizing and managing the DAS cluster for an interesting and rewarding conference this past fall. Kudos (and advance thanks) to our "incoming" officers: Council members Andrea Hupman and Valentina Ferretti, and VP/President-elect Emanuele Borgonovo (continuing his DAS service in a new role). Thanks to all of our candidates for their readiness to serve and their continued support of the Society. We are fortunate to have a rich and fascinating body of work to examine when considering our annual DAS awards. Thanks to all of you for continuing to make the job of our selection committees challenging, and kudos to all the 2018 winners.

And finally, thanks for your time, and please enjoy this edition of ***Decision Analysis Today.***

Karen Jenni
U.S. Geological Survey
kjenni@usgs.gov
And: karenjenni@gmail.com
DAS President, 2018-2020

Letter from the Editors

Debarun Bhattacharjya and Cameron MacKenzie

Dear reader,

We hope your summer is going well, and we hope that you are staying cool with the various heat waves going around the world. We thank everyone involved in organizing and planning Advances in Decision Analysis (ADA) in Milan for such an excellent conference! Also, we welcome the new leadership team—Karen Jenni and Emanuele Borgonovo—as well as newly elected Council members—Andrea Hupman and Valentina Ferretti.

We are excited to include numerous interesting columns in this newsletter edition. Pat Leach in his DA Practice column questions the extent to we should let artificial intelligence make decisions in place of humans. As artificial intelligence becomes more complex and less transparent, how should we allow it to make decisions around the most important issues in our lives? These sorts of concerns are primarily what is resulting in a lot of interest in "explainability" and "interpretability" in artificial intelligence and machine learning communities.

Seth Baum, Executive Director of the Global Catastrophic Risk Institute, authors a unique Research column examining how to think about global catastrophic risks and how this type of analysis can influence good risk management decisions. We thank column editor Roshanak Nateghi and Seth for bringing this column to the newsletter.

In DA Around the World, Mavis Chen and Shijith Kumar use text mining and machine learning techniques to analyze keywords in every *Decision Analysis* journal article since 2004. They analyze these keywords according to the authors' countries of origin, making several interesting observations about how topics in the decision analysis literature vary in time and across geography. We thank Mavis and Shijith for spending all the time it must have taken to conduct this extensive analysis.

We also encourage you to look at the Articles in Advance for the *Decision Analysis* journal. Plus, don't miss the book announcement for *The Art and Science of Making Up Your Mind* by Rex Brown, who passed away in 2017; this book will likely be of interest to most readers.

We hope you enjoy this issue. As always, thank you to all of the contributors. This newsletter could not be completed without their efforts. We are interested in learning about how DAS members use the newsletter and will reach out to the community to learn how to improve it—stay tuned for updates.

Happy reading,

Cameron and Debarun

Upcoming Conferences

July 22-25, 2019
Conference on Uncertainty in Artificial Intelligence (UAI)
Tel Aviv, Israel
http://www.auai.org/uai2019

July 27-29, 2019
INFORMS Healthcare
Cambridge, Massachusetts
http://meetings2.informs.org/wordpress/healthcare2019/

October 20-23, 2019
INFORMS Annual Meeting
Seattle, Washington
http://meetings2.informs.org/wordpress/seattle2019/

November 23-25, 2019
Decision Sciences Institute Annual Meeting
New Orleans, Lousiana
https://decisionsciences.org/annual-meetings/national-dsi/

December 8-12, 2019
Society for Risk Analysis Annual Meeting
Arlington, Virginia
https://www.sra.org/events/2019-sra-annual-meeting

December 8-11, 2019
Winter Simulation Conference
National Harbor, Maryland
http://meetings2.informs.org/wordpress/wsc2019

March 31 – April 3, 2020
SDP Annual DAAG Conference
Houston, Texas
https://www.decisionprofessionals.com/

April 26-28, 2020
INFORMS Business Analytics Conference
Denver, Colorado
http://meetings2.informs.org/wordpress/analytics2020/

May 30 – June 2, 2020
Institute of Industrial and Systems Engineers Annual Conference & Expo
New Orleans, Louisiana
https://www.iise.org/Annual/

Book Announcement

The Art and Science of Making Up Your Mind: 1st Edition (Paperback) - Routledge

By Rex V. Brown

Description

The Art and Science of Making Up Your Mind presents basic decision-making principles and tools to help the reader respond efficiently and wisely to everyday dilemmas.

Although most decisions are made informally (whether intuitively without deliberate thought or based on careful reflection), over the centuries people have tried to develop systematic, scientific, and structured ways in which to make decisions. Using qualitative counterparts to quantitative models, Rex Brown takes the reader through the basics, like “what is a decision” and then considers a wide variety of real-life decisions, explaining how the best judgments can be made using logical principles.

Combining multiple evaluations of the same judgment (“hybrid judgment”) and exploring innovative analytical concepts (such as “ideal judgment”), this book explores and analyzes the skills needed to master the basics of non-mathematical decision making, and what should be done, using real world illustrations of decision methods.

The book is an ideal companion for students of Thinking, Reasoning and Decision-Making, and also for anyone wanting to understand how to make better judgments in their everyday lives.

Reviews

- “This highly readable book aims to teach the reader to obtain a superior second opinion—from herself. A wealth of often provocative examples reveals the wisdom of a master of applied decision theory.” - **Daniel Kahneman, Nobel Prize Winner, Author of "Thinking, Fast and Slow," Eugene Higgins Professor of Psychology Emeritus at Princeton University**

- “A trailblazing pioneer in decision education, Rex Brown has provided us with invaluable examples, tools, evidence and arguments for everyone to invest in their decision skills.” - **Chris Spetzler, Executive Director, Decision Education Foundation**

- “Rex Brown spent all his working life thinking about how to make decisions well. This book is the culmination of his thought. He concentrates on the essential ideas, illustrated with many practical examples, and avoiding most of the mathematics that surrounded the subject when originally formed. This excellent text is to be recommended to all who want a readable and straightforward introduction to the analysis of any decision.” - **Professor Stephen Watson, Life Fellow, Emmanuel College, Cambridge University**

- “Did you ever think hard about a decision and later feel like it had been a bad mistake? After decades advising others and trying to avoid making more mistakes myself, I urge you to read this book now. No psychology or statistics, just lots of pithy, how-to-do-it, real examples. You'll make better decisions easily, and you'll have fewer regrets.” - **Andrew Kahr, Founder: First Deposit Corporation, formerly Assoc. Prof. of Business Administration, Harvard Business School**

- "Rex Brown's *The Art and Science of Making Up Your Mind* is a non-technical textbook and guidebook for how to understand and use basic principles and tools of applied decision theory to deal effectively with everyday decisions and difficult dilemmas. Rex Brown is a wise mentor and reliable companion. He guides the reader through basic questions such as 'what is a decision' and 'what is an ideal judgement.' He then draws on personal, family and friends' decisions in professional life, health, illness, and voting to elucidate how quantitative decision aids and qualitative considerations help clarify problems and lead to sound decisions. An indispensable resource for students and all decision makers." - **Professor Leon Mann AO, Professorial Fellow, Melbourne School of Psychological Sciences, University of Melbourne**

- "This gem of a book synthesizes the life's work of one of decision science's seminal thinkers. Rex Brown's passion was making the tools and thought processes of applied decision theory accessible to ordinary people. This clear and entertaining how-to manual uses examples from his own life and the lives of his family to help readers grasp how we can realize our potential for better everyday decision-making and, ultimately, greater satisfaction in our lives."- **Kathryn Blackmond Laskey, Professor of Systems Engineering and Operations Research, George Mason University**

Table of Contents

About the Author

Rex V. Brown's 50-year career combined research, consulting and teaching on helping people and organizations to make better decisions. He wrote five previous books on decision science and more than 80 papers and articles. Brown was a Distinguished Senior Fellow in the School of Public Policy at George Mason University. He held faculty appointments at Harvard Business School, University College London, London School of Economics and the Universities of Michigan, Carnegie-Mellon, George Mason and Cambridge. He spent 20 years as Chairman of Decision Science Consortium, Inc. in Reston, Virginia, and was a founding Council Member of the Decision Analysis Society.

