Stony Brook University

On Using Risk-Benefit Analysis to Assess and Manage Controversial Research

A Dissertation Presented

by

Daniel James Rozell

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Technology, Policy and Innovation

Stony Brook University

December 2015

Stony Brook University

The Graduate School

Daniel James Rozell

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation.

Sheldon J. Reaven – Dissertation Advisor
Associate Professor, Department of Technology and Society

David L. Ferguson - Chairperson of Defense
Distinguished Service Professor and Chair, Department of Technology and Society

David J. Tonjes
Research Associate Professor, Department of Technology and Society

Scott Ferson
Senior Scientist, Applied Biomathematics

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

On Using Risk-Benefit Analysis to Assess and Manage Controversial Research

by

Daniel James Rozell

Doctor of Philosophy

in

Technology, Policy and Innovation

Stony Brook University

2015

Science research in the twenty-first century increasingly results in powerful and rapidly disseminated technologies that raise concerns over technological risk and public safety. A common recommendation for assessing and managing controversial research is to perform a quantitative risk-benefit analysis (RBA) in order to help determine an appropriate course of action. However, these recommendations lack guidelines on how to perform such an assessment. This dissertation lays out a detailed argument why a quantitative RBA should not be used as a basis for making complex policy decisions. Rather, RBA should be regarded as a tool for exploring the impacts of technology on society. This argument is made in four parts: (1) a comparison of the various legitimate, but incomparable ways that research benefits are assessed; (2) a guide to the many value-laden assumptions made in risk assessments that undermine unqualified claims of objectivity and neutrality; (3) a clarifying discussion of the technological risk attitudes that underlie research controversies; and (4) a discussion of how this inherent subjectivity favors particular risk management techniques. In applying these arguments to the use of RBA for controversial research, the concern is that RBA is unlikely to build consensus

because the results of even the most mathematically sophisticated assessments tend to be too epistemically narrow or ethically controversial to resolve science policy disputes. These arguments are applied to a contemporary case study – controversial "gain-of-function" research involving highly pathogenic avian influenza. My analysis argues that the debate should place less emphasis on attempts to quantify current risks and benefits which are, in themselves, controversial. Rather, the least problematic resolution requires shifting away from traditional biosafety and biosecurity measures and towards more inherently safe research techniques that accomplish the same goals. This analysis is applicable to other contemporary research controversies such as those surrounding synthetic biology and geoengineering.

Table of Contents

1 Introduction 1

1.1 The H5N1 Debate 1

1.2 The Changing Face of Technological Risk 5

1.3 Biosafety Concerns 7

1.4 The Role of Self-Regulation 9

1.5 Problem Statement 9

1.6 Purpose 10

1.7 Overview of Chapters 12

2 Assessing the Benefits of Research 13

2.1 Methods of R&D Valuation 15

2.1.1 R&D as a jobs program 15

2.1.2 Econometric valuation 16

2.1.3 Valuation by knowledge output 20

2.1.4 Valuation by multiple metrics 25

2.1.5 Value-of-Information analysis 28

2.1.6 Qualitative assessment by expert opinion 33

2.2 Implications of Multiple R&D Valuation Methods 37

2.3 From R&D Valuation to Allocation 39

2.3.1 Ethics of allocation 40

2.4 Conclusion 42

3 Values in Risk Assessment 45

3.1 Categorizing Risk Assumptions 46

3.2 Selection of Topic 47

3.3 Defining System Boundaries 48

3.4 Method and Units of Assessment 49

3.4.1 Unit of assessment 49

3.4.2 Value of life 50

3.4.3 Discount rate 51

3.4.4 Other methodological considerations 52

3.4.5 Treatment of error 54

3.5 Model Selection 55

3.5.1 Theoretical versus empirical choices 56

3.5.2 Level of complexity 58

3.6 Data Collection 59

3.6.1 Data screening 60

3.6.2 Data for rare events 61

3.6.3 Expert opinion 62

3.7 Accounting for Uncertainty 63

3.7.1 Deterministic versus probabilistic 63

3.7.2 Objective versus subjective probabilities 63

3.7.3 Hybrid probabilistic methods 64

3.7.4 Non-probabilistic methods 65
3.7.5 Selecting an uncertainty representation 66
3.7.6 Treatment of variable dependence 68
3.7.7 Model uncertainty 69
3.7.8 Unknown unknowns 72
3.8 Comparing Risks 73
3.8.1 Incommensurability 74
3.8.2 Risk ranking 76
3.9 A Value Assumption Roadmap 77
3.9.1 Example: risk analysis of farmed salmon 78
3.10 Discussion 80
3.10.1 Choice of ethical framework 81
3.10.2 Ignoring the normative 84
3.11 Conclusion 85
4 Technological risk attitudes 88
4.1 Technological Optimism and Skepticism 89
4.1.1 Technological optimism 90
4.1.2 Technological skepticism 91
4.1.3 Explaining the differences 92
4.1.4 Variability in technology attitudes 96
4.2 Principle of Controversial Research 99

4.3 The principle in action 102

4.4 Conclusion 107

5 Managing Controversial Research 108

5.1 Common risk management strategies 108

5.1.1 Engaging multiple perspectives 109

5.1.2 Critique of traditional risk management 112

5.2 Inherent safety 113

5.3 Recommendations 117

5.3.1 Example: Gene drive risk assessment 120

5.4 Conclusion 123

6 The gain-of-function research debate 124

6.1 Timeline 124

6.2 The Risk-Benefit Assessment 125

6.2.1 Quantitatively estimating the risks and benefits 130

6.2.2 Quantitative ≠ objective 135

6.2.3 Ebola: lessons from a current epidemic 137

6.3 Underlying Risk Frameworks 139

6.4 The Deliberative Process 141

6.5 Inherent Safety 144

6.6 Potential Outcome 147

6.7 Lessons on Managing Controversial Research 150

6.7.1 Post hoc management is difficult 151

6.7.2 Research controversies are predictably unpredictable 152

6.8 Relevance 152

6.9 Future Work 155

6.10 Conclusion 156

7 References 158

List of Figures

Figure 1: Ways of assessing the benefits of research ordered by increasing range of benefits that can be considered and also by the increasing uncertainty and subjectivity of estimates.. 15

Figure 2: Ways of valuing R&D 31

Figure 3: In this idealized scenario, option A is optimal for the most probable risk, but option B is more robust having a better design score for a wider range of outcomes. 54

Figure 4: Convolution of *n* uniform distributions with bounds (0,1) 67

Figure 5: NIH proposed deliberative outline (from Sharples et al., 2015) 125

Figure 6: RBA guidelines (modified from NSABB, 2012a) 128

List of Tables

Table 1: A process map of value judgments in risk assessment. Details for each question are discussed in the previous sections..... 78

Table 2: Accuracy of EVD predictions made in late 2014 138

List of Abbreviations

BSL – biosafety level

CDC – U.S. Center for Disease Control and Prevention

EVD – Ebola virus disease

FDA – U.S. Food and Drug Administration

GDP – gross domestic product

GOF – gain of function

HHS – U.S. Department of Health and Human Services

IUS – Innovation Union Scoreboard

IVF – *in vitro* fertilization

NASA – U.S. National Aeronautics and Space Administration

NIH – U.S. National Institutes of Health

NRC – U.S. National Research Council

NSF – U.S. National Science Foundation

NSABB – U.S. National Science Advisory Board for Biosecurity

OTA – U.S. Office of Technology Assessment

PPP – potentially pandemic pathogens

PSAC – President's Science Advisory Committee

RBA – risk-benefit analysis

R&D – research and development

VOI – value of information

WHO – World Health Organization

Acknowledgments

I would like to thank my dissertation committee for their insightful comments and friendly encouragement throughout this process. During our many interactions, both in and out of the classroom, each member has been generous with his time. I would also like to specifically thank: Dr. David Tonjes for his practical advice regarding graduate studies, Dr. David Ferguson for shepherding me (and other graduate students) through the early years of a new doctoral program, and Dr. Scott Ferson for introducing me to the field of quantitative risk analysis and finding me a job at an opportune time. Finally, I would like to thank my advisor, Dr. Sheldon Reaven, for the countless hours spent discussing technology assessment, risk analysis, the philosophy of science, and opera.

1 INTRODUCTION

The media coverage of the latest scientific discoveries and technological innovations often optimistically describes research that satisfies our curiosity, improves our health, increases our productivity, and raises our standard of living. However, sometimes the results of research also inspire critical public reactions over ethical concerns or fears of unintended consequences, accidents, or intentional misuse. Let us consider a recent example in the popular media.

1.1 The H5N1 Debate

In September of 2011, Dutch virologist Ron Fouchier announced that his research team had recently engineered a version of the H5N1[1] avian influenza virus that appeared to have high mammalian transmissibility while retaining its original lethality (Herfst et al., 2012). A similar study headed by Yoshihiro Kawaoka of the University of Wisconsin-Madison was announced shortly thereafter (Imai et al., 2012) and the subsequent media coverage raised concerns among the general public and politicians and set off a renewed discussion and policy review of how dual-use[2] research is assessed, funded, and managed.

The initial response to the H5N1 research debate demonstrates the difficulty of assessing and managing controversial R&D. The two engineered avian influenza virus studies (Herfst et

[1] The formal naming convention for influenza viruses includes: the antigenic type (A, B, or C); the originating host (e.g., swine); the geographical origin; year of identification; strain number; and for type A viruses, the specific subtypes of two surface proteins, hemagglutinin (H) and neuraminidase (N) (e.g., H5N1) (Assaad et al., 1980). Given the complexity of formal names, popular shorthand names, such as Spanish flu, Swine flu, and H1N1 can potentially be referring to the same influenza virus. The World Health Organization (WHO) is attempting to improve shorthand names to make them less stigmatizing and more informative (Kupferschmidt, 2015).

[2] References to the term "dual-use" prior to the twenty-first century commonly denoted a technology with both military and civilian uses (e.g., Molas-Gallart, 1997). However, in this century, U.S. security policy has partially shifted from international conflict towards non-state threats. Dual-use policy discussions have likewise refocused on technologies that have both societal benefits and potential malicious uses by terrorists (van der Bruggen, 2012).

al., 2012; Imai et al., 2012), both funded by the National Institute of Health (NIH), intended to investigate the ease and potential mechanism by which the H5N1 virus would naturally become a more serious public health threat. However, many scientists and security experts became concerned that the papers detailing the experiments could be a blueprint for skilled bioterrorists. In an initial assessment of the papers, the U.S. National Science Advisory Board for Biosecurity (NSABB) recommended redacting sections describing the methodology of each paper to limit reproducibility. This was the first recommendation of publication restriction since the board's formation in 2005. A subsequent review by a panel of experts at the World Health Organization recommended full publication (WHO, 2012). After a second review four months later, the NSABB reversed its position and recommended full publication (NSABB, 2012b) and the U.S. government simultaneously released a new policy clarifying what was dual-use research of concern (Malakoff, 2012). Critics claimed this reversal was based on the unpopularity and impracticality of redaction rather than legitimate public safety considerations (Maher, 2012a, 2012b).

Following this appearance of inconsistency in the review process, a voluntary H5N1 research moratorium was agreed upon by prominent influenza research laboratories (Fouchier et al., 2012) until new guidelines could be put in place. After almost a year, new guidelines were released by the U.S. Department of Health and Human Services (HHS) (Malakoff and Enserink, 2013). The HHS policy for funding H5N1 "gain-of-function"[3] research proposals required that: (1) the engineered virus could have naturally evolved; (2) the research addressed an important public health concern; (3) there were no less risky alternative methods to accomplish the same

[3] Gain-of-function research purposefully induces mutations in an organism or disease agent to add new functions or amplify existing functions (e.g., increasing transmissibility or virulence).

goal; (4) biosafety issues for laboratory workers and the public were addressed; (5) biosecurity issues were addressed; (6) the research results were intended to be widely available, i.e., not classified; and (7) the funding would have risk reduction oversight mechanisms (Patterson et al., 2013). The review process was later extended to avian influenza H7N9 research (Jaffe et al., 2013). Like the various policies before it, the HHS policy did not address the method by which the risks and benefits of a particular research proposal were to be assessed or how mitigating actions were to be selected. In 2006, the NSABB made a similar move when it asked authors, institutional review boards, and journal editors to perform a risk-benefit analysis (RBA) before publishing dual-use research, but did not provide detailed guidance on how to perform such an analysis (Bhattacharjee, 2006; Rappert, 2014).[4] Likewise, when the Dutch government required Fouchier to apply for an export license before publishing his research (Enserink, 2013), the application was granted within days and determination was based on an unspecified risk-benefit criterion (Butler, 2013).

The range of expert opinion regarding gain-of-function research on highly pathogenic influenza virus is substantial. Some scientists claim the work is critical to public health and relatively low-risk (Morens et al., 2012, 2013; Palese and Wang, 2012; e.g., Fouchier et al., 2013) while others claim the approach has no predictive ability useful to pandemic preparedness and presents an unnecessary public health risk (e.g., Mahmoud, 2013; Rey et al., 2013).

While useful as a tool for discussion, given the considerable uncertainty and subjectivity surrounding the risks and benefits of gain-of-function influenza research, there is considerable doubt that the various stakeholders can agree on a formal risk-benefit analysis (Chubin, 2013;

[4] Five years later, an NIH random survey of 155 life science journals found that less than ten percent had a written dual-use policy or reported reviewing dual-use manuscripts in the previous five years (Resnik et al., 2011).