2018 INFORMS *DAS Awards*

Frank P. Ramsey Medal Award

Max Henrion (Lumina Decision Systems) has been named as the 2018 Ramsey Medalist.

Picture: Ramsey Medal Award winner: Max Henrion (R) with committee chair Eric Bickel (L)

DAS Practice Award

The winners of the 2018 DAS / SDP Practice Award are Corinne Carland, Jarrod Goentzel, and Gilberto Montibeller for their work titled: "Modeling the Value of Agents in Supply Chains of Malaria Rapid Diagnostic Test Kits with Decision Analysis."

Picture: Practice Award winners: Gilberto Montibeller (L) and Jarod Goentzel (C) with committee co-chair Mike Runge (R)

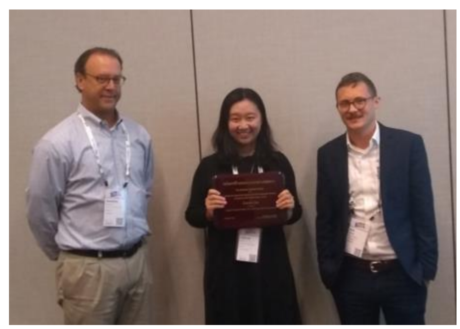

DAS Student Paper Award

The winner of the 2018 student paper award is Xiaojia Guo for her paper, "Quantile Forecasts of Product Lifecycle Cycles Using Exponential Smoothing."

Picture: Student Paper Award winner: Xiaojia Guo (C) with committee co-chairs Joe Hahn (L) and Asa Palley (R)

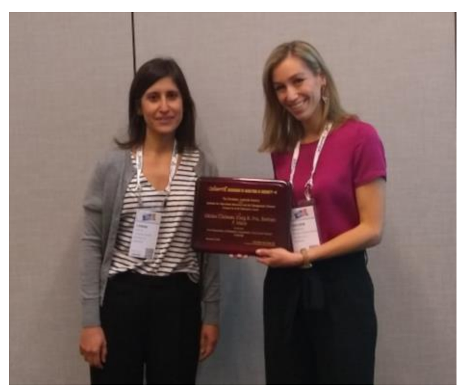

DAS Publication Award

The winners of the 2018 publication paper award are Gülden Üklümen, Craig R. Fox, and Betram F. Malle for their work entitled, "Two Dimensions of Subjective Uncertainty: Clues from Natural Language."

Picture: Publication Award winner: Gülden Üklümen (R) with committee member Canan Ulu (L)

Decision Analysis June 2019 Issue

http://pubsonline.informs.org/toc/deca/16/2

Gulf Coast Port Selection Using Multiple-Objective Decision Analysis

Rivelino R. De Icaza, Gregory S. Parnell, and Edward A. Pohl

In today's competitive global markets, ports play a vital role in global supply chain operations. A port selection model was developed to demonstrate the potential to aid shipping line companies with decisions to select the best ports in the U.S. Gulf Coast. Because port selection is complex, is dynamic, and includes multiple objectives, a multiple-objective decision analysis quantified the value of the various ports using industry data and an industry expert's knowledge. In addition, a corresponding cost model was developed using available cost data for each port. Monte Carlo simulation analyzed the uncertainties in the value and cost models. For the demonstration model scenario assumptions, the results show that the Houston port would be the best alternative for shipping lines in the Gulf Coast. The value model also identified opportunities for improvement for each port compared with the best West Coast port. The demonstration models show that port data exist to support the development of a multiple-objective decision analysis model and a cost model that can provide useful insights to port selection decision makers. Furthermore, these models can be easily tailored to support port selection in other regions by tailoring the objectives and measures to the decision-maker objectives.

https://doi.org/10.1287/deca.2018.0381

Principal–Agent Theory, Game Theory, and the Precautionary Principle

Kjell Hausken

Principal–agent theory and game theory are applied to the precautionary principle (PP) to open up a new research agenda. Principals assess whether the threat is uncertain above a threshold. If it is, the principals choose, pay, and command agents to decrease the uncertainty below the threshold. The agents perform the action. The process is repeated through a feedback loop impacting the threat, after which the process is renewed. The four dimensions of the PP, that is, threat, uncertainty, command, and action, are described. Games and game characteristics in the four dimensions are recognized. Games are possible between natural, technological, and human factors causing the threat and between principals, agents, and external actors. Moral hazard and adverse selection in principal–agent theory related to the PP are considered. Twelve kinds of uncertainty are identified for principal–agent theory in the PP, that is, the natures of the threat, uncertainty, and threshold; states of nature, technology, knowledge, and information; whether a game is played; players; which game is played; strategy sets; utilities; beliefs; incomplete information; imperfect information; risk attitudes; and bounded rationality.

https://doi.org/10.1287/deca.2018.0380

The Myopic Property in Decision Models

Manel Baucells and Rakesh K. Sarin

We examine conditions under which decisions made in isolation provide an optimal strategy for the multiperiod problem. We focus on investment decisions with constant returns to scale. We first consider the framework of subjective expected utility. Under minimal assumptions (i.e., without assuming utility is concave), we prove that *only* log utility is myopically optimal when returns are serially correlated. When returns are serially independent, we generalize Mossin's result [Mossin J (1968) Optimal multiperiod portfolio policies. *J. Bus.* 41(2):215–229.] that only log and power, including linear and convex, possess the myopic property. Finally, we extend the inquiry when probabilities are uncertain and the decision maker uses the recursive smooth model of ambiguity to identify an optimal strategy. We show that with serial correlation, preferences including ambiguity concerns cannot be myopic and optimal. Without correlation, we identify the exact pairs of utility and ambiguity functions that permit myopic decision rules.

https://doi.org/10.1287/deca.2018.0384

On the Geometry of Nash and Correlated Equilibria with Cumulative Prospect Theoretic Preferences

Soham R. Phade and Venkat Anantharam

It is known that the set of all correlated equilibria of an nn-player non-cooperative game is a convex polytope and includes all of the Nash equilibria. Furthermore, the Nash equilibria all lie on the boundary of this polytope. We study the geometry of both these equilibrium notions when the players have cumulative prospect theoretic (CPT) preferences. The set of CPT correlated equilibria includes all of the CPT Nash equilibria, but it need not be a convex polytope. We show that it can, in fact, be disconnected. However, all of the CPT Nash equilibria continue to lie on its boundary. We also characterize the sets of CPT correlated equilibria and CPT Nash equilibria for all 2×2 games, with the sets of correlated and Nash equilibria in the classical sense being a special case.

https://doi.org/10.1287/deca.2018.0378

Decision Analysis Special Recognition Award, 2018

The Decision Analysis Special Recognition Award is given annually for the paper most worthy of recognition published in *Decision Analysis*. For more information: http://pubsonline.informs.org/page/deca/decision-analysis-special-recognition-award

2018 Winner

Bin Mai, Shailesh Kulkarni, **When Hackers Err: The Impacts of False Positives on Information Security Games**

Volume 15, Issue 2, June 2018

https://doi.org/10.1287/deca.2017.0363

2018 Finalist

Robin L. Dillon, William J. Burns, Richard S. John**; Insights for Critical Alarm-Based Warning Systems from a Risk Analysis of Commercial Aviation Passenger Screening**

Volume 15, Issue 3, September 2018

https://doi.org/10.1287/deca.2018.0369

2018 Award Selection Committee: Victor Richmond Jose (chair), Eva Regnier, and Matthias Seifert

Attention INFORMS Decision Analysis Society Members!

By special arrangement with the Decision Analysis Society Council,
dues-paying regular members of the DAS receive a

subscription to the journal as part of their membership dues.

The DAS is a subdivision of INFORMS.
For information on DAS: https://www.informs.org/Community/DAS .

Decision Analysis is a quarterly journal dedicated to advancing the theory, application, and teaching of all aspects of decision analysis. The primary focus of the journal is to develop and study operational decision-making methods, drawing on all aspects of decision theory and decision analysis, with the ultimate objective of providing practical guidance for decision makers. As such, the journal aims to bridge the theory and practice of decision analysis, facilitating communication and the exchange of knowledge among decision analysts in academia, business, industry, and government. *Decision Analysis* is published in March, June, September, and December by the Institute for Operations Research and the Management Sciences (INFORMS) at 5521 Research Park Drive, Suite 200, Catonsville, Maryland 21228. Please visit our website at http://pubsonline.informs.org/journal/deca.

DA Around the World

Column Editors: Chen (Mavis) Wang and Shijith Kumar

In this column we introduce Decision Analysis communities around the world with the purpose of promoting their visibility and strengthening the ties between DA researchers and practitioners across borders. In the current issue, we apply two state-of-the-art text mining methods, *word embedding* and *topic modeling*, to the papers published in *Decision Analysis* since its inception (from 2004 to 2018), and analyze the research interests of authors from different countries.

To start with, we downloaded 284 research articles in total (including introductions to two special issues) from the journal website https://pubsonline.informs.org/journal/deca. Countries of each author's first affiliation are summarized as follows. Researchers from the United States authored or co-authored 218 of these papers. Besides the United States, the top nine countries are: Germany (13), Canada (12), the United Kingdom (11), Spain (8), France (8), Finland (7), Singapore (5), and China (5). Furthermore, there are authors from 23 other countries/regions, including the Netherlands (4), Israel (4), Australia (4), Sweden (3), Norway (3), Italy (3), Switzerland (3), New Zealand (3), Brazil (2), Thailand (2), Portugal (2), Belgium (2), Denmark (2), Ireland (2), South Africa (1), Austria (1), Mexico (1), Azerbaijan (1), India (1), Turkey (1), Hong Kong (1), Ecuador (1), and Chile (1), contributing to the high geographic diversity of *Decision Analysis* authorship. In the current issue, we aim to provide a brief summary of the research interests of authors from several top countries.

Figure 1 displays the top 30 high-frequency author-supplied keywords in the 284 research articles. For the convenience of presentation, we merged keywords with nearly identical meanings (e.g., by replacing *multi-objective* with *multiattribute* whenever appropriate). The selected words and phrases cover a variety of methodological topics such as utility theory, Bayesian statistics, decision tree, game theory, value of information, forecast aggregation, and influence diagram, and application areas such as auction and counterterrorism. We can also observe the transition of research interests over time. For example, decision tree, influence diagram, and Bayesian network were used most often during 2004-2008, while game theory, auction, terrorism, and group decision making drew more attention after 2009. However, when analyzing each keyword separately, we lose sight of the interconnections between them. Instead of relying on experience for categorization, we apply the recent machine learning method of *word embedding* to capture their interdependencies and cluster them into meaningful topics.

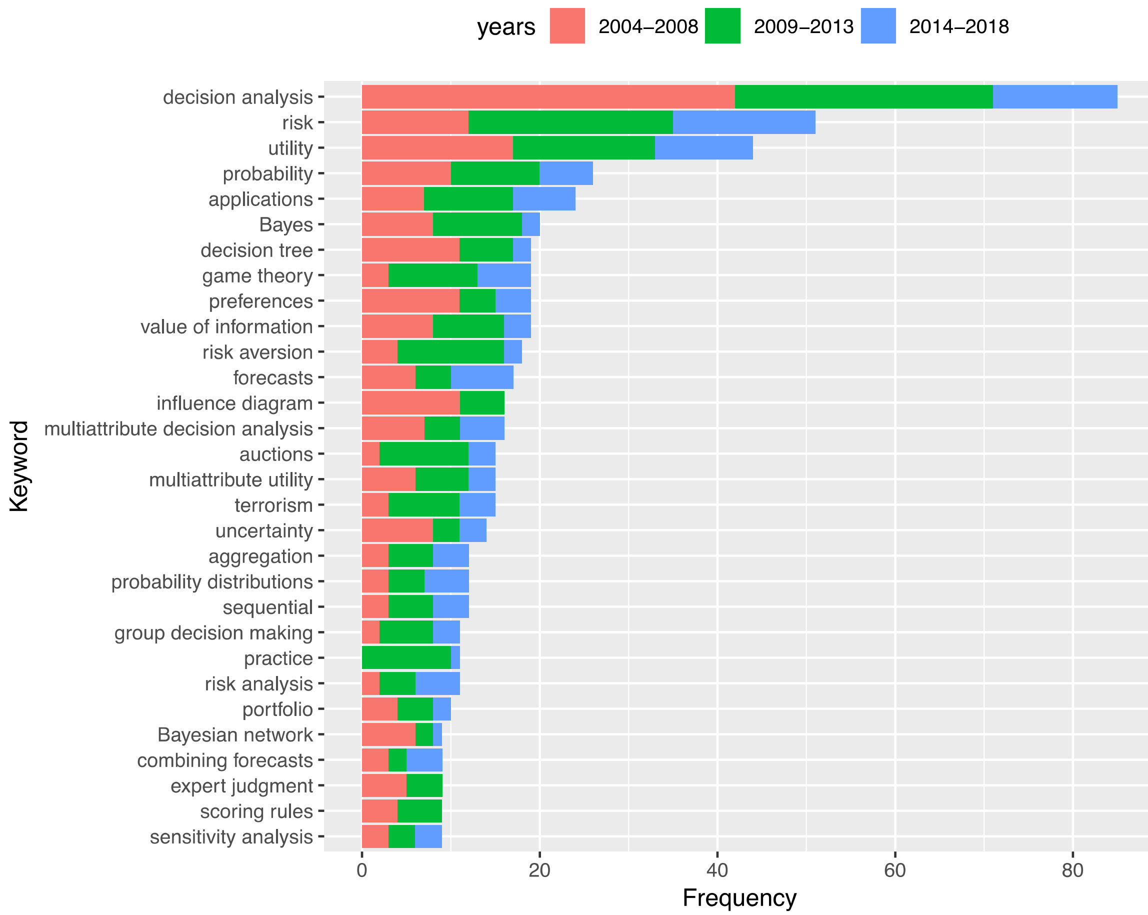

Figure 1 Top 30 High-frequency Key Words

Word embedding is a popular method for natural language processing in which words are represented by multi-dimensional vectors in a latent space (e.g., Grover and Leskovec 2016). The conditional probability of the appearance of a word given the context composed of other words is specified by a neural network. Proximity in the latent space indicates similar semantic functions in terms of predicting the occurrence of other words. We implemented the node2vec algorithm of *word embedding* (using the python code provided by https://github.com/aditya-grover/node2vec) on all of the author-supplied keywords. We set the dimension of the latent space to be four, and employed principal component analysis for further dimensionality reduction. Figure 2 presents the first two principal components of those key words that show up more than three times in the 284 research articles. The first two principal components account for 72% of the variability in data.