Uhlenhaut et al., 2013; Casadevall and Imperiale, 2014b). Even the NSABB, when first recommending an RBA informed by scientific expertise, acknowledged that the final evaluation would be subjective (NSABB, 2007).[5] A review of biosecurity risk assessments (Rappert, 2014) also raises suspicion about the efficacy of risk assessments. For example, a three-year sample of 74,000 biological research manuscripts submitted to Nature Publishing Group resulted in only 28 flagged and no rejected manuscripts for biosecurity concerns (Boulton, 2012). In light of previous publications, such as the papers detailing the recreation of the polio virus (Cello et al., 2002) and the 1918 Spanish flu virus (Tumpey et al., 2005), the H5N1 papers are most notable in that they actually instigated a public debate and government action on biosecurity policy (Faden and Karron, 2012; Rappert, 2014).[6]

So, it is not surprising that the new HHS policy did not resolve the controversy and three papers published in 2014 renewed the debate surrounding gain-of-function flu research. One study (Sutton et al., 2014) made the H7N1 flu strain, which was not covered by the new HHS rules, transmissible in ferrets. A second study by Fouchier's lab (Linster et al., 2014) expanded on earlier H5N1 work. The third paper, published by Kawaoka's lab (Watanabe et al., 2014), detailed the engineering of a virus similar to the strain responsible for the 1918 flu pandemic in order to argue that another major pandemic could arise from the existing reservoir of wild avian flu viruses. Many critics were particularly disturbed by this paper because the University of Wisconsin-Madison biosafety review of the proposal failed to classify the work as dual-use

[5] More recent NSABB documents associated with the H5N1 controversy have partially reframed the issue by using the terms "qualitative" and "quantitative" rather than "subjective" and "objective."

[6] The insensitivity of the review process suggests that either very little dual-use research is being proposed or, more likely, there are fundamental problems with defining and recognizing dual-use research of concern (Wolinetz, 2012).

research of concern despite a consensus among biosecurity experts that it clearly was (Butler and Maher, 2014). Collectively, these results again raised concerns that an intentional or accidental release of an engineered virus from gain-of-function influenza research could be the source of a future pandemic (Kaiser, 2014c) – an ironic and deadly self-fulfilling prophesy.[7]

1.2 The Changing Face of Technological Risk

The H5N1 engineered virus research is representative of a new and growing category of R&D – relatively accessible, yet very powerful technologies. Whereas chemical technology was the main focus of dual-use concerns in the first half of the twentieth century and nuclear technology in the last half, biotechnology arguably constitutes the most significant challenge to dual-use research policy in the twenty-first century. The primary reason was recognized decades ago; many hurdles in the biosciences are more theoretical than practical (Cairns-Smith, 1978). That is, material resources are rarely the limiting factor. Laboratories around the world (cf. Kant and Mourya, 2010), including an increasing number of small unaffiliated labs (Gewin, 2013), can now create or order with relative ease what was only possible in a handful of state-of-the-art research facilities a few years before (Suk et al., 2011; Trevan, 2012; Adleman, 2014). The critical component is knowledge. Once the information is publicly available, the ability to create both immensely beneficial and harmful biotechnology becomes almost ubiquitous. Even practical difficulties caused by limited tacit knowledge (e.g., assumed professional knowledge, ignored or hidden details, and essential laboratory skills) are decreasing over time as the number

[7] It should be noted that the debate over research on potentially pandemic pathogens has taken place during a time of falling federal support for biodefense research. From 2010 to 2013, funding specifically for pandemic influenza and emerging infectious disease programs dropped from $953 to $585 million. For comparison, proposed funding for pandemics research in 2015 was less than half of the requested funding for radiological/nuclear preparedness (Boddie et al., 2014). More generic health security preparedness funding is relatively robust (Reardon, 2014b).

of experienced biotechnologists increase and modern electronic communication lowers the cost of detailed knowledge transfer (cf. Engel-Glatter, 2013; Revill and Jefferson, 2013). For this reason, managing potentially harmful biotechnology requires a fundamental shift in thinking because traditional nuclear weapons nonproliferation policy has focused on controlling materials as much as knowledge (Moodie, 2012; Tucker, 2012).

The H5N1 research example also hints at the difficulty of reactively managing controversial research. After research has been conducted and technologies developed, responses tend to be ad hoc and difficult to implement. This raises important questions of how to proactively manage technologies that have a high potential for negative consequences, but are difficult to govern because the technologies are already widespread (e.g., Miller et al., 2009) and knowledge, rather than materials, is the limiting factor. Suggested approaches for these cases concentrate on informal governance measures that rely on moral suasion such as codes of conduct, educational outreach, transparency, and whistle-blowing (Tucker, 2012). However, self-regulation appears to be an inadequate response to threats of intentional misuse (Kaiser and Moreno, 2012) as it requires virtuous behavior from scientists and engineers who, by definition, have less than virtuous intentions.[8] Given that there are few viable options for managing low-governability technologies, involuntary monitoring of R&D activities (e.g., electronic spying) by government intelligence agencies is a more likely scenario (Rozell, 2013). For example, Italian police have already investigated one potential case of scientists illegally selling bird flu virus to animal vaccine makers for personal profit (Margottini, 2014).

[8] For example, the deadly 2001 anthrax letters originated from an experienced research scientist (DOJ, 2010).

1.3 Biosafety Concerns

Despite the difficulties of addressing biosecurity concerns, self-regulation has an important role in reducing risk from emerging low-governability technologies by reducing unintentional harm. While dual-use security experts concentrate on intentional malevolent misuses of technology, the risk of harmful effects caused by accidents and other unintended consequences may be even greater (Berg, 2012; Hulme, 2013). For example, the 1977 Russian influenza outbreak was an H1N1 strain genetically identical to a virus that caused a 1950 epidemic (Rozo and Gronvall, 2015). The lack of mutations in the intervening 27 years suggests a frozen virus source – presumably an unintended laboratory release (Webster et al., 1992; Lipsitch and Galvani, 2014).[9]

Research accidents continue to occur even in the most modern and technologically advanced research facilities. A series of accidents at the U.S. Center for Disease Control and Prevention (CDC) in Atlanta led to the temporary closure of the anthrax and influenza research labs in July 2014 (McNeil Jr, 2014a) and the resignation of the head of a bioterrorism lab (McNeil Jr, 2014b). Incidents included: accidental exposure of researchers to the anthrax-causing *B. anthracis* (Biosafety, 2014; Tavernise and McNeil Jr, 2014), accidental contamination of a weak flu sample with a dangerous flu strain (Butler, 2014a), and the discovery of six vials of live smallpox virus in storage at a U.S. Food and Drug Administration (FDA) lab in Bethesda, MD (Fountain, 2014).[10] A review of U.S. reports of theft, loss, or release of select agents and toxins

[9] Another Russian accident occurred in April of 1979 when approximately 100 unintentional deaths from anthrax occurred in Sverdlovsk. In this case, the accidental release, caused by human error and a failed ventilation system, came from a facility secretly producing weaponized *Bacillus anthracis* in clear violation of the 1972 Biological and Toxin Weapons Convention (Leitenberg et al., 2012). It would appear that casualties (so far) from the massive bioweapons program of the former Soviet Union have been largely accidental rather than intentional.

[10] After smallpox was globally eradicated in 1980, only two secure labs (the CDC and Novosibirsk, Russia) were supposed to have *Variola* samples (Butler, 2014a). Proponents of

compiled by the CDC between 2004 and 2010 found no reports of theft, 88 reports of loss (of which only one was never resolved)[11], and 639 cases of accidental release (Henkel et al., 2012). The releases, fourteen of which occurred at biosafety level 4 facilities,[12] resulted in eleven laboratory acquired infections, but no fatalities or secondary transmissions.

These mishaps at the CDC led to the U.S. Office of Science and Technology Policy and the HHS to impose another temporary moratorium on NIH-funded gain-of-function research for influenza and two coronaviruses in October of 2014 (Kaiser, 2014d). The moratorium was intended to last until the NSABB[13] and National Research Council could assess the risks and benefits of these lines of research and make recommendations (Reardon, 2014c). The accidents also increased public attention to scientists calling for more deliberation regarding gain-of-function research on potentially pandemic pathogens (PPP) – most notably the Cambridge Working Group led by Harvard epidemiologist Marc Lipsitch (Kaiser, 2014c). The accidents even prompted the CDC Director, Thomas Frieden, to propose closing many BSL-3 and BSL-4 labs. There is no definitive list of these labs for public review or government oversight (Young and Penzenstadler, 2015). However, since the anthrax attacks of 2001, it is estimated that the number of labs working with PPPs has tripled and the number of workers with access or

delaying the eventual planned destruction of these official samples argue that other unknown samples or even intentional synthesis (Adleman, 2014) warrant preservation to test vaccines and antivirals in the future (Arita and Francis, 2014; c.f. Damon et al., 2014; Henderson and Arita, 2014; Reardon, 2015).

[11] A case of *Coccidioides immitis,* a pathogenic fungus that causes Valley fever, appears to have been destroyed during shipment.

[12] Biosafety levels (BSL), as defined by the CDC, range from BSL-1 to BSL-4 with the latter having the strictest protocols and equipment for handling potentially fatal infectious agents for which there is no vaccine or treatment.

[13] The first NSABB meeting occurred the same week after the moratorium announcement. It was the board's first meeting in two years and was preceded by the replacement of the remaining inaugural board members – half of the 23-member panel (Cohen, 2014).

exposure risk has increased by at least an order of magnitude – as did the number of reported lab accidents. According to a 2013 Government Accountability Office assessment, the increased number of labs and workers unintentionally increased rather than decreased national risk (Begley and Steenhuysen, 2014).

Subsequent events suggest that appropriate laboratory safeguards for reducing human error had yet to be implemented. In December of 2014, it was discovered that live Ebola virus was sent to the wrong lab within the CDC (Grady and McNeil Jr, 2014). In May of 2015, it was discovered that a U.S. Army lab in Utah may have shipped live, rather than inactivated, *B. anthracis* samples to approximately 30 U.S. and foreign labs over several years (Mcleary, 2015).

1.4 The Role of Self-Regulation

Because intentional harms are within the purview of security agencies, security specialists are trained to focus on signs of malicious intent. Conversely, it is in the area of unintended harm where self-regulation appears to be most appropriate. Signs of impending accidents and other unintentionally harmful R&D activities may best be detected by those fully immersed in the field who understand the formal procedures and tacit knowledge associated with the research (i.e., the members of the research community). This may be one reason that the NSABB, whose members are primarily science research experts, has been increasingly concerned with biosafety risk issues in the H5N1 debate even though the board was created to deal with biosecurity concerns (Fink, 2004).

1.5 Problem Statement

The HHS policy on gain-of-function influenza research is a typical response to an R&D controversy – an ad hoc directive to assess and monitor research in its early stages without specifics on just how to assess the risks and benefits of a R&D activity (cf. Renn, 2008).

Strangely, the lack of clear methods for conducting a risk-benefit analysis extends to even relatively narrow and uncontroversial research questions that commonly come before regulatory agencies, such as medical drug efficacy comparisons (Holden, 2003). This suggests many, if not most, risk-benefit assessments are in key respects incomparable (Ernst and Resch, 1996) because the scientists and engineers performing them are dealing with methodological vagueness by folding in the past experience, education, and R&D culture particular to their field and region. Risk-benefit analysis clearly has value insofar as it focuses attention on the public impact of research and discerns risk scenarios that may not have been studied or even thought of. However, it is not obvious that merely asking scientists to consider the risks and benefits of their work will result in due consideration and communication of risks in ways that satisfy policy makers and the general public (Slovic, 1987; Fischhoff, 1995; Slovic et al., 2004). RBA is often recommended as a policy mechanism for mitigating technological risk, but it is still unclear what its practical value is to science policy. This leads to the fundamental question of this work. What role does risk-benefit analysis have in *resolving* research controversies?

1.6 Purpose

This dissertation has several objectives. The primary purpose is to argue that the reason that no detailed RBA methodology exists is because no consensus can be found regarding exactly how to categorize and quantify the risks and benefits of any controversial research program. This objective is accomplished by analyzing the many legitimate, but incomparable ways in which benefits and risks can be measured and the various value judgments inherent in the process which render any claim of objectivity or neutrality for a controversial research assessment contentious itself. That is, disagreements over technological risk are compounded by disagreements over how to *assess* technological risk – making consensus even more elusive.

The practical aim of this work is to provide the most comprehensive and detailed argument so far as to why expectations of RBA as a quantitative decision tool should be tempered; it is better viewed as a tool for exploring and communicating the social impact of research and technology, providing insight, uncovering potential problems, and suggesting solutions previously unnoticed. Likewise, the discussion of the practical and philosophical difficulties underlying RBA should also help scientists and engineers perform assessments that better show specific areas of consensus and disagreement among researchers, the general public, and policymakers.[14]

Second, I argue that, in the absence of a convincing assessment, pre-existing technological risk attitudes guide technological risk management decisions. I note a general trend of permissive management of controversial research and suggest that, when possible, redesigning research to avoid controversy is the most viable way forward. Lastly, as a practical application, I apply this policy analysis to the highly pathogenic avian influenza virus research debate.

It should be noted that I apply the term "controversial" to any research that may have substantial ethical or physical risk concerns. While there are some important distinctions within this broad use of the term, which will be discussed in the last chapter, I use an umbrella classification because, from a policy perspective, it is not important whether public objections are based on the perception of immorality or danger. Regardless of the source, public objections require that important policy decisions be made in order to proceed with science research or technology development.

[14] While consensus within science can be detrimental to the process of science (cf. Feyerabend, 1978), some degree of consensus within democratic societies is necessary in order to implement science and technology policy.

1.7 Overview of Chapters

The second chapter reviews the popular methods of gauging the societal benefits of publicly funded R&D. Public funding is the primary source of recent controversial research and is the area where questions of social risks and benefits are most likely to arise. Also, the majority of discussions regarding the benefits of research have been in the context of research funding allocation. Discussing the strengths and weaknesses of each assessment approach, I argue that there is currently no reliable or universally acceptable way of valuing research. In the third chapter, the important epistemic and ethical value judgments inherent in assessing technological risk are detailed as they arise in the process. I argue that open acknowledgment of the value judgments (of both types) made in any assessment increases its usefulness as a risk communication tool. In the fourth chapter, underlying technological risk attitudes are discussed. The implications of the previous chapters for managing controversial research are discussed in the fifth chapter. The sixth chapter analyzes the avian influenza gain-of-function debate as an example of controversial research. I draw some general lessons from the analysis, discuss implications and qualifications, and suggest future work.

2 ASSESSING THE BENEFITS OF RESEARCH

Hard times invariably lead to tough questions. The global economic downturn that began in 2008 strained national budgets and increased scrutiny on how governments spend their limited revenues. Government funding of research and development (R&D)[15] is one category of discretionary spending that has received considerable attention. For example, in the United States (US), federal R&D spending dropped over 16 percent from 2010 to 2013 (Morello, 2013) and there were congressional efforts to limit funding for basic research without immediate potential economic returns (Lymn and Houston, 2013). The response from the scientific community was a call for renewed investment in basic research (Augustine and Lane, 2014; Mervis, 2014d) with congressional science proponents claiming an "innovation deficit" (Mervis, 2014b). However, without any clear justification of what constitutes optimal R&D support, proposed funding goals can appear arbitrary (Goldston, 2012) and claims of shortages or impending crises are met with skepticism (Teitelbaum, 2014). A number of recent US administrative initiatives (e.g. the National Center for Advancing Translational Sciences) are intended to improve the rate at which basic research becomes innovation (Sarewitz, 2013). However, in order to tell whether any of these efforts are successful, we must first agree upon some way to estimate R&D public benefit. The common goal of these efforts is to fund worthwhile research, but "worthwhile" is a subjective term that leads back to the fundamental question. How do we value R&D?

Trying to measure the value of R&D in some qualitative or quantitative fashion is a long-standing and contentious activity among science policy analysts. Ideally, our method of measurement reflects our values. For example, we perform economic analyses because we value

[15] According to the U.S. Office of Management and Budget, "Research and development (R&D) activities comprise creative work undertaken on a systematic basis in order to increase the stock of knowledge, including knowledge of man, culture and society, and the use of this stock of knowledge to devise new applications" (OMB, 2014).

economic efficiency. However, it is not always clear what all our values are or their priorities and if there are specific, practical ways to quantify either. Privately funded R&D undertaken by an organization such as a corporation can have highly uncertain outcomes with potential benefits and costs that are difficult to quantify. Government funded R&D is even more difficult as the benefits and costs to society can be even more diffuse. Yet, pursuit of these questions is a worthwhile endeavor. Government funded R&D constitutes less than one third of total R&D spending in the US, yet public funding of R&D does not merely augment or even displace private investment in R&D (Czarnitzki and Lopes-Bento, 2013). Rather, case studies suggest that public funding is critical to early-stage, high-risk research that the private sector is unwilling to fund; the result is a disproportionate contribution of government funding to subsequent technological innovation (Mazzucato, 2011). For example, while the private sector funds the majority of clinical trials in the U.S., public funding accounts for most pharmaceutical basic research (Stevens et al., 2011). Government involvement is widely accepted as necessary because the public benefit of research greatly exceeds its private value (Jones and Williams, 1998). The inability of private enterprise to profit from the full value of R&D has strengthened the case for direct government funding as well as subsidies for R&D in order to mitigate potential private underfunding.

This chapter reviews some common ways in which policymakers can assess the value of R&D: as a simple source of employment, as an engine for economic growth assessed through complex econometric methods, as knowledge and technology output, and even as an endeavor with social implications requiring assessment by ethical standards (Figure 1). The strengths and weaknesses of each approach are discussed and compared. Finally, some implications of the

various R&D valuation methods are presented to argue that there is no universally acceptable method for assessing the benefits of research.

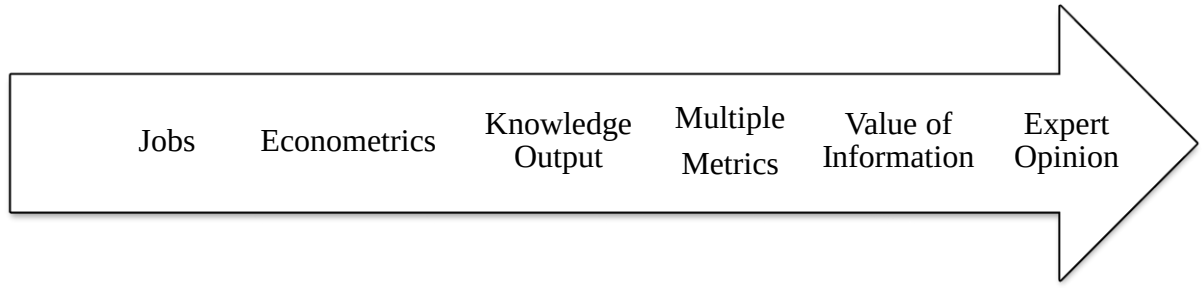

Figure 1: Ways of assessing the benefits of research ordered by increasing range of benefits that can be considered and also by the increasing uncertainty and subjectivity of estimates.

2.1 Methods of R&D Valuation

2.1.1 R&D as a jobs program

In a simple input-output model of R&D, spending on salaries, equipment, and facilities associated with research has an analogous output of research jobs, manufacturing jobs, construction jobs, etc., but the long-term impact of the actual research output is neglected (Lane, 2009). While this is a very simplistic way to view R&D, it is popular for two reasons. First, it is relatively quantifiable and predictable compared to other methods, to be discussed later, which assess the value of research output. For example, the STAR METRICS[16] program started in 2009 was intended to replace anecdotes and case studies of individual success stories with data that could be analyzed to inform the "science of science policy" (Largent and Lane, 2012). However, the first phase of STAR METRICS only attempted to measure job creation from federal spending (Lane and Bertuzzi, 2011; Weinberg et al., 2014). The jobs-only approach is also popular

[16] Science and Technology for America's Reinvestment – Measuring the EffecTs of Research on Innovation, Competitiveness, and Science

because job creation and retention is a primary interest of government policymakers. Even though the many long-term implications of R&D spending are frequently discussed by elected officials, the short-term impacts on jobs for their constituents are a more immediate concern. As US Representative George E. Brown Jr. (1999) said, "we hear how Congress needs to develop a rational priority-setting process for science and technology funding. Congress does have a rational priority-setting system. Unfortunately it is largely ZIP-code based: Anything close to my district or state is better than something farther away."

One outcome of viewing R&D in terms of immediate job creation is that any R&D program may be seen as a benefit to society because all R&D creates jobs. However, this neglects the economic efficiencies created by some R&D which, through innovation in automation and productivity improvements, eventually eliminates jobs. Likewise, when R&D funding is heavily influenced by the desire to retain science and engineering jobs in a particular electoral district, it can diminish the perceived legitimacy of an R&D program – for example, the cynical perception that the U.S. National Aeronautics and Space Administration (NASA) is a southern states jobs program (Berger, 2013; Clark, 2013).

2.1.2 Econometric valuation

Because the jobs perspective is very narrow, many analyses of R&D outcomes attempt to measure benefits using broader economic indicators. Econometric methods have attempted to measure the value of R&D by either a microeconomic or macroeconomic approach. The microeconomic approach attempts to estimate the direct and indirect benefits of a particular innovation by historical case study to generate a benefit-cost ratio (Hertzfeld, 1985). The historical case study approach has been very useful in assessing the value of particular technologies, but it is time- and resource-intensive. While the depth of insight available from a

case study approach is often underappreciated (Flyvbjerg, 2006), its detailed qualitative nature does not lend itself to decontextualized quantification.[17] Furthermore, actual benefit-cost ratios or rates of return for case studies tend to be valid only for the industry and the time period studied; as a result, they can be a poor source for forming generalizations about R&D activities (Griliches, 1979). Additionally, past innovations have often been based on serendipitous research (Ban, 2006) from unrelated fields with unforeseen benefits which further complicates attempts to directly correlate specific R&D directly with economic productivity (OTA, 1986).

The macroeconomic approach attempts to relate past R&D investments to a future economic indicator, such as gross domestic product (GDP). This approach is more useful for evaluating a broader range of research activities. In the macroeconomic approach, the value of research is the total output or productivity of an organization or economy attributed to past research investments. There are two primary weaknesses of this method. First, an organization's output is influenced by many factors other than prior research (e.g., macroeconomic conditions or political decisions). Due to the many potential confounding factors, it is difficult to demonstrate a causal relationship for any correlation of productivity or profit with R&D. Second, econometric methods are less effective when analyzing research intensive organizations (e.g., research universities) that have a poor delineation between research and output, poorly defined output, or output that is not directly valued in the market (Hertzfeld, 1985).

Unfortunately, some advocacy groups have issued reports that imply a simple direct relationship between scientific investment and economic growth that is not supported by

[17] However, this has not stopped big-data enthusiasts from trying. For example, keyword text-mining was performed on a 7,000-case-study audit of research impact in the United Kingdom. Ironically, the point of the audit was to add complementary context to a quantitative assessment (Van Noorden, 2015).

historical data (Lane, 2009). For example, Japan spends a higher percentage of GDP on R&D than most countries, but has not experienced the expected commensurate economic growth for the past two decades (Goldston, 2012). Likewise, in the year 2000, R&D spending in the US was more than ten times higher than China which resulted in the US contributing 27 percent of the global scientific literature and China contributing three percent; however, in the subsequent decade, China's economy expanded ten times faster than the US economy (Macilwain, 2013b).[18] While the exact relationship between R&D spending and economic growth remains unclear, the strong consensus is that R&D is beneficial. For that reason, the Bureau of Economic Analysis finally recognized the economic impact of R&D in 2013 by changing accounting procedures for GDP such that R&D is now considered an investment rather than an expense (Mcculla et al., 2013). This accounting change increased the GDP growth rate for the period 1995-2007 by about seven percent and the GDP in 2007 by almost $400 billion (Aizcorbe et al., 2009; Lee and Schmidt, 2010).

Three important factors have been noted when attempting an economic analysis of the value of R&D (Griliches, 1979):

1. The time lag between when research is conducted and when its results become useful to the organization defines the timeframe of the analysis. Depending on the research, the time lag from investment to implementation may take years or decades (Lane, 2009).
2. The rate at which older research becomes obsolete as it is replaced by newer technology and processes should be considered. Until recently, economic analyses often assumed an R&D depreciation rate of 15 percent (Li, 2012) - a gross simplification. The knowledge

[18] This has led to speculation that a robust research community is a result rather than a cause of national wealth (Macilwain, 2013b).

depreciation rate should be higher for a rapidly changing technology than for basic science research. For example, expertise in vacuum tubes became substantially less valuable after the invention of the transistor. Conversely, the value of a mathematical method commonly used in computer science might never decrease.

3. Since there is a "spillover" effect in research, the amount of similar R&D being conducted by competing organizations has an impact on the value of an organization's own research. This effect might be small for unique research that is unlikely to be used elsewhere. The spillover effect can vary in size and be positive or negative based on an organization's size, the current pace of innovation within the field, or the amount of competition within the field (Kafouros and Buckley, 2008). One reason for the complexity of the spillover effect is the influence of an organization's "absorptive capacity" or ability to make use of R&D output that was developed elsewhere (Cohen and Levinthal, 1989). Even without performing substantial research on its own, by keeping at least a minimum level of research capability, an organization can reap the benefits of the publicly available research output in its field. Using this theory, R&D is likely to yield much more for an organization than the direct output of its R&D activities. Presumably, this is a non-linear relationship. Increased R&D funding in a country with many well-trained researchers should cause a smaller relative increase in its absorptive capacity than one with a negligible R&D system.

In general, quantifying any of the above factors is easier for applied research than for basic research. Likewise, it is easier to quantify private benefit to a particular organization than social benefit (Hertzfeld, 1985). Another factor that prevents easy identification of the economic value of R&D is the general lack of data variation. Research funding rarely changes abruptly over time so it is difficult to measure the lag between research investments and results (Lach and

Schankerman, 1989). Taken collectively, these factors suggest that attempting an econometric assessment of R&D is indeed difficult.