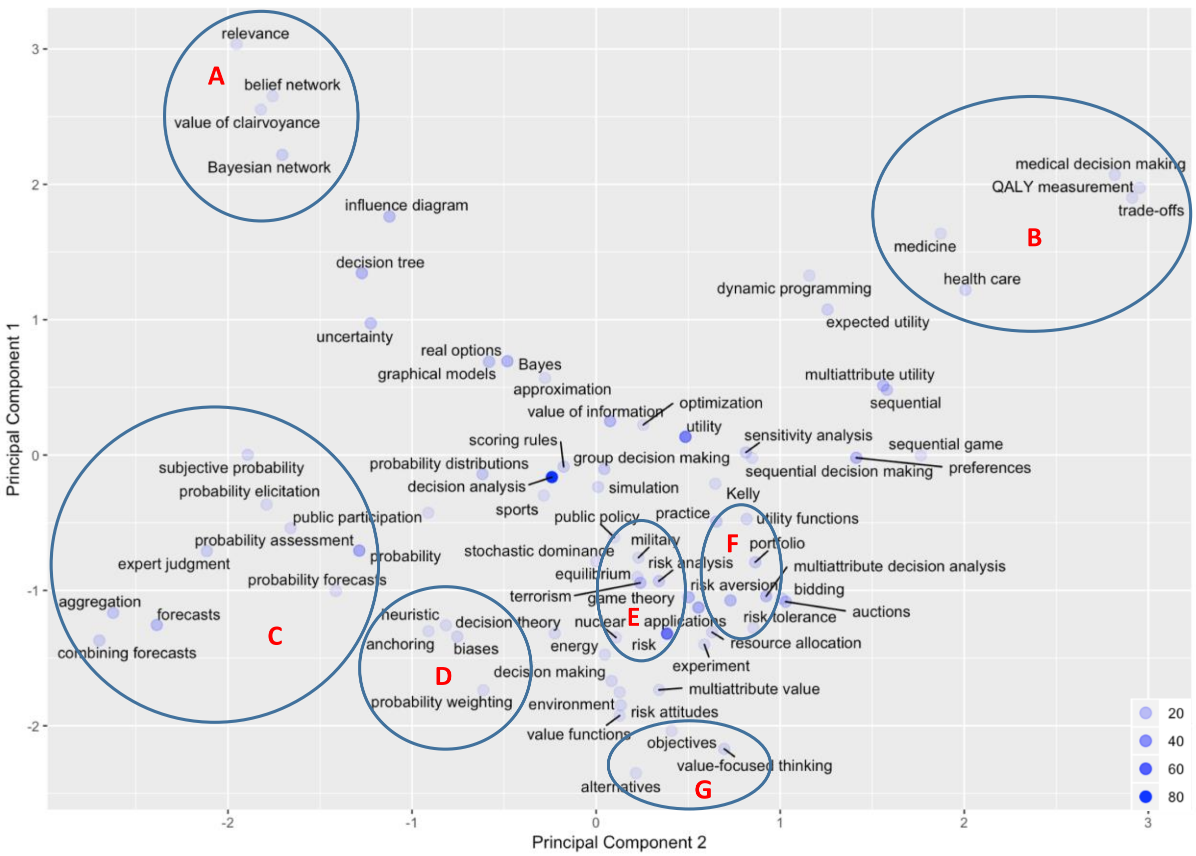

Figure 2 Latent-space Representations of Key Words

The latent-space representations provide a map for the contextual relationships between all keywords. If two words or phrases are located close to each other, they either co-occur together frequently or have similar co-occurrence patterns with other words and phrases, both implying conceptual similarity. For example, the area marked with “A” in Figure 2 is related to Bayesian network; the area marked with “B” involves trade-offs in medical decision making; the “C” area covers the topics of probability elicitation and aggregation; the “D” area is about biases and heuristics; the “E” area covers key words in counterterrorism and military; the “F” area contains topics about risk attitudes and multiattribute decision analysis; and the “G” area is about value-focused thinking. Moreover, relative locations of the topics reflect their semantic distances. For example, the “C” area (on probability elicitation and aggregation) is closer to the “D” area (on biases and heuristics) than the “A” area (on Bayesian network).

Figure 3 shows the distributions of the key words chosen by authors from the top four countries, including the United States, Germany, Canada, and the United Kingdom. From these maps we can get a brief idea about how different the research interests are among different countries. Papers from the United States spanned the entire map of keywords. Authors based in Germany supplied keywords that covered most topics except for forecast aggregation and trade-offs in medical decision making. Comparatively, authors based in Canada focused less on Bayesian networks, dynamic decisions in health care, and aggregation. Authors based in the United Kingdom used fewer keywords related to health care and medical decision making.

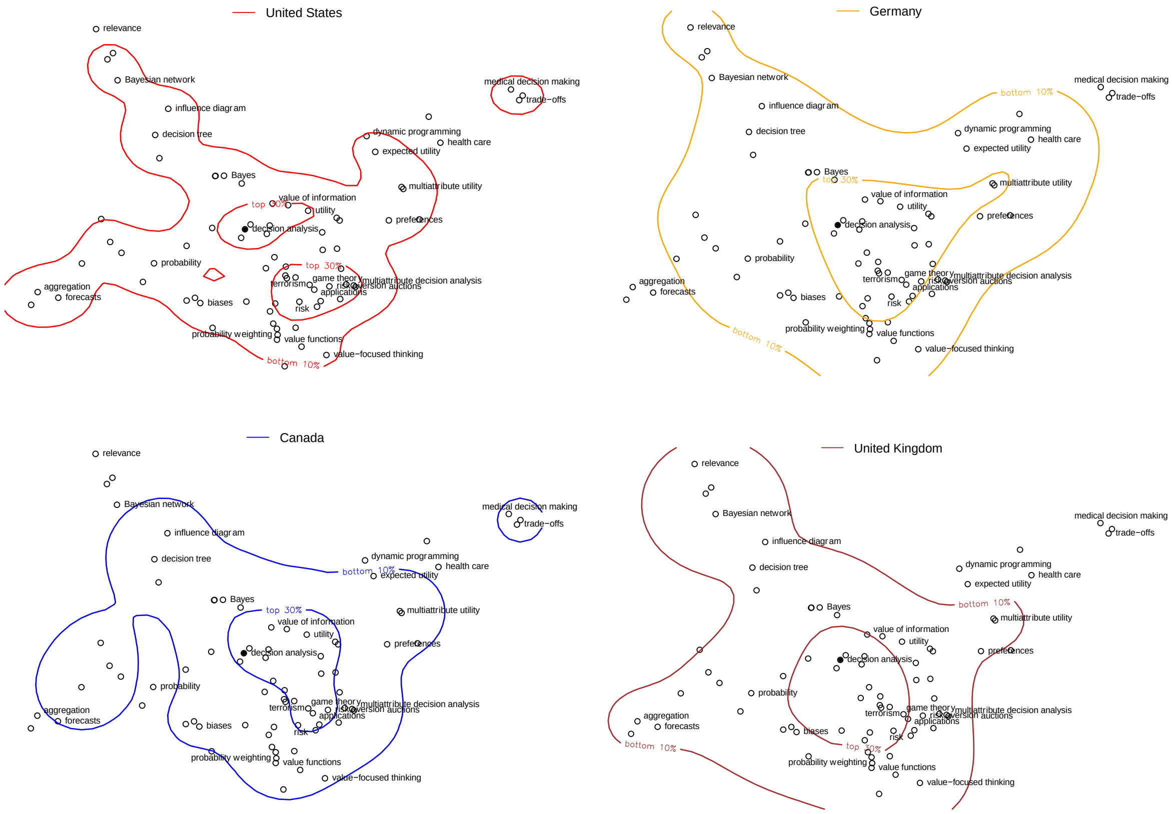

Figure 4 Distributions of Key Words for Top Countries

One caveat of this analysis is that keywords are selected by authors, who may have different preferences when deciding whether a keyword should be more generic or specific. For example, in the field of counterterrorism resource allocation, some authors prefer generic keywords such as "game theory," while other authors include keywords for specific solution approaches such as "Karush–Kuhn–Tucker conditions." It is also possible that some authors talk about "QALYs" (i.e., quality-adjusted life years) a lot in the paper, but never mention it in the list of keywords. Therefore, the *word embedding* results need to be treated with caution. Even so, the latent-space representations are still useful for organizing the keywords into topics and demonstrating closeness between the various topics, as reflected in the authors' choices of keywords.