A comprehensive review of attempts to determine the economic rate of return for R&D (Hall et al., 2009) finds a variety of methods being employed. The most common is growth accounting where R&D is assumed to produce all economic growth not accounted for by other inputs such as labor and capital. This unaccounted-for growth is often referred to by economists as the Solow residual (Solow, 1957). Hall et al. (2009) review 147 prior R&D studies that used either an individual business (59 studies), an industry (57 studies), or a region or country (31 studies) and found most rates of return ranging from 0 to 100 percent[19] – a wide interval which portrays the difficulty in quantifying the value of R&D. Not surprisingly, Hall et al. (2009) observe that R&D return is not constant across fields, countries, or time, so any estimates from one study should be used cautiously elsewhere. Likewise, it is important to distinguish technological progress from R&D. While technological progress may account for most of the Solow residual, a non-negligible amount of innovation occurs outside of formal public and private R&D programs.

2.1.3 Valuation by knowledge output

Given the difficulties of econometric valuation of R&D, economists have also explored alternative methods of assessing research investments that avoid the difficulty of monetizing research benefits and the private versus social benefit distinction. Instead of using monetary value, the assessment tool is academic publications in a particular field over a given time period, lagging the research investment in question, and weighted by the number of researchers working

[19] One outlier study conservatively (and perhaps humbly) estimated rates between -606 and 734 percent.

in the field (Adams, 1990). Despite the relative simplicity and ease of assessment compared to economic growth, there are still many difficulties with this approach. First, comparisons are complicated because scientific publication is not equally valued among all fields and organizations. Second, the method does not assess the relative value or visibility of individual publications although this issue is partially addressed by using the number of citations rather than the number of publications. However, citations are not a clear sign of quality research; citations are commonly made to reference methods, provide basic background information, or even criticize prior work (Werner, 2015). Furthermore, despite the personal advantage of publishing in prestigious journals, an ever increasing portion of research is being disseminated outside standard academic journals via the internet (Reich, 2013). It is unclear how non-traditional open source information sharing should be measured. Additional shortcomings are similar to traditional econometric approaches: what is the appropriate lag time for which to consider publications and what time window should be considered for counting publications based on the depreciation rate of scientific knowledge (Adams and Sveikauskas, 1993)?

A study of over 2800 academic and professional journals finds that the time it takes for half of an article's lifetime downloads (half-life) ranges from two to five years with health sciences having the shortest half-life and humanities and mathematics having the longest (Davis, 2013). Less than three percent of journals have average article half-lives of less than one year, while 17 percent are more than six years. In a separate model (Wang et al., 2013) of the ultimate number of citations for a particular article, influencing factors included: citations accruing faster to papers that already have many citations, a log-normal decay rate for future citations, and a general factor that accounts for the novelty and importance of a paper. However, the model requires 5 to 10 years of citation history to make projections and the difficulty of properly

calibrating the model limits its utility (Wang et al., 2014a, 2014b). Conversely, an extensive study (Ke et al., 2015) that looked at 22 million papers published over the timespan of a century in the natural and social sciences found that citations histories are heterogeneous and unpredictable. Some extreme papers, labeled "sleeping beauties" would accumulate very few citations for decades and then suddenly peak suggesting that using short-term citations as a metric for assessing research is problematic.

A similar non-monetary approach for measuring R&D benefits is to measure the number of patent citations in a particular field (Griliches, 1979; Jaffe et al., 1993). This method has the benefit of better assessing the practical value of R&D activities and capturing the technological innovation component of research that is likely to have high social benefit. However, this method also shares some of the drawbacks of the publication approach as well as a few unique drawbacks of its own. Griliches (1994) observes that a US productivity peak in the late 1960s was followed by a decline in patents granted in the early 1970s and that both events were preceded by a decline in the proportion of GDP devoted to industrial R&D spending in the mid-1960s. Whether productivity and patents followed a five-to-ten-year lag behind R&D spending is difficult to determine given that among other factors: the number of patents per research dollar also declined during the same time period, an energy crisis occurred during that time period, and other countries suffered similar productivity losses without the drop in R&D funding.

Likewise, fluctuations in patent generation may be due to the national patent office itself. For example, the 2011 Leahy-Smith America Invents Act, which took effect in 2013, changed the unique US "first to invent" patent system to a more standard "first to file" system. This makes comparisons before and after the new system more difficult. Similarly, stagnant or declining funding of a patent office could limit the throughput of the department or prevent it

from keeping up with growing patent application submissions (Griliches, 1994). This very phenomenon appears to have occurred in the US since the innovation boom of the Internet age (Wyatt, 2011). These various influences on the patent system can be partially accounted for by looking at a particular sub-field as a percentage of all patents (Popp, 2002), using patent applications rather than patents awarded (Ernst, 2001), or performing international patent count comparisons (de Rassenfosse and van Pottelsberghe de la Potterie, 2009). However, the increasing globalization of research efforts (Wagner, 2014) may limit international comparisons in the future as international research teams complicate the clear delineation of R&D spending and knowledge output within individual countries.

In the end, patents remain a very limited and non-representative measure of R&D benefits. There is poor correlation between patents and public benefit because most benefits come from a small subset of all patents and only about half of all patents are ever used and fewer are ever renewed (Scotchmer, 2004). Not all organizations patent their inventions at the same rate because the value of a patent is distinct from the value of the invention (Bessen, 2008). Nonetheless, patents are still widely used as a measure of R&D value for lack of a better alternative.

While economists view the measurement of R&D to be problematic, but possible, other academic traditions suggest that the problem is intractable or at least not quantifiable in any honest way. Feyerabend (2011) argues that a careful study of the history of science shows that the truth or usefulness of any particular scientific theory or line of research may not be appreciated for decades or even centuries. An extreme example is the theory proposed by the ancient Greek philosopher Parmenides of Elea (fifth century BCE) that all matter has the same fundamental nature. The theory was abandoned for over two thousand years before being revived

by particle physicists in the twentieth century (Feyerabend, 2011). Another more recent example is the theory of continental drift, first proposed in 1596 by Flemish cartographer Abraham Ortelius. The theory was revived in 1912 by meteorologist Alfred Wegener who championed the idea for two decades, but was met with ridicule and dismissal by a near-unanimous consensus of scientists. The idea was eventually incorporated into the theory of plate tectonics in the 1960s and is now a cornerstone of modern geoscience (Frisch et al., 2010).

Of course, economists are generally trying to measure the societal benefits of more tangible and immediate R&D. However, efforts to select an appropriate lag time for measuring practical R&D value will still likely miss truly important findings. There were decades between the development of quantum physics and technologies based on quantum mechanics (e.g., transistors, lasers, and magnetic resonance imaging (MRI)) and new technologies based on quantum theory are still in development (e.g., quantum computers). Even a field of research that has yet to yield useful results should not be dismissed as long as it still has intellectually inspirational value, for one never knows what is yet to transpire. Selecting a lag time by a cutoff function that is designed to capture most of the citations, patents, or economic growth based on past research is based on the questionable assumption that only the intended outcome of applied research is of interest. However, the history of technology suggests that secondary unintended discoveries, both good and bad, are very important.[20] Thus, selecting a time period for the evaluation of R&D investments may capture the intended primary applied research to some degree, but will likely miss the secondary serendipitous discoveries. More importantly, these

[20] In the pharmaceutical industry, drugs are commonly repurposed when they are found to treat a disease other than their intended target. For example, while aspirin has been used for centuries as an analgesic, only in the past few decades has it also been found to benefit cardiovascular disease treatment.

secondary factors are frequently large enough to substantially change or even reverse the primary assessment. Thus, any claim that an assessment is mostly complete or representative of the true social benefit is difficult to justify.

2.1.4 Valuation by multiple metrics

The various econometric and knowledge-output metrics discussed here appear to be poor measures of the social benefits of R&D. They are popular primarily because they make use of the available data, not because they are ideal. Furthermore, using a single metric or narrow set of metrics can lead to researchers changing their behaviors in potentially non-productive ways in order to improve their "scores" under those very metrics. These "reactive effects" of measurement are widely recognized in the social sciences and are difficult to avoid (Stone, 2002). For example, if patents become a preferred metric of research productivity, many researchers will knowingly generate patents that are of questionable licensing value in order to improve their likelihood of securing future funding. Likewise, the frequent practice of using the number of publications as a metric has led to academic complaints and jokes about "salami-slicing" research and the "least publishable unit." There are already indications that quantitative assessments of research output in the United Kingdom, Australia and New Zealand are having various unintended consequences such as pushing researchers away from high-risk basic research and towards more conventional, short-term, applied projects to improve their rankings (Owens, 2013; McGilvray, 2014). Likewise, increasing the pressure on academics to publish (which was already substantial) may account for some of the discovered intentional fraud (Labbé and Labbé, 2013) in journals and conference proceedings. While quantitative ranking is intended to act as an objective merit-based system, the competition may also be selecting undesirable

traits, such as self-promotion, rather than creativity and collaboration (Polka and Krukenberg, 2014).

Despite these critiques of individual approaches to R&D valuation, there has been some hope that using a family of complementary metrics would yield an improved estimate. For example, a combination of publication citations to capture basic research and patents to capture technology development would appear to be complementary measurements. Likewise, a multi-metric approach is intended to have fewer unintended consequences in the sense that it is more difficult for those being measured to game the system. At the urging of John Marburger (2005), then Director of the Office of Science and Technology Policy and presidential Science Advisor, the STAR METRICS program was created to measure the impact of US federally funded R&D using a multi-dimensional approach (Cragin et al., 2012). Phase II of the STAR METRICS program was intended to expand measurement to a comprehensive set of metrics. Suggested indicators were grouped into four domains and included: economic (e.g., number of patents, number of start-up companies, economic value of start-up companies over time, future employment of student researchers, and impacts on industry from research); workforce (e.g., the number of researchers employed); knowledge creation (e.g., publications and citations); and social outcomes (e.g., long-term health and environmental impacts) (Federal Demonstration Partnership, 2013). While the STAR METRICS proposed comprehensive approach avoids some of the limitations of individual metrics previously discussed, it is questionable how many of the proposed metrics could be measured in practice or how representative the final set of metrics would be. In particular, accounting for the total social benefit of R&D is a daunting task and separating the public benefits attributed to R&D from other factors is still problematic. In 2015, it was announced that the STAR METRICS program would end and researchers could perform

their own multi-metric analyses using an open access database, Federal RePORTER (Rockey, 2015).

A more successful program, the Innovation Union Scoreboard (IUS) (Hollanders and Es-Sadki, 2014) has been evaluating European Union member states since 2007. The IUS encompasses twenty-five R&D metrics including multiple indicators for: educational outcomes, scientific publications, patents, public R&D investments, private R&D investments, employment, and other economic indicators. As with similar programs, the IUS is by necessity restricted to indicators for which there are data. Unquantifiable benefits are necessarily missed. Even where data are available, it is difficult to describe and account for all the possible facets of social benefit.

Despite the difficulties of quantitatively valuing R&D, the era of big-data has inspired an entire alphabet soup of research assessment systems, none of which can be easily compared. Besides STAR METRICS and IUS, the United Kingdom has the Research Excellence Framework (REF) (cf. Atkinson, 2014) which replaced the previous Research Assessment Exercise (RAE), New Zealand has the Performance-Based Research Fund (PBRF), and Australia has the Excellence in Research (ERA) program to name a few (Sheil, 2014). Detractors have argued that these broad quantitative measurement tools are just as non-representative and easily gamed as the many popular, but widely derided college ranking schemes (Macilwain, 2013a). It has yet to be seen if any of these multi-metric systems will improve research or how (outside of their own definition) success will be determined. The rush to quantitative assessment is not universal. Since 2011, the Chinese Academy of Sciences has been moving away from an existing multi-metric national research ranking system that used 24 indicators to a qualitative system

based on peer review (Kun, 2015). The motivation is a desire to place emphasis on the real social value of research rather than on easily measured surrogates.

2.1.5 Value-of-Information analysis

When performing a cost benefit analysis of any research program, econometric methods can appear to be an obvious choice based on convenience and methodological familiarity. However, as discussed above, this is a difficult task even for research that has already been conducted. Estimating the value of future research is even more uncertain as it requires the questionable assumption that the future will be much like the past. This is a difficult assumption to defend because history suggests that the progress of technology is inconsistent and unpredictable. For example, computer technology has exceeded most predictions made in the twentieth century while nuclear fusion power plants have stubbornly remained a promise of the future.

For future research decisions, an alternative to the econometric or knowledge output approach is to use value-of-information (VOI) analysis where the value of the research is measured by estimating its expected value to a particular decision and weighing it against the cost of obtaining that information (Morgan et al., 1990; Pratt et al., 1995; Hammitt and Shlyakhtel, 1999; Fischhoff, 2000).[21] For example, knowing the transmissibility of a particular pathogen has value for public health officials in their decision of how to prepare for a potential future pandemic. This value can be measured in any agreeable units – money, lives saved,

[21] VOI literature often uses the term "expected value of perfect information" which is simply the difference between the value of the decision made with complete information and with existing information. Restated, this is the value of removing uncertainty from the decision process.

response time, etc.[22] The primary strength of this approach is that it deals directly with the value of R&D to the decision maker (Claxton and Sculpher, 2006). By comparison, high quality research, as measured by knowledge output methods, has no clear correlation to societal benefit – only an assumed link (Bornmann, 2012).

Likewise, because VOI is a forward-looking predictive method of valuation rather than a backward-looking reflective method, it sidesteps the issue of making direct comparisons between past and future research. VOI does not require the assumption that the benefits of past research will be equivalent to similar future research. However, it does not entirely escape the past. VOI must place a value on proposed research and that valuation is at least implicitly informed by how valuable past research was to decision-makers.

Another strength is that VOI analysis is a more theoretically complete and consistent method of R&D valuation. Performing a cost benefit analysis using a family of economic, knowledge, and social metrics can use collected data, but that data will generally be an incomplete measure of the total value of R&D and will often consist of proxies for the characteristics that we would prefer to measure. Conversely, a VOI approach can place a value on factors that are difficult to monetize: aesthetic, intellectual, or even the cultural significance of a scientific discovery[23]. Thus, VOI is comprehensive in the sense that any recognized benefit can be included in the analysis. The thoroughness of the VOI approach comes at the price of subjective estimates and value judgments – especially for estimating the value of highly uncertain basic research. That is, VOI is a productive decision tool when one can reasonably

[22] Important ethical assumptions regarding what is worth optimizing are embedded within the choice of units (cf. Robberstad, 2005)

[23] Consider the social implications of world-view changing theories such as evolution or the Big Bang.

estimate the value that obtaining information will have to a decision. For that reason, VOI is often applied to business, engineering, and applied science decisions (Keisler et al., 2013). For example, VOI could more or less objectively estimate whether a particular medical test has value to decisions about patient treatment or patient outcomes. However, it is harder to justify using VOI, as suggested by Fischhoff (2000), to determine the value of more basic research, such as the discovery of the Higgs boson. VOI is most subjective when it is measuring the most subjective things. Like all decision methods, it cannot spin straw into gold.

The thoroughness of the VOI approach also makes analysis difficult due to the richer array of potential social benefits that might be considered. While VOI is simple in concept, it can be quite complex in practice (e.g., Yokota and Thompson, 2004). For this reason, the VOI approach is often used in conjunction with an influence diagram – a visual representation of a decision process that represents variables as nodes and interactions between variables as arrows (Clemen, 1991; Howard and Matheson, 2005). The influence diagram serves as an aid to elucidate and organize the often complex interaction of factors that can affect the value of basic research.

For example, organizing the various ways in which R&D can be valued in an influence diagram, Figure 2, shows how they are related. Each form of valuation is represented as a node with arrows indicating if the method informs another method. For example, job creation is often concurrent with economic growth (but not always) so we would expect these two R&D valuation methods to be closely related. Likewise, both jobs and economic growth can be used in a multiple metrics approach or in an expert opinion approach. Knowledge output, in the form of citations and patents, can also be used in a multiple metrics approach and is similar to the VOI approach in that both are non-monetary and can more easily characterize the value of basic

research with no immediate practical applications. Expert opinion, discussed in the next section, is the most comprehensive approach in that it can make use of all the other methods of valuation. However, in practice, expert opinion and managerial review can range from superficial to comprehensive.

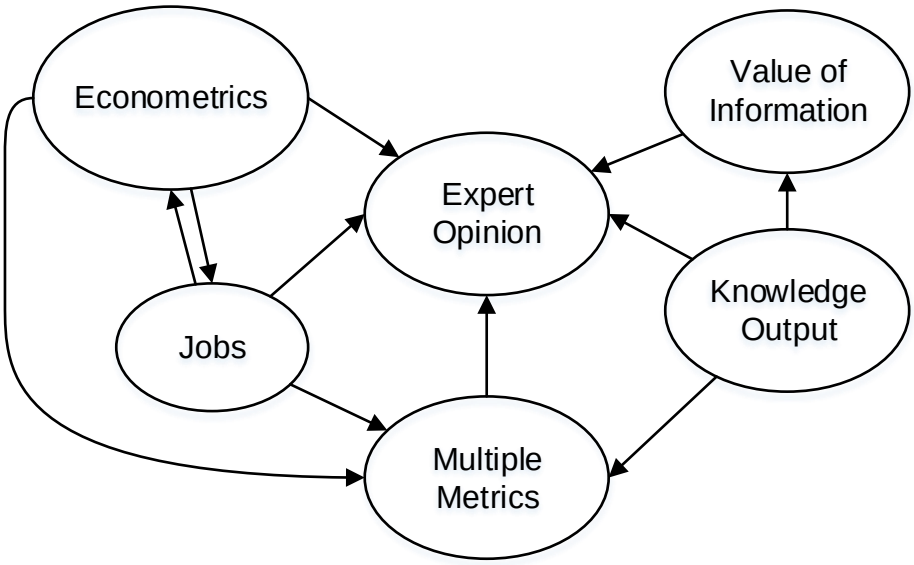

Figure 2: Ways of valuing R&D

Although less formal than the VOI approach, Sarewitz and Pielke Jr. (2007) outline a similar process by which the supply of research and the demand for research are separately assessed and compared. This is done by detailing the information required by policymakers for policy formulation (the demand) through regular communication between researchers and policymakers in the form of workshops, surveys, interviews and committees. Using the same process, regular assessments are made regarding whether the R&D output is actually being used. Rather than placing a common value on the information, the intent is only to re-align R&D priorities such that funded research is focused on producing social benefit. This process is much less subjective than VOI in the sense that it does not attempt to quantitatively compare R&D

programs. However, it is also time consuming in that it seeks input from all stakeholder groups, with their own disparate research agendas, and can be difficult to achieve when contentious issues preclude a consensus on the "demand" for science. Furthermore, it is difficult to predict what research will actually yield the most social benefit; focusing only on immediate applied research would miss important basic research that eventually yields important technology.

2.1.5.1 Extending the value of information approach for R&D assessment

In standard VOI literature, benefit is derived from additional perfect or imperfect knowledge and it is assumed that the value of information can never be negative because a decision maker can always choose to ignore low value information (Blackwell, 1951; Marschak, 1954, 1959). However, experiments suggest that decision makers are often unable to ignore unhelpful information once it is known resulting from a "curse of knowledge"[24] (Camerer et al., 1989). Furthermore, decision makers are often unaware when information is unhelpful based on their willingness to pay for unhelpful information (Loewenstein et al., 2003). This questions the basic assumption that the value of information is never negative because it can be ignored without cost.

We can extend this concept of negative value of information to include R&D that may yield knowledge that has potential for public harm as well as public benefit. The idea that publicly funded R&D should serve society is a reasonable and widespread belief (Sarewitz and Pielke Jr, 2007; Sarewitz, 2013). However, there is much less discussion regarding the potential public harm of R&D. Some notable exceptions are the discussions surrounding "dual-use"

[24] For example, the curse of knowledge is frequently experienced in teaching. It is extremely difficult to imagine one's own state of mind before a concept was understood. This leads to teachers often overestimating the clarity of their instruction and the comprehension in their students (Weiman, 2007).

research (i.e., research that also has obvious use by military, terrorists, or criminals). Without the negative VOI concept, the worst R&D is only wasted effort. With the idea of negative VOI, some R&D programs may yield information we might prefer not to know. This introduces the potential to account for important ethical considerations into the R&D valuation process.

Additionally, there is research that yields harm to the public because it is erroneous. It can be difficult to assess the benefits of research that has implications for public policy because some research may be of low quality and may inappropriately influence public behavior (Boyd, 2013). For example, the 1998 Lancet paper linking the MMR vaccine with autism, although later discredited and retracted, fueled a wave of suspicion regarding the safety of childhood vaccination; subsequent outbreaks of preventable diseases and multiple fatalities have been seen in the U.S. that appear to be linked to this anti-vaccination movement (Gross, 2009).

2.1.6 Qualitative assessment by expert opinion

A 1986 Office of Technology Assessment (OTA) report (OTA, 1986) reviewed a variety of quantitative methods for determining the value of R&D and the prevalence of such methods in industry and government. The OTA found that the majority of government and industrial managers perceived quantitative methods to be simplistic and misleading and preferred traditional peer review as the best method of assessing the value of research. "Any effort to substitute formalistic quantitative models for the judgment of mature, experienced managers can reduce rather than improve the quality of R&D decision making. The resistance of R&D managers to the use of quantitative decision tools is, to some degree, a rational response to the complexity and uncertainty of the process" (OTA, 1986).

Given the issues described above with various quantitative methods, it is unsurprising that expert opinion is still the standard in estimating the value of R&D. However, qualitative

peer review is also problematic, especially for evaluations that require the comparison of expert opinions across a wide range of research specialties such as during R&D funding discussions. One problem encountered with using subject experts in peer review is the inherent appearance of conflict of interest. Specialists often have a financial or emotional stake in their field of expertise which leads to bias. The value people attach to an activity is increased by the effort they expend. This phenomenon, referred to as "effort justification" (Festinger, 1957) or the "IKEA effect" (Norton et al., 2012) can lead experts to unintentionally overestimate the value of the research with which they have been most involved. Even the appearance of conflict between what is in the best interest for the general public and what is in the best interest for the experts themselves decreases credibility and can reframe budgeting discussions as special interest political lobbying.

Making science policy in this fashion can be undesirable as the political process may not lead to an optimal research agenda as defined by “well-ordered science” (Kitcher, 2001) (c.f. Keren, 2013). For example, the 2014 NASA budget request included funding for a project that would ultimately capture a small asteroid from deep space and put it in a near-earth orbit. However, many planetary scientists expressed doubts that the purported benefits of the project, asteroid mining and planetary defense research, had any reality or utility beyond attracting attention from politicians (Kerr, 2013).

One way to partially compensate for potential expert bias is to actively seek competing views. Kitcher (2001) recommends an “enlightened democracy” where well-informed individuals widely representing society set science research agendas. This ideal is set as a middle ground between a “vulgar democracy”, where science suffers from the “tyranny of the ignorant”, and the existing system where a struggle for control over the research agenda is waged between scientists (“internal elitism”) and a privileged group of research funders (“external elitism”).

Some influences on science research agenda, such as the focused lobbying by well-informed advocates for rare diseases (Reardon, 2014a), defy this idealized distinction between a scientific elite and the uninformed public. Nonetheless, the struggle to maintain a balanced and representative set of R&D policymakers is real. An example of this struggle, US President Eisenhower created the President's Science Advisory Committee (PSAC) to provide cautious science policy analysis during the American pro-science panic that occurred after the launch of Sputnik in October, 1957. PSAC's criticism of President Kennedy's manned space program and President Johnson's and President Nixon's military programs led to PSAC's ultimate demise (Wang, 2008) which suggests that balanced counsel from technological optimists and skeptics can be difficult to maintain.

Furthermore, expert opinion is time-consuming (Arns, 2014),[25] expensive (Mervis, 2014a) and has a poor historical record of predicting R&D outcomes. For example, the current National Science Foundation (NSF) method of using peer review panels to assess research proposals based on significance of goals, feasibility, the investigator's track record, etc., may not be capable of reliably making even relative rankings of future research impact. For a study of 41 NSF projects funded a decade prior, using the number of publications and citations as a measure of research success, panelists' predictions of future research success were found to have no significant correlation with actual research success (Scheiner and Bouchie, 2013). Similar larger studies of the NIH peer review processes have found that there is no appreciable difference between high and low ranked grant proposals in their eventual number of publications per grant, number of citations adjusted for grant size, or time to publication (Mervis, 2014c). However, a

[25] There is also a concern that the peer-review process, as it currently operates, is being overwhelmed by data – an explosion of journals with increasingly data-rich papers (Siebert et al., 2015).

more recent study of 137,215 NIH grants awarded between 1980 and 2008 found the highest rated grant proposals yielded the most publications, citations, and patents (Li and Agha, 2015). For example, a proposal with a review score one standard deviation above another generated 8% more publications on average. Critics have questioned the cause of this correlation considering that journal publications are also based on peer-review; any correlation may be indicative of measuring the same reputational system.

The journal peer-review system was the subject of another study that followed the publication history of 1,008 submissions to three top medical journals (Siler et al., 2014). Of the 808 manuscripts that were eventually published, the lowest rated submissions tended to receive the least eventual citations. However, the top fourteen papers were all rejected at least once which suggests that the most innovative, high-impact work is often unappreciated by the peer-review process. [26]

Additionally, the history of science is littered with examples of substantial scientific research that went unappreciated by expert opinion to an extent that is almost comical in hindsight. For example, Lynn Margulis' paper proposing that mitochondria and chloroplasts in eukaryotic cells evolved from bacteria (Sagan, 1967) was originally rejected by over a dozen journals (Lake, 2011). Over a decade later, DNA evidence confirmed the theory and Margulis was eventually elected to the National Academy of Sciences and given various awards including the National Medal of Science.

In another example, Dan Shechtman, the first to identify the existence of quasicrystals (Shechtman et al., 1984), needed two years to get his paper published, was asked to leave a

[26] While most rejections were desk rejections (i.e., rejections by the journal editors), this is, in practice, part of the peer-review process. This conservatism may be a sign of "normal science" (Kuhn, 1962) in action.

research group, and was met with considerable opposition from the scientific community including Nobel laureate Linus Pauling (Jha, 2013). Ironically, for this work, Shechtman was eventually awarded the Nobel Prize in Chemistry in 2011. Perhaps for these reasons, there has been discussion at the NIH to skip the standard peer review which appears to be less effective at selecting and promoting innovative high-impact research than funding based solely on a researcher's reputation (i.e., funding people rather than projects)[27] (Kaiser, 2014b).