Next, apart from the possibly biased keywords, we look at the contents of the 284 papers to extract research topics directly. Here we apply another machine learning method called *topic modeling* to fulfill this task (Blei et al. 2003). *Topic modeling* is capable of automatically discovering main themes from a collection of documents. Each topic is constructed by assigning occurrence probabilities to each of the words in the vocabulary. Topics are differentiated from each other by the varying probability distributions. Each document is allowed to contain multiple topics, which is depicted by a discrete distribution over all topics. We implemented the Latent Dirichlet Allocation approach of *topic modeling* on the 284 research articles (using the python module "lda", see https://pythonhosted.org/lda), where estimation was done by variational Bayes inference. For the ease of interpretation, we set the number of topics to be 11. Figure 5 displays the top 20 words with the highest probabilities in each of the 11 topics.

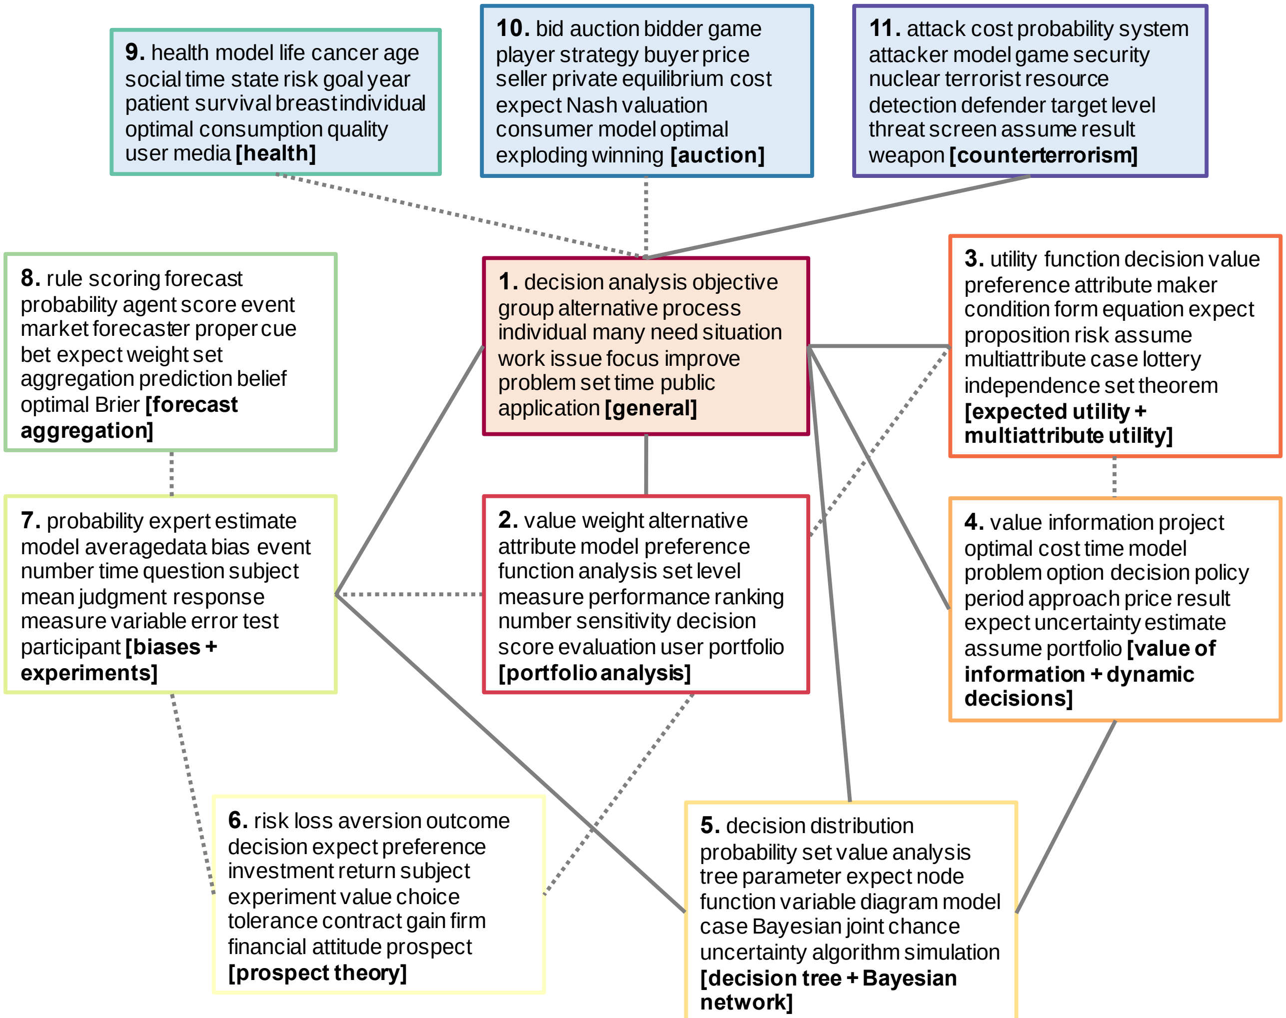

Figure 5 Top Words in the 11 Topics

Topic #1 assigns high probabilities to general concepts such as decision, analysis, objective, group, alternative, process, and individual. Almost every paper published in *Decision Analysis* mentions these words fairly often. Moreover, because of the limited number of topics, some research areas with similar word usage (e.g., value-focused thinking, group decision making, public decisions) are merged within this topic implicitly. Topic #2 also consists of quite general terms such as value, weight, alternative, preference, ranking, sensitivity, and portfolio. To differentiate this topic from Topic #1, we label this topic as "portfolio analysis" for convenience. Topic #3 is about utilities, including expected utility and multiattribute utility. Topic #4 involves value of information and dynamic decisions. Topic #5 places high weights on words related to uncertainty modeling, decision tree, Bayesian network, and Bayes' decisions.

Topic #6 emphasizes elements of prospect theory, as well as concepts of risk preferences in general. Topic #7 covers words about judgment biases and experimental studies. Topic #8 is related to probability elicitation and aggregation. Topics #9, #10, and #11 correspond to applications of decision analysis in health care, auction, and counterterrorism, respectively. Links between each pair of topics indicate how often they simultaneously appear as top-three topics in a document. For example, Topic #1 (the general topic) connects with all other topics except for Topic #6 (prospect theory) and Topic #8 (forecast aggregation). Topic #6 (prospect theory) is connected to Topic #7 (biases and experiments).

Figure 6 shows the change of topics over time. We only considered the top three topics of each document in this analysis and treated the presence of the remaining topics as random noises. The results are generally consistent with those obtained using author-supplied keywords. During 2004-2006, the most popular topics were general decision analysis, decision tree (and other graphical models), biases and experiments, and value of information and dynamic decisions. During 2010-2012, utilities, auction, and counterterrorism started to attract more researchers. During 2016-2018, *Decision Analysis* authors became more interested in experimental studies, judgment biases, and prospect theory.