2.2 Implications of Multiple R&D Valuation Methods

The optimal funding levels of R&D and the types of research that should be publicly funded are perhaps unresolvable questions. However, a comparison of the various R&D valuation methods provides some insight into science policy. For example, Figure 1 shows the various methods ordered from the narrowest to the broadest conception of benefits. This is also the same ordering of the most objective to the most subjective approach. R&D policy decisions based on job creation or economic growth are the most data-driven, while VOI and expert opinion approaches require more personal judgment. The choice of approach requires a compromise between being objective and being comprehensive.

[27] The purported weakness of this approach is that it may make it harder for younger researchers to obtain funding. This is a concern for the biomedical science community given that the median age of first NIH grant increased by six years between 1980 and 2014 (Kaiser, 2014a). However, during the same time the median age of the total U.S. population increased by over seven years. This demographic trend coupled with the widely discussed glut of biomedical researchers (Meredith, 2012) relative to funding opportunities likely explains the increase in age of first NIH grants. There is also no convincing evidence that this is hurting science. The association of youth with innovation appears to be a cultural myth (Agan, 2013). Over the past century, the average age of the Nobel Prize winner has increased (Jones, 2010) likely indicating that scientists are both living longer (the Nobel is not awarded posthumously) and major innovative ideas in science increasingly require more training. Even in the supposed innovation youth paradise of high-tech startups, twice as many founding entrepreneurs are over 50 than under 25 (Wadhwa et al., 2008). Without further justification, attempting to lower the age of the first award of research grants is simply ageism (cf. Wray, 2014).

One attempt to bridge the differences between objective and more subjective ways of R&D valuation calls for expanded econometrics. The limitations of GDP and similar measures of economic growth are widely known and even expanded measures (e.g., Muller, 2014) that use monetary valuation schemes are problematic (Adams, 2014). Alternatives, such as the "Genuine Progress Indicator" or "World Values Survey," have been proposed to account for factors important to well-being that are ignored by GDP (Costanza et al., 2014). Expanded econometric approaches can incorporate factors relevant to R&D valuation, but this requires general acceptance of a new more complex system of metrics and agreement that the new system is actually measuring the desired characteristic (cf. Krueger and Stone, 2014). Ultimately, any attempts to capture the more intangible benefits of research and broaden a benefits assessment necessarily increases its uncertainty and subjectivity.

One basic observation regarding the different approaches is that some types of research are more amenable to a particular form of assessment. This suggests that scientists and engineers involved in "blue skies" basic research, that has only job creation as an immediate quantifiable benefit, should avoid getting locked into econometric valuation debates. When science is treated as a mere economic engine, the weaknesses rather than the strengths of curiosity-driven research are emphasized resulting in weak justifications such as the previously mentioned NASA proposal to capture an asteroid. Rather, curiosity science should be honestly argued on intellectual, aesthetic, and even moral grounds if support from the general public is expected.

Staying with examples from NASA, in 1970, Ernst Stuhlinger, a scientist and NASA administrator, responded to a letter from Sister Mary Jucunda questioning the expenditure of billions of dollars for manned space flight given the plight of starving children in Africa. Stuhlinger's response (Usher, 2013) is an eloquent defense of the value of R&D generally, but a

rather weak defense of space exploration programs based on several proposed practical benefits (none of which are actually dependent on manned space flight): satellite data to improve agricultural output, encouraging science careers, increasing international cooperation, and serving as a more benign outlet for Cold War competition. However, Stuhlinger wisely closes the letter with a reference to an enclosed photograph of the Earth from the Moon and hints at its worldview changing implications. The 1968 picture now referred to as "Earthrise" was later described by nature photographer Galen Rowell as "the most influential environmental photograph ever taken" (Henry and Taylor, 2009).

2.3 From R&D Valuation to Allocation

Most of the discussion so far has focused on assessing the benefits of past research. Because R&D explores the unknown, estimating the benefits of future research is, by its nature, extremely difficult. Reviews of academic studies by economists have found that there is no consistent set of criteria that will predict whether a particular R&D project will succeed. The list of contributing factors is extensive and there is even disagreement among studies regarding the magnitude and direction of influence of each contributing factor (Balachandra and Friar, 1997). Thus, it would appear to be ill-advised to set R&D budgets based on the perceived viability of individual projects. This does not imply that policymakers must treat and fund all R&D requests equally. Rather, they should fund a broad range of individual projects based on best current science and the intended goals for funding. Given the serendipity inherent in R&D and the potential for "black swans," the optimal investment strategy appears to be relatively stable investment over time (Press, 2013). [28]

[28] Stability has an additional benefit. The rapid doubling and subsequent stagnation of the NIH budget between 1998 and 2003 created a painful market instability among biomedical scientists

It is also unwise to abandon basic research in favor of seemingly more predictable short-term applied R&D. How could one have predicted that the germ theory of disease, widely accepted by the end of the nineteenth century, would be the impetus for drinking water disinfection techniques responsible for much of the increase in average life expectancy in the twentieth century? While our inability to predict the future may seem obvious, predicting the future is a natural, if unintended, assumption of many attempts to improve societal benefit of publicly funded R&D. For example, the US High Quality Research Act proposed in 2013 attempted to guarantee economic benefit from R&D by requiring the NSF Director to certify that any funded research would be ground-breaking and "advance the national health, prosperity, or welfare" (Lymn and Houston, 2013). Likewise, the bill required certification that any funded proposal "is the finest quality, is ground breaking, and answers questions or solves problems that are of utmost importance to society at large." The failed bill would have been less disquieting if it were the only inelegant attempt by Congress to "improve" science research or if it had not been proposed by the head of the House Committee on Science, Space and Technology. While well-intentioned, predicting the outcome of a specific research proposal is logically not possible; otherwise, there would be no reason to perform the research in the first place. In this limited sense, the benefits of R&D are unmeasurable.

2.3.1 Ethics of allocation

Although we cannot predict the future, this does not mean that no assessment is possible and no setting of research priorities is desirable. The general public and science policymakers clearly do have R&D priorities. The argument that the 2013 discovery of the Higgs boson will

---

that may have ultimately dissuaded potential future scientists from careers in research (Freeman and Van Reenen, 2009; Shen, 2014).

someday be useful outside the field of particle physics (Prewitt, 2013) is likely true, but the public does not value all basic research equally. According to a directory[29] of over 7,000 charities operating in the US, in the beginning of 2014 there were over 100 charities that included cancer research as part of their mission, but none for particle physics. While the intellectual pleasures of verifying the Standard Model of particle physics are real, medical science, with its more immediate application to quality of life issues, attracts considerably more public attention. This precise allocation issue was recognized fifty years ago by philosopher Stephen Toulmin who wrote that "the choice between particle physics and cancer research becomes a decision whether to allocate more funds (a) to the patronage of the intellect or (b) to improving the nation's health. This is not a technical choice, but a political one" (Toulmin, 1964). My purpose here is not to argue over whether medical research is more worthy than particle physics. Rather, it is to highlight how different methods of valuing R&D have ethical and pragmatic dimensions that affect science policy. A jobs-only valuation might prefer funding particle physics research for the immediate construction and engineering jobs it supports, an econometric approach might prefer medical research based on historical growth rates in the pharmaceutical sector, and a knowledge output approach might be ambivalent between the two options.

Clearly, even with their explicit consideration, the expression of public values[30] in science policy is not assured in the near term. For example, if a nation chose to scale back on curiosity science, it is not clear that displaced scientists and engineers would necessarily start

[29] http://www.charitynavigator.org/

[30] Bozeman (2007) defined "public values" as the ethical consensus of society on what constitutes the rights, freedoms, and duties of individuals, organizations, and society. His definition also acknowledges that public values are not necessarily fixed, monolithic, or entirely compatible with each other (e.g., valuing both liberty and security).

working on applied projects that would more directly minimize human suffering. Scientists and engineers are not fungible commodities nor are they devoid of personal preferences. There is no reason to assume that R&D is a zero sum game. Likewise, public R&D funding is generally small compared to many other government expenditures which may have more or less societal benefit. For example, in 2014, the NSF budget was $7.6 billion and the NIH budget was $30.2 billion. In the same year, according to the US Office of Management and Budget, spending on national defense was $620.5 billion.

2.4 Conclusion

William H. Press (2013), then president of the American Association for the Advancement of Science, stated that "[a] skeptical and stressed Congress is entitled to wonder whether scientists are the geese that lay golden eggs or just another group of pigs at the trough." This is a return to questioning the social value of science (Bernal, 1939) – an attitude about US science policy that fell out of favor for several decades when Vannevar Bush (1945) rather successfully argued that science funding should be insulated from the political process. R&D allocation decisions have always been predicated on an expectation of societal benefit (Sarewitz and Pielke Jr, 2007), and the standard belief has been that R&D funding directly translates into knowledge, innovation, services, and jobs (e.g., Leshner, 2013; Porter, 2014).

When attempting to value R&D, it is important to consider the various approaches to assessing benefits and how they relate to each other. A federal report on the "science of science policy" (NSTC, 2008) categorized three broad methods by which research benefits and effectiveness could be evaluated: peer review, output metrics, and case studies. The distinction among methods was framed as quantitative versus qualitative. However, as illustrated in Figure 1, there are (at least) six distinctive common ways to assess the benefits of research.

Furthermore, there is an unavoidable tradeoff when using any of these approaches. The most "objective" methods use the narrowest and most tangible conception of benefits. In order to include more types of benefits requires the use of imperfect surrogates (e.g., counting registered patents rather than the utility of new technologies). Using the broadest conception of benefits requires considerable use of subjective estimates and the incorporation of ethical value-judgments. That is, one can use these various methods to obtain answers that are either precise and incomplete, or comprehensive and subjective.

In summary, the consensus among economists is that R&D activities are valuable (Salter and Martin, 2001), but difficult to quantify. Given the uncertainty, it is unwise to only use traditional econometric techniques to assess the benefits of R&D because they constitute only a portion of the total benefits of research. Other benefits, such as intellectual, social, and aesthetic value are poorly measured or ignored by econometric analyses.

Each of the previously discussed methods of assessing the value of publicly funded R&D has its own particular utility and shortcomings. A jobs-only perspective is immediate and understandable, but simplistic. Econometric and knowledge output metrics are more comprehensive, but they are also more data-intensive and can only attempt to project the value of future R&D based on similar past R&D. A value of information approach is a more comprehensive method, capable of valuing even the most intangible social benefits, but it is also difficult to perform and suffers from even more subjectivity than data-driven methods. Finally, the expert opinion and managerial review approach has the benefit of familiarity and the potential to incorporate all the other methods of R&D valuation, but due to the vagaries of the political process, has an unconvincing historical record of wisely allocating public R&D funding.

While only the VOI and expert opinion methods are capable of including ethical issues related to R&D funding, this does not imply that ethical considerations are always taken into account when using VOI or traditional expert opinion methods.

A primary intent of this discussion of R&D valuation is to help policymakers understand that there is no objective *and* comprehensive method for measuring the benefits of research. However, while the various methods are imperfect, they are reasonable to deploy when used with their limitations in mind. Failure to consider the characteristics of each assessment approach risks letting the method shape the goal of the assessment – the reverse of what constitutes good policy making.

3 VALUES IN RISK ASSESSMENT

The field of risk analysis, which emerged in the latter half of the twentieth century (cf. Otway, 1987), can be categorized into three main components: risk assessment, risk management, and risk communication (Paté-Cornell and Cox, 2014). Risk management (i.e., how to respond to risk) is widely recognized to include the most value judgments (Aven and Renn, 2010; Aven and Zio, 2014; Doorn, 2015). Conversely, there is a general assumption that good risk assessments are relatively free of individual biases and value judgments (Hansson, 2010b; cf. Hansson and Aven, 2014). However, despite quantitative rigor, the outcomes of risk assessments typically reflect the many value judgments implicit in the assumptions of the analysis – even the supposedly value-free components of a risk assessment (Ruckelshaus, 1984; Thompson, 1985; MacLean, 1986; Shrader-Frechette, 1986, 1991; Cranor, 1997; Hansson and Rudén, 2006; Duckett et al., 2015). Such misconceptions of objectivity in risk assessments raise the concern that unrecognized but pervasive value judgments will transform seemingly value-neutral risk assessments into unintentional stealth policy advocacy (Pielke Jr, 2007; Calow, 2014). This is particularly troublesome given an increasing trend, in the view of many, of intentional politicization in science (Macilwain, 2014).

In their description of the "ten commandments" of good policy analysis, Morgen et al. (1990) require the identification of all significant assumptions. This may be harder than most analysts realize, as a systematic discussion of the many value assumptions inherent in risk analysis has yet to be completed (Hansson, 2009a). Another "commandment" of good policy analysis (Morgan et al., 1990) is to perform a sensitivity analysis where input values and assumptions are varied to see the effect on the output. Clearly, a full sensitivity analysis is not possible if many of the assumptions go unrecognized. Likewise, delineating the assumptions may

lead to insight by the analyst on how to minimize unnecessary assumptions and appropriately account for the remaining assumptions in a more transparent manner.

Rather than attempt the Sisyphean task of an exhaustive accounting of all possible value assumptions, the intent here is to organize and discuss a number of the most common and contentious value judgments for illustrative purposes. The primary goal is to help guide subject experts attempting a risk assessment and policymakers using the results. The field of risk analysis is broadly multidisciplinary (Aven and Zio, 2014) with contributions coming from many academic disciplines across the sciences and humanities. Creating a general roadmap of value assumptions benefits those accustomed to working within the practices of one academic tradition. In addition to providing a tool for analysts, this chapter tries to make a convincing argument that all risk assessment involves unavoidable and important value assumptions that, if ignored, decrease the credibility of a risk assessment and its utility in informing public policy.