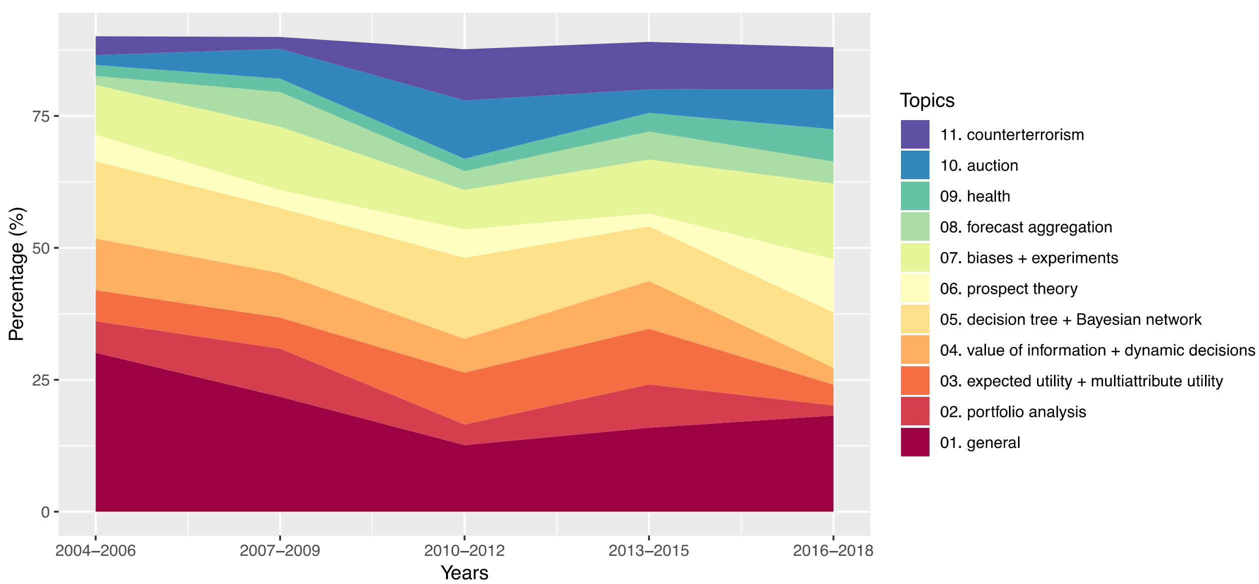

Figure 6 Topic Dynamics of *Decision Analysis* Papers

Figure 7 presents topic breakdowns for each of the top seven countries. Again, only the top three topics of each document were retained in our analysis. We can observe pronounced differences in research interests among these countries. Interests of the authors from the United States were roughly evenly distributed among the various topics, with relatively more attention paid to general decision analysis, decision tree (and other graphical models), biases and experiments, value of information, and counterterrorism. By contrast, authors from Germany were more interested in biases and experiments, prospect theory, general decision analysis (especially value-focused thinking), and health care. Studies based in Canada focused on auction, forecast aggregation, and utility theory. Authors from France also favored the topics of auction and utility theory, as well as value of information and dynamic decisions. Papers by authors from the

United Kingdom wrote more about public decisions, decision tree (and other graphical models), Bayes' decisions, and utility theory. Authors based in Spain published more work on decision tree (and other graphical models), auction, and counterterrorism. Authors from Finland focused on general decision analysis, multiattribute utility, and portfolio analysis. We validated these results by manually checking the original papers, and found good consistency between the automatically extracted topics and the true contents. Of course, if we use a larger number of topics, some topics will be split into subtopics which are more easily differentiable from each other.

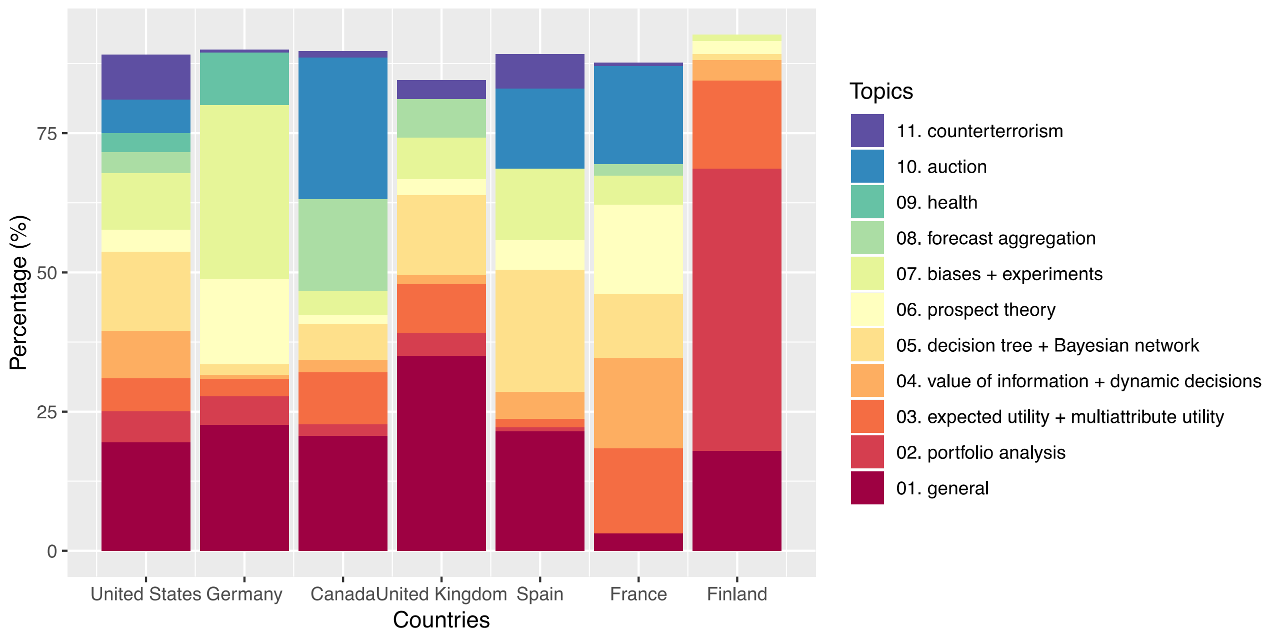

Figure 7 Distributions of Topics for Different Countries

With the development of natural language processing methodology, various techniques for constructing knowledge graphs for the academic achievements in decision analysis are available. In the current issue, we applied *word embedding* and *topic modeling* to analyze key words and topics in the journal of *Decision Analysis*, and compared the research interests of authors from different countries. Using these analysis tools, we can expand the investigation to a wider range of decision analysis publications including books and journals. Moreover, these tools also make it feasible to trace how a given topic develops, from its origin to variations to potential directions for evolution, across disciplines and across borders. We hope to continue with this exploration in future issues.

References

Blei, D.M., Ng, A.Y. and Jordan, M.I., 2003. Latent dirichlet allocation. *Journal of Machine Learning research*, *3*(Jan), pp.993-1022.

Grover, A. and Leskovec, J., 2016, August. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 855-864). ACM.

DA Practice

Column Editor: Pat Leach

Machine Learning, Machine Decisions?

In November, I attended a Decision Quality Conference in London, jointly put on by the European Decision Professionals Network, the London Decision Quality Group, and the Society of Decision Professionals. It was a great conference, filled with interesting presentations, interactive learning experiences, and stimulating discussions.

At one point, the subject of artificial intelligence (AI) and machine learning came up. This is the area of research in which a computer is programmed with algorithms that enable it to learn from its own experiences and thus improve its performance over time, rather than needing a human being to supply it with instructions on how to react in almost any situation. Human coaches are needed initially to provide context and guidance as the machine learns, but eventually it gets to a point where it has internalized enough of the broader framework to be able to continue, and even accelerate, the kind of performance improvement that we all recognize as "learning."

Perhaps the most famous of these computers is AlphaGo, developed by Alphabet Inc.'s Google DeepMind in London. This machine was taught the basics of the board game, Go, and then trained itself to become better by playing literally millions of games against itself in a very short period of time. It then went on to defeat a human master Go player, something many people thought would either never be possible or would only happen years into the future.

"So what?", you might ask. "Didn't IBM's Deep Blue defeat Garry Kasparov in chess some years back?" Yes, it did—but Go is a much more complex game than chess. The number of possible moves is infinitely larger, the variety of strategies is much broader, the game is far less formulaic (especially the opening), and the objectives are far more nebulous than in chess. The consensus among gamers and computer scientists alike was that teaching a computer how to succeed at Go would be far more difficult than teaching it how to win at chess. And so, the AlphaGo team didn't try to teach the computer; they let it teach itself.

But back to the London conference. I asked a question of the gentleman who had talked about AI, essentially, "How far do we go with this? To what degree are we comfortable delegating any but the most routine decisions to machines, regardless of how clever they are, regardless of how sophisticated the algorithms may be that allow them to learn and improve much faster than humans can?" The gentleman replied that we are always uncomfortable with new technologies until we achieve a certain level of familiarity with them. He made an analogy to touch-and-go credit cards, about which people were initially nervous (what information is my card really sending out here?), but which people don't think twice about anymore.

I did not find his answer very satisfactory, and I still don't. There are people, real human beings, who can tell me exactly what happens when I use a touch-and-go credit card, what information is being transferred, and what information remains encrypted on the card. They can predict with great accuracy what will happen the next time I use my card. I may not understand how this technology works, but there are those who do.

Not so with machine learning. In the course of defeating the human Go master, AlphaGo made several moves that no one in the room could explain—not Go masters, not the people who built the machine, not the people who guided, trained, and coached it during its learning phase—no one. At first, these out-of-the-blue moves looked crazy, possibly even erroneous. Only as the game progressed and AlphaGo developed a clearly advantaged position did the Go masters realize that the earlier moves had been sheer genius. Even then, they could not explain exactly what reasoning had driven AlphaGo to make these moves, nor could they file these new insights away in their brains for use in their own future games. The logic was simply too complex. As such, the Go masters could never predict how the machine would play in the future, either.

No one had taught AlphaGo how to reason like this; it had developed this logic on its own. That's not much of a concern when you're just playing a board game, but most decisions we make have more serious consequences than games of Go. I still have the same nagging question that I asked in London: To what extent are we comfortable delegating decisions to self-taught computers, machines that use their own reasoning, their own logic, and no one on Earth truly understands how they arrive at their conclusions?

I strongly suspect I'm not the only person who is disquieted by the thought of taking this step. After all, a very common theme in apocalyptic science fiction is that machines become more and more intelligent, to the point where they realize that the Earth would be far better off without human beings (which is very likely true; with the possible exception of cockroaches, dairy cows, and domestic dogs, pretty much every species on the planet would be better off if humans went extinct). I'm not seriously predicting a *Terminator*-like future; I'm just saying that putting important decisions into the hands (circuits?) of self-taught computers requires a leap of faith that I'm not sure I'm ready to make.

Now, one could argue that keeping important decisions in the hands of flesh-and-blood human beings isn't such a great idea, either. We're horribly biased, we see what we want to see, and our capacity to process data is extremely limited compared to modern computers. But we can deal with abstractions like ethics, fairness, beauty, wonder, empathy, social cohesion, and morality. Most of the really big decisions we face as a society are rooted in abstractions like these. I'm not ready to trust AlphaGo and its descendants with these decisions, regardless of how good they are at board games.

Some of you are probably rolling your eyes at my Luddite-ism, so let's put it into Decision Quality terms. One key feature of a good decision-making process is transparency. Any model should be auditable, and one should be able to follow the logic in that model. Any time an individual or a team chooses A over B, they should be able to explain why. In the world of decision analysis, black boxes are a no-no. But a black box is exactly what a self-taught machine is. Are we comfortable with that?

Research

Column Editor: Roshanak Nateghi

The Challenge of Analyzing Global Catastrophic Risks

Seth Baum

Global catastrophic risks are risks of the highest severity, regardless of their probability. Exact definitions vary, but they all point to catastrophes more severe than anything in recent human history. I tend to define global catastrophe as the collapse of modern global civilization. Some people in the field focus specifically on risks to human extinction, often using the term existential risk.

The concept of global catastrophe has a long history in theology, such as in the notion of apocalypse. The modern scientific study of global catastrophe can be traced to 1940s research on the threat from asteroids and a Manhattan Project study of the possibility of nuclear detonations igniting the atmosphere. Today, the field covers threats of natural origin including asteroids, comets, and volcanoes, as well as anthropogenic threats such as nuclear war and global warming. There is considerable interest in threats from future technologies, in particular biotechnology and artificial intelligence (AI).

Much of the interest in global catastrophic risk comes from big-picture thinking about human civilization and its role in the world and the universe. Indeed, people drawn to global catastrophic risk often have backgrounds in philosophy and cosmology. The essential idea is that a global catastrophe could be a discontinuity in the course of human history. Noting the prospect for humans or our long-term descendants to expand into outer space, the stakes may be literally astronomical.

Similarly, the case made for the importance of global catastrophic risk typically centers on a moral concern for future generations. A global catastrophe could substantially diminish future generations or eliminate them entirely in the event of human extinction. This extreme severity could make global catastrophic risk an unusually large class of risk, even if the probability is low. However, there has been little work comparing the long-term consequences of global catastrophes to the long-term consequences of smaller events. This is an important open area of research.

The Analytical Challenge

The prospect of comparing long-term consequences speaks to a wider theme: the global catastrophic risks are quite challenging to analyze. While long-term consequences are difficult to analyze for any risk, even the short-term consequences of global catastrophe are opaque. How resilient is modern global civilization to extreme catastrophes? If civilization collapses, what happens next? These questions do not have clear answers, but they are vital for evaluating the severity of global catastrophe.

The probability of global catastrophe is also hard to pin down. There is an inherent lack of data: modern global civilization has never previously been destroyed. Compounding the challenge is a phenomenon known as the observation selection effect: if a global catastrophe had previously occurred, we might not be alive to observe that fact. Rigorous estimates of the probability of global catastrophe must account for the lack of data and the observation selection effect, both of which make the analysis more difficult than it is for many other risks.

For some risk management decisions, it is not essential to rigorously quantify the global catastrophic risks or be able to compare them to other possible outcomes. To take one very simple example, it is probably good for us to turn the light off when we leave the room. Doing so could ever so slightly reduce the risk from global warming and have other benefits such as saving money on energy. The “cost” of turning the light off is usually fairly trivial—just the tiny effort it takes to flip the switch. One does not need a rigorous decision analysis accounting for long-term consequences and observation selection effects to reach the conclusion to flip the switch.

For other decisions, quantitative decision analysis is more important. For example, how much should society focus on global catastrophic risk relative to other issues? This question is essential for decisions on the allocation of scarce resources such as policymaker attention and philanthropic funding. Also, what should be done when some action poses a trade-off between a global catastrophic risk and something else, or between multiple global catastrophic risks? These trade-offs are quite common and worth illustrating with a few examples.

First, consider the decision of whether states should disarm nuclear weapons. This is currently a major point of contention for the international community. Disarmament advocates often emphasize the outsized severity of nuclear detonations, while disarmament opponents often emphasize the advantages of nuclear deterrence. Essentially, nuclear weapons could increase the severity of major war and decrease its probability. Nuclear disarmament could have the reverse effect. For nuclear disarmament, would the decrease in severity be large enough to offset any increase in probability?

Second, consider the decision to build nuclear power plants. Nuclear power can reduce the risk from global warming by shifting electricity production away from fossil fuel. It can also increase the risk of nuclear war by facilitating proliferation—this is seen especially clearly in the international concern over the nuclear program of Iran. Nuclear power can have other effects as well, including benefits for local air pollution, the risk of meltdown, and the displeasure of anti-nuclear citizens. Accounting for all these factors, under what circumstances would nuclear power be of net benefit to society?

Third, consider the decision to restrict the development of advanced AI. Forward-thinking scholars of AI consider the possibility that AI systems may eventually be able to outsmart humans across a wide range of domains, similar to how they can now outsmart us in a few specific ones like chess and the board game Go (and simpler domains like multiplication). There may even be some nonzero probability of runaway AI scenarios in which humans lose control, with potentially catastrophic consequences. However, restricting the development of AI could deprive society of many benefits that come from this technology, potentially including reductions in other global catastrophic risks. Therefore, in what ways, if at all, should the development of AI be restricted?