3.1 Categorizing Risk Assumptions

Value judgments in risk assessment can be categorized in multiple ways. One popular distinction is between two general classes of values: epistemic and non-epistemic (e.g., ethical or aesthetic values) (Rudner, 1953; Jeffrey, 1956; Colglazier et al., 1987; Reaven, 1989a; Rooney, 1992; Douglas, 2000; Hansson, 2012). Normative value assumptions in an assessment presuppose what ought to be and ethical distinctions can be made regarding the acceptability of risks (Slovic, 1987; Gillette and Krier, 1990; Cranor, 2009). Epistemic value assumptions presuppose what can be known and how best to know it. Epistemological assumptions, while within the purview of the philosophy of science, are nonetheless values (Hempel, 1960; Feleppa,

1981) in that they embody judgments of relative merit.[31] While the epistemic/ethical distinction has philosophical importance, it is less critical for risk analysis practitioners since, as discussed later, some assumptions can be justified by either epistemic or ethical reasons. Rather, organizing value assumptions by where they arise in the analysis process (e.g., methodological choices, collection and treatment of data, etc.) is more instructive to risk analysts. Momentarily ignoring that the process is iterative, the following is a discussion of the value assumptions that arise in each step of a risk assessment.

3.2 Selection of Topic

Perhaps the first value judgment made in any risk assessment is the choice of topic if it has not already been predetermined for the analyst. Similarly, when tasked with evaluating risk within a large pool of potential hazards, a screening mechanism or set of criteria is needed (e.g., MacGillivray, 2014) which involves some epistemic and/or ethical value judgments to prioritize efforts. Assessments are generally performed for situations that are perceived to be dangerous rather than benign. Thus, risk perception gives rise to important judgments on topic choice.[32]

Like many everyday activities, human risk assessment employs heuristics (i.e., mental shortcuts) that lead to biases[33] (Tversky and Kahneman, 1974). For example, technological hazards are perceived as more controllable than natural hazards, but also more dangerous and

[31] Epistemic value judgments (often uncontroversial) exist even at the most fundamental level. Scientific inductive reasoning, typically in the form of hypothesis construction, requires some basic beliefs about what is knowable, such as: correlations being potential indicators of cause and effect or that the world is predictable enough for that observable phenomena to be generalized into theory or law.

[32] Risk perception is frequently discussed in terms of how the public reacts to hazards, but less often in terms of how analysts and experts are affected.

[33] The term "bias" is widely used in risk perception literature, but the negative connotation of "bias" may obscure some of the positive qualities of heuristics to be discussed later. Montibeller and von Winterfeldt (2015) provide a detailed list of the many biases in risk analysis along with recommended "debiasing" techniques.

more likely (Xie et al., 2011). If the public perceive anthropogenic risks to be more threatening than natural risks of the same magnitude (Baum et al., 1983; Brun, 1992; Williams and Hammitt, 2001), the result should be a general propensity to make technological hazards the subject of risk assessments while overlooking equally or more risky natural hazards.[34]

Likewise, the timing of a risk assessment is often influenced by how easily we can imagine a risk, i.e., the "availability heuristic" (Tversky and Kahneman, 1973). Funding for risk assessment appears when a threat is fresh in the minds of a funding organization rather than when the risk may be greatest (assuming actual risk is even knowable). Thus, assessments tend to follow recent events such as a natural disaster, attack, or even a well-timed visit by an advocate lobbying for a particular issue. For example, newsworthy meteorite impacts or near-misses periodically renew interest in assessing the risk of major impacts from near earth objects (e.g., Broad, 2013). However, given the rarity of catastrophic impacts compared to other serious natural and anthropogenic threats, such risks tends to feel theoretical and unreal. As a result, funding for surveillance and mitigation research is inconsistent (Watson, 2015). The value judgments involved are both epistemic and ethical. Epistemic in the sense that there is an assumption that the hazard is now more "real" and knowable after an event and ethical in the sense that the hazard is now viewed as more worthy of analysis than other hazards.

3.3 Defining System Boundaries

For any assessment, the boundaries of the analysis must be selected. This includes time frame, spatial scale, and relevant populations. While the boundaries are primarily determined by the topic, some pragmatic epistemic value judgments are involved. First, the analyst must believe

[34] This phenomenon may be partially responsible for inadequate natural disaster preparedness. A related question is whether many nations allocate too many resources to traditional national defense compared to natural disaster mitigation.

that it is possible to distinguish what factors can be left out of an assessment without appreciably affecting the results. Likewise, the analyst must believe that meaningful boundaries can be defined for a system.[35] Ideally, expanding an assessment to add secondary factors will make incrementally smaller changes to the results thereby showing the assessment to be convergent (Wilson and Crouch, 2001), but this is not always the case (cf. Reaven, 1984, 1985a, 1985b).

The analyst's educational background and experience may also have an important influence on choices of time scale. Moser et al. (2012) argues that training in the natural sciences tend to increase the perception of predictability for long-range risk assessments. That is, in a survey, scientists, such as physicists and geologists, were more likely to believe that performing a risk assessment of a long-term nuclear waste repository was a reasonable task, while those in the social sciences and humanities, such as anthropologists and philosophers, were considerably more skeptical.[36]

3.4 Method and Units of Assessment

3.4.1 Unit of assessment

A standard assessment in modern[37] risk analysis entails maximizing expected utility via probabilistic risk analysis (Hermansson and Hansson, 2007). An analyst using this approach must decide whether the risks will be expressed in monetary units or another unit, such as

[35] The idea of inherent interconnectedness has long existed in the philosophical traditions of Buddhism and Taoism, but was an uncommon idea in western thinking until the 20th century when the emerging field of ecology led to sentiments such as, "When we try to pick out anything by itself, we find it hitched to everything else in the Universe," (Muir, 1911) and, "Everything is connected to everything else" (Commoner, 1971).

[36] Conversely, one could also argue that thoughtful scientists make the best skeptics in the sense that their training helps them to perceive flaws in analyses rather than treat technical complexity with deference.

[37] The 1975 Reactor Safety Study, WASH-1400, a.k.a. the Rasmussen Report, was an influential work in quantitative risk methodology and perhaps best marks the start of modern formal risk analysis (cf. Reaven, 1988).

injuries/year or quality adjusted life years. The units are important in that a different unit can change the perception of the risk (Wilson and Crouch, 2001). Depending on the situation, a relative risk can appear to be much larger than an absolute risk or vice versa. For example, a WHO (2013) report regarding the 2011 Fukushima nuclear accident estimated that the lifetime risk of thyroid cancer for nearby infant females increased from about 0.75 to 1.25 percent. Depending on the desired effect, media reported the findings as a 70 percent relative rate increase or a 0.5 percent absolute rate increase.

Selection of a unit involves an epistemic value judgment regarding the measurability of a characteristic; presumably, an analyst would not pick an unmeasurable unit. The choice of units is also accompanied by important, and often unwitting, ethical assumptions. For example, the units of lives saved, life years gained (LYs), quality adjusted life years (QALYs), and disability adjusted life years (DALYs) all differ in the populations that preferentially benefit from the use of the unit (cf. Robberstad, 2005). Likewise, the disparate nature of various risks and benefits often requires the use of assumption-laden conversion factors and equivalencies to make comparisons.

3.4.2 Value of life

The primary benefit of a monetary unit is the ability to integrate the results into a larger economic analysis. However, the conversion to monetary units requires some important value assumptions. The most controversial is the frequent need to monetize the value of life. Attempts to quantify the value of life often use willingness-to-pay measurements or expert opinion (e.g., Roman et al., 2012). A defense of value of life estimates is that the number does not represent the worth of a life, but rather the rational amount society should be willing to pay to save a life. However, measurement techniques can confuse willingness-to-pay with ability-to-pay – a much

less ethical measure in that it would seem to undervalue the lives of the poor. Likewise, even individual behavior is inconsistent because it has been found that the amount an individual will pay to *avoid* a risk is not the same as the amount the same individual must be paid to *take* a risk (Wilson and Crouch, 2001). Furthermore, an individual's willingness-to-pay to save a life appears to vary depending upon whether the choice is made directly or through an open market (i.e., market interactions may degrade moral values) (Falk and Szech, 2013).

The valuation issue extends to other elements in an assessment that are frequently quantified via willingness-to-pay methods. MacLean (2009) suggests two essential difficulties. First, public well-being and willingness-to-pay are not always equivalent as some individuals have preferences that are counter to society's, and even their own, well-being. Second, economic valuation is an incorrect way to measure many abstract and essential values (e.g., duties associated with religion or community). Evidence suggests that individual behavior is inconsistent when asked to put a price on deeply held values and such requests are frequently met with moral outrage (Tetlock, 2003).

Using a non-monetary unit of risks and benefits is particularly useful in assessments focused on public health. Measures such as number-needed-to-treat and number-needed-to-harm (Holden, 2003) are commonly used in pharmaceutical, healthcare, and epidemiological risk assessments. However, it becomes more difficult to compare these assessments with non-medical priorities and to make policy recommendations that consider limited financial resources.

3.4.3 Discount rate

Another contentious issue regarding the use of monetary units is the choice of discounting rate. The time range of the assessment determines how much the discount rate will

affect the assessment. For example, the long range of climate change assessments[38] has raised concerns regarding intergenerational fairness: accounting for unrepresented future generations and assumptions regarding their values, technology, and wealth. Selecting lower discount rates gives more weight to future costs and future generations. Whether to use a constant or variable discount rate and its precise value reflect individual ethical value judgments (Revesz et al., 2014). Even the basic idea that future generations deserve ethical consideration is not universally accepted (cf. Visser't Hooft, 1999). However, Davidson (2009) notes that international laws increasingly acknowledge that physical distance is no longer an excuse for exclusion from risk considerations; eventually, separation in time may no longer be an acceptable reason for "empathic remoteness."

3.4.4 Other methodological considerations

One of the first methodological decisions is whether to perform a qualitative or quantitative risk assessment. Selecting a qualitative (i.e., non-numerical) assessment may indicate deep uncertainty, lack of confidence in available data, or even mistrust of available quantitative methods (cf. Pilkey and Pilkey-Jarvis, 2007). For example, the FDA issued new guidelines in 2013 that rejected a purely quantitative approach for performing comprehensive risk-benefit assessments for new drug approvals (FDA, 2013). Likewise, a quantitative assessment may be selected for good reasons, such as strong past performance (cf. Armstrong, 2001), or bad ones, such as to use numbers to imbue an assessment with an air of authority.[39]

[38] The Stern Review on the Economics of Climate Change (Stern, 2006) is one example that has been controversial (Nordhaus, 2007; Stern and Taylor, 2007).

[39] The non-intuitive nature of even relatively simple mathematics can be exploited by savvy advocates. A classic example of non-intuitive simple math is the potato paradox where a sack of impossibly soggy potatoes is 99% water by weight. If the potatoes are left in the sun to dry until they are only 98% water, the sack will lose half its weight! Much of the non-intuitiveness of

Hall (1988) notes "a frequent confusion of mathematical rigor with scientific rigor" which is partly based on the academic prestige associated with mathematical analysis.

The form of assessment will also depend on the exact definition and treatment of risk used (Hansson, 2005; Aven, 2012; Andretta, 2014; Aven and Zio, 2014). Möller (2012) outlines five common meanings of the term "risk": an undesirable event, the *cause* of an undesirable event, the *probability* of an undesirable event, the *expectation value* of an undesirable event, or a decision made with known probability outcome (i.e. Knightian risk). These definitions of "risk" range from broad and vague to narrow and precise. The use of any particular definition generally depends on the discipline in which it is used (Althaus, 2005). Decision theorists and economists are more likely to use the last two definitions, but even those academic fields do not use "risk" consistently.[40] In the relatively young field of risk analysis, the expectation value interpretation of risk has become widespread along with the use of probabilistic risk analysis (Rechard, 1999).

The expectation value interpretation has been criticized (Hansson, 2012) for failing to account for the seriousness of a particular risk. That is, the expectation value for a low probability, high consequence event can be the same as the expectation value for a high probability, low consequence event. For example, the expectation value for a 1-in-200,000 chance of losing $1 million is the same as for a 1-in-2 chance of losing $10.

A related concern is with the criterion of economic efficiency – an assumption that accompanies the use of expectation values. The economic goal of optimization assumed in most analyses can come at the cost of missing less efficient, but more robust options (Ben-Haim,

---

math may be a function of how concepts are presented. When posed visually, there appears to be an innate understanding of probabilistic questions regardless of education (Fontanari et al., 2014)

[40] This inconsistency contributes to unwitting public confusion and exacerbates the gap created by formal and common definitions of risk (Hansson, 2012).

2012). The idea is illustrated in Figure 3 where option A may be more efficient, but option B is more robust. There is an ethical component involved in selecting an assessment criterion; is it better to favor efficiency to avoid waste or is it better to favor robustness to increase safety (e.g., in building codes)?