Each of these three decisions could be of profound consequence for society, yet none of them are easy to evaluate. They involve trade-offs between multiple global catastrophic risks, or between global catastrophic risks and other important values. Evaluating these trade-offs can require a quantitative analysis of the global catastrophic risks themselves and the relative importance of global catastrophic risk to other values. Such analysis is difficult, but it is important to pursue in order to inform these and other decisions.

Quantitative Analysis of Global Catastrophic Risk

There is a nascent effort to apply rigorous risk and decision analysis to quantitatively analyze the global catastrophic risks. Leading this effort are the Garrick Institute for the Risk Sciences at UCLA (https://www.risksciences.ucla.edu) and my own group, the Global Catastrophic Risk Institute (http://gcrinstitute.org). Our aim is help society understand the global catastrophic risks and make better decisions on them.

Central to this effort is the synthesis of relevant information. While there is no data on past global catastrophes, there is nonetheless a lot of information that can help evaluate the global catastrophic risks and the opportunities to reduce them. There are historical near-miss events that may have gone partway to global catastrophe, such as the Cuban missile crisis. There is information on the mechanisms underlying global catastrophes, such as the physics of asteroid-Earth collision. There are indicators of potential changes in risks, such as the decline in Russia-U.S. relations since the recent Ukraine crisis. And there are the informed opinions of subject matter experts. All of these sources of information can be used to quantify global catastrophic risk, even in the absence of historical event data.

My group has been especially active in quantifying the risk of nuclear war. As with many risk analyses, we proceed by building models of how nuclear war could occur and what its consequences could be, by synthesizing available information of relevance to model parameters, and by crafting probability distributions to depict the uncertainty about model parameters.

Our model of the probability of nuclear war (available at https://ssrn.com/abstract=3137081) distinguishes between two primary types of nuclear war. First, there is "intentional" nuclear war, in which state leadership makes the decision to initiate nuclear war. This can come either from escalating a conventional (non-nuclear) war or from a non-war crisis, such as the Cuban missile crisis. Second, there is "inadvertent" nuclear war, in which state leadership mistakes some event as a nuclear attack and launches nuclear weapons in what it believes is retaliation but is in fact the first strike. The other event could be a non-war nuclear detonation, such as a nuclear terrorist attack, or a false alarm of nuclear detonation, such as a warning system malfunction.

To inform the quantification of the probability model parameters, we compiled a dataset of 60 historical near-miss incidents that may have threatened to turn into nuclear war. (These incidents are data points in the class of nuclear war near-misses, not the class of nuclear wars, though there is one nuclear war data point, World War II.) Post-WWII incidents span from the 1946 Azerbaijan crisis—in which U.S. President Harry Truman allegedly threatened the Soviet Union with nuclear attack when it initially refused to leave Iranian Azerbaijan after WWII—to the 2018 Hawaii missile alert, in which the Hawaii Emergency Management Agency accidentally sent out a mass text message warning of an incoming ballistic missile.

A next step for this work is to assess these near-miss incidents in terms of how close they came to nuclear war. It is clear that some, such as the Cuban missile crisis, came closer than others, such as the Hawaii missile alert. For the Hawaii missile alert to have resulted in nuclear war, someone with nuclear launch authority would have needed to receive and act on the warning within the 38 minutes before a correction message was sent. There is no indication that anything along these lines came close to occurring.

Quantifying how close these incidents came to nuclear war is a greater challenge. The historical record is ambiguous, and scholars of nuclear war disagree on how close the incidents came. This disagreement resembles disagreements on the interpretation of other types of near-miss events documented by analysts such as Robin Dillon-Merrill and Catherine Tinsley. Their work can be a valuable resource for the study of nuclear war near-misses and similar events of relevance to global catastrophic risk. More generally, the base of knowledge built up over the decades in risk and decision analysis has much to offer to the analysis of global catastrophic risk.

Global Catastrophic Risk Decisions

The difficulty of analyzing global catastrophic risk is only one challenge to effective decision making about global catastrophic risk. The other major challenge is ensuring that careful analysis is used by the decision makers.

In many cases, risk and decision analyses are commissioned by decision makers who want the analysis to inform their decisions. In these cases, the findings of the analysis do not necessarily weigh heavily into the decision, but there is at least a clear path from analysis to decision. However, in global catastrophic risk, this is generally not the case. Risk management institutions are typically charged with focusing on risks that are local in scale with higher probability, while decision makers on important aspects of global catastrophic risk (for example, political leadership of nuclear-armed states) tend not to approach their decisions in risk and decision analytic terms.

Meanwhile, analysis of global catastrophic risk is often motivated by an intrinsic concern about global catastrophic risk held by the analysts. That is certainly the case for myself: in reflecting upon my own moral values, I find the reduction of global catastrophic risk to be important, and my work proceeds accordingly. Though perhaps admirable, this situation lacks a built-in audience for the risk and decision analysis.

In response to this situation, my colleagues and I have developed two approaches. One is to promote the idea that global catastrophic risk reduction is important. While it is not our place to tell others what they should care about, we can at least explain why we have reached the conclusion that global catastrophic risk reduction is important, in particular the outsized severity of global catastrophe and the availability of opportunities to reduce the risk. In some cases, this “direct” approach is enough to generate interest.

The other approach seeks opportunities to reduce global catastrophic risk that are consistent with whatever other values the decision makers hold. This approach is “indirect” because it does not directly emphasize the importance of global catastrophic risk. It is inspired by the concept of “mainstreaming” developed by the natural hazards community. To mainstream is to take a less prominent issue, such as natural hazards or global catastrophic risk, and embed it into a more prominent (i.e., mainstream) issue, such as economic development. Doing so makes it easier for decision makers to factor those other issues in by not forcing them to restructure their agenda. In general, people do prefer to reduce their vulnerability to natural hazards and to reduce global catastrophic risk. Often, it's just a matter of finding opportunities that make sense from their perspective.

To that end, we find that global catastrophic risk decision analysis should include the social science of decision makers alongside the technical analysis of the risks and decision options. The social science includes understanding specific people who make the decisions as well as the social and institutional contexts in which the decisions are made. Including the social science does add considerably to the complexity of what is already a challenging set of risks to analyze. However, in the interest of actually reducing the risks, we believe it is vital.

Editorial Team

DAS Officers

President:
Karen Jenni
Science and Decisions Center
U.S. Geological Survey
kjenni@usgs.gov

VP/President-Elect:
Emanuele Borgonovo
Department of Decision Sciences
Bocconi University
emanuele.borgonovo@unibocconi.it

Past President:
Jason Merrick
Department of Statistical Sciences &
Operations Research
Virginia Commonwealth University
jrmerric@vcu.edu

Secretary-Treasurer:
Melissa A. Kenney
Environmental Decision Analysis and
Indicators
University of Maryland
kenney@umd.edu

Social Media Officer/Webmaster:
Jay Simon
Kogod School of Business
American University
jaysimon@american.edu

DAS Council

Léa Deleris
Risk Artificial Intelligence Research
BNP Paribas
Lea.deleris@bnpparibas.com

Valentina Ferretti
Department of Management
London School of Economics
v.ferretti@lse.ac.uk

Andrea Hupman
College of Business Administration
University of Missouri-St. Louis
hupmana@umsl.edu

Eva Regnier
Graduate School of Business and Public Policy
Naval Postgraduate School
eregnier@nps.edu

Heather Rosoff
Sol Price School of Public Policy
University of Southern California
rosoff@usc.edu

Matthias Seifert
Operations and Technology Department
IE Business School
Matthias.Seifert@ie.edu