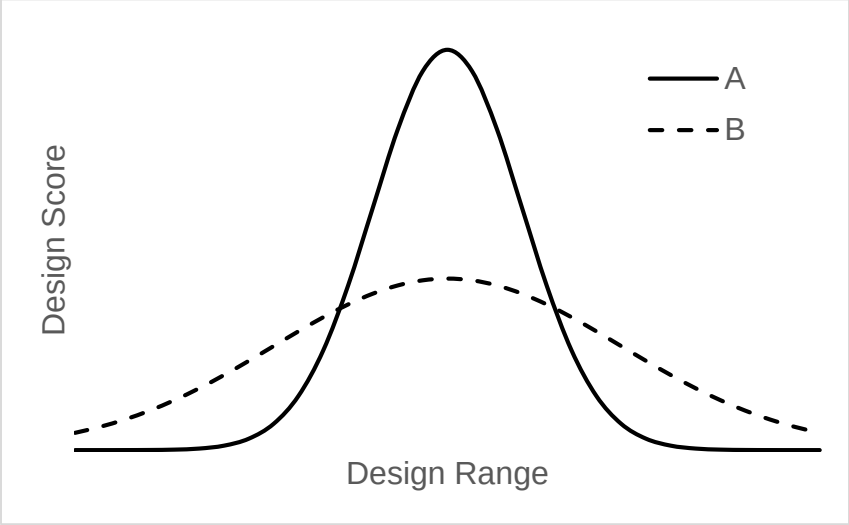

Figure 3: In this idealized scenario, option A is optimal for the most probable risk, but option B is more robust having a better design score for a wider range of outcomes.

3.4.5 Treatment of error

In any scientific statement or statistical test, there are two errors that can be made (cf. Kanji, 2006). A Type I error is finding an effect or phenomenon where it does not exist (i.e., incorrectly rejecting a null hypothesis). A Type II error is failing to find an existing effect (i.e., accepting a false null hypothesis). Traditionally, scientists have focused on Type I errors (false positives) because the emphasis is on avoiding the addition of false theories to the corpus of science or prematurely closing out a scientific debate (Hansson, 2012). However, in risk assessment and public policy, Type II errors (false negatives) may become more important (Douglas, 2000). For example, emergency managers will usually prefer to err on the side of

caution when making a decision regarding evacuation during a natural disaster. In these cases, it becomes questionable whether the epistemological values of traditional scientific inquiry are appropriate. In other words, for academic questions and questions regarding benefits, it seems appropriate to minimize false positives, while for questions of harm it may be more appropriate to be precautionary and err on the side of minimizing false negatives. For example, we would like our analyses to reject medical treatments that do not offer any benefit, but we would also like them to find any carcinogenic effects of synthetic chemicals used in consumer products. Accomplishing both goals effectively may involve selecting different criteria of what constitutes acceptable evidence.[41] It is not always as simple as reanalyzing data although that is sometimes an option. Detractors believe that an emphasis on Type II errors strays too far from defensible science (Sunstein, 2005).

From a policymaker's perspective, overly optimistic underestimates of potential harm (Type I error) may be worse than overly conservative estimates (Type II error). Overly conservative estimates of potential harm can lead to avoidance of potential benefits, over-preparing for harm, and less than optimal allocation of resources. However, underestimating risk is potentially much worse as it could lead to catastrophic loss of life. This is commonly referred to as an "asymmetric loss function" (e.g., Bijak, 2010).

3.5 Model Selection

Selection of the risk model can be the most contentious component of a risk assessment. Even simple models require a variety of value judgments. For example, if using a standard dose-response curve, how should data be fit to the curve: maximum likelihood, non-linear least

[41] The precautionary principle is a qualitative attempt at minimizing Type II errors within a body of science that was generated using Type I error minimization criteria.

squares, or even piecewise linear interpolation (MacGillivray, 2014)? While some decisions are purely epistemic, some are a mix of epistemic and ethical. For example, when selecting a low-dose cancer risk model, an analyst might choose a linear no-threshold model, a non-linear model, or even a hormesis (i.e., beneficial effects at low dose) model (Calabrese and Baldwin, 2003; Feinendegen, 2005) based on an epistemic value judgment. However, the selection process and whether the model accounts for susceptible populations (Bogen, 2014a, 2014b; Finkel, 2014) are also ethical value judgments reflecting the analyst's inclinations regarding whether a model should strive to be as scientifically "accurate" as possible or should err on the side of being conservatively protective (Nichols and Zeckhauser, 1988; MacGillivray, 2014).

As another example, epidemiologists may choose to use a continuous or discrete mathematical model to represent the behavior of a single-outbreak avian flu epidemic. A discrete model is useful in that it is easier to compare to epidemiological data (which is usually collected at discrete times) and is easier for non-mathematicians to use (Brauer et al., 2010). However, discrete models sometimes exhibit very different dynamics than the continuous models they are intended to duplicate when used outside of a narrow set of parameters (Mollison and Din, 1993; Glass et al., 2003). Researchers must make judgments regarding the relative value of "correctness" versus tractability when modeling a potential flu pandemic. Some of the same concerns arise when selecting between an analytical model and an approximated numerical model.

3.5.1 Theoretical versus empirical choices

The basis of a risk assessment model may be primarily theoretical or empirical.[42] Any distinction between the two types is somewhat artificial as most models are somewhere on the

[42] In the era of Big Data, empirical models are becoming increasingly popular.

spectrum between the two extremes. Because risk estimates involve projections into the future, they all have a theoretical component – even empirical trend extrapolations rely on the basic theory that future behavior will mimic past behavior (Wilson and Crouch, 2001). Nonetheless, there are special considerations for models that are more empirical or theoretical in nature. For example, defining and characterizing uncertainty can be difficult for theoretical models, while empirical models must assume that the data available is representative of the full range of possibilities and will not miss the "tail risks."

Opinions on empiricism constitute an important value assumption that rests on the perceived relationship between data and theory. Data are interpreted through theory; that is, facts and observations are "theory-laden" (cf. Feyerabend, 1975). Lyttleton (1978) summarizes the point well: "And it is here that one comes up against the strangely paradoxical nature of science, for the observations of phenomena are first needed to inspire someone to imagine an appropriate theory, yet they (the observations) cannot be claimed to be properly understood until a formal theory of them is available […] There can be no 'facts', no reliable 'evidence', until there are hypotheses and theories to test out." Data and theory are intertwined. Misinterpreting a lot of data with the wrong model only improves the precision of the error. Likewise, data mining is helpful only if the appropriate data are analyzed.[43] Data are important, but not a panacea to all risk assessments. As Einstein's conceptual thought experiments (*Gedankenexperiment*), used in the creation of relativity theory, demonstrated, empiricism is not always the best method for generating scientific knowledge.

[43] We try to measure only what we believe is valuable, but then we tend to value only what we've measured.

3.5.2 Level of complexity

Selecting the level of model complexity in a risk assessment is an epistemic value judgment similar to the selection of the assessment's boundaries. Complexity of analysis is often equated with thoroughness and appropriate representation of reality. Meanwhile, proponents of simplicity often invoke Occam's razor – the principle that simpler explanations are preferred. Additionally, analysts face the often completing goals of representativeness and usefulness. A broad, comprehensive assessment that accounts for the many complex consequences of a new technology may appropriately describe the risks and benefits, but may also not lend itself to the clear comparisons needed for a policy decision. Thus, the choice of complexity can be a tradeoff. While a simple assessment may be easier to explore and explain, it runs a higher risk of missing critical relationships. Likewise, a complex assessment might have a better chance of capturing all the salient components of a system, but also is harder to check, evaluate, understand, and compare to competing assessments (cf. von Winterfeldt and Edwards, 2007).

One might presume that complex systems generally require more complex analysis. However, there are often reasons to perform simple modeling of complex systems such as, lack of theory or known lack of predictive ability of complex models. For example, the time between certain catastrophic events (e.g., large earthquakes) follows an exponential distribution (Abaimov et al., 2007) which is "memoryless" – the time between events is independent of previous events (Solow, 2005). In this case, simple average time between events is as useful as a complex model (Cox, 2012a).

Pilkey and Pilkey-Jarvis (2007) make the further distinction that complex quantitative modeling, which has a successful history in the engineering sciences, has been frequently

misapplied to poorly-understood complex systems (e.g., ecosystem modeling).[44] These applications often have little predictive value due to uncertain input parameters and ordering complexity (i.e., there is often inherent uncertainty regarding the order, strength, and timeframe in which various model parameters will interact (Pilkey and Pilkey-Jarvis, 2007)).

Although the goals of the assessment and the subject investigated should dictate the level of complexity, the preferences of the analyst are influential. Mathematical sophistication tends to imbue an air of legitimacy to an analysis (Wachs, 1989; Stone, 2002) even when the rational basis for the numerical valuation is weak or intangible factors that are difficult to estimate and publicly defend are excluded from the analysis. Choosing complexity can be a conscious or unconscious consideration for models that will be used as "fig leaves, shields, and clubs" (Pilkey and Pilkey-Jarvis, 2007). Likewise, while simplicity has some practical appeal, it can also be an aesthetic value judgment. For example, Bondi (1978) notes that an aesthetic bias exists in the field of physics such that physicists prefer comprehensive or unified theories over theories that well describe a more limited range of phenomena. Explanations for this preference for "complete" theories rather than collections of "incomplete" theories, such as claims of mathematical elegance, are in themselves also aesthetic criteria.

3.6 Data Collection

There are many value judgments inherent in the process of data collection and selection. Some of the value judgments happen beyond the control of the risk analyst. For example, there is often considerable bias in available data due to widespread publication bias (Dickersin, 1990; Young et al., 2008) caused by the common value judgment that null results are not valuable

[44] The overuse is attributed to the ubiquity of computers and policies requiring cost-benefit ratios.

information (Rosenthal, 1979; Franco et al., 2014). For example, randomized clinical trials testing the efficacy of medications and psychological treatments for depression are more frequently published when a larger effect is found – leading to an overestimation of the efficacy of both treatment options (Turner et al., 2008; Driessen et al., 2015). This preference for positive effects[45] extends to the over-reporting of false positive findings (Ioannidis, 2005; Sarewitz, 2012; Nuzzo, 2015). For example, an effort to reproduce 100 psychology experiments found that the mean effect size of the reproduced studies was half of the original and the studies with significant results dropped from 97% to 36% (Open Science Collaboration, 2015). Likewise, developing regions appear to suffer from a publication bias that limits data availability (Sumathipala et al., 2004; Singh, 2006). However, many value decisions are made by the analyst.

3.6.1 Data screening

In situations where there are copious data or data of variable quality, a screening process may be used. For small studies where data are evaluated by hand, epistemic judgments are made regarding whether each datum should be incorporated into the risk assessment. In larger studies that use an automated screening mechanism, judgments must be made regarding the appropriate cutoff criteria. Criteria might include: adherence to good laboratory practices, findings with a certain statistical significance,[46] or similarity of surrogate data to the risk in question (MacGillivray, 2014). While criteria such as statistical significance are epistemic value decisions, they can have substantial ethical implications regarding the level of precaution in a risk analysis (cf. Douglas, 2000). For example, harmful effects below the current detection limit are often assumed to be acceptable risks (Hansson, 1999). While this practice may be pragmatic,

[45] This bias was first discussed by Francis Bacon in the *Novum Organon*, 1620.
[46] See Nuzzo (2014) for a discussion of the common misuse of P-values and Leek and Peng (2015) regarding the epistemic assumptions throughout statistical analysis.

it can affirm the implicit assumption that the low-dose threshold is well-established and uncontroversial.

3.6.2 Data for rare events

Some risk events are so rare that little to no data may be available as a basis for estimating risk. Several limited methods have been used to address rare event risks including: statistical methods, such as bootstrapping; expert opinion and variations, such as the Delphi method; using analogous situations; and bounding with scenarios (Goodwin and Wright, 2010).[47] Sometimes assumptions are used to simplify rare data analysis. For example, Gallucci (2012) uses a (questionable) zero-failure-rate approximation for the likelihood of an internally-caused nuclear power plant core meltdown accident. This approximation is based on no such events in recent decades[48] and assumed process improvements since the 1980s. Statistical methods for estimating the probability of events which have not yet occurred exist (e.g., Winkler et al., 2002; Quigley and Revie, 2011), but are far from being universally accepted.

One interesting subset of rare events is the low probability, high consequence risk (e.g., Baum, 2014). This includes known low probability events as well as improbable, but unrefuted, theories that suggest catastrophic risk (Cirković, 2012). These risks tend to get considerable attention in the media. One interpretation of the "social amplification of risk" theory (Kasperson et al., 1988) is that the public often overreacts to these risks. However, Coeckelbergh (2009) argues that inflated concern for these risks is actually a moral judgment by the public that catastrophic risks should not be reduced to a probabilistic expectation value.

---

[47] Historical records (e.g., old Tsunami markers in Japan) are an excellent source for scenario bounding of natural risk events. Even many anthropogenic risks (e.g., in bioengineering and geoengineering) have analogous natural correlates.

[48] The most recent were Three Mile Island in 1979, Saint Laurent Des Eaux in 1980, and Chernobyl in 1986.

From an epistemic perspective, Ort et al. (2010) argue that in these low probability, high consequence cases, any probability of occurrence is actually the probability of the event conditioned on the theory being correct. Thus, if the model is wrong, the rare event might not be as rare as experts estimate. It is important to note that the very idea of attempting to predict rare catastrophic events is controversial (e.g., Taleb et al., 2009). However, there remains considerable interest in quantifying the likelihood of catastrophic events and human extinction (Rees, 2013; e.g., Tonn and Stiefel, 2014)

3.6.3 Expert opinion

If the assessment is to use expert opinion, the method by which the expert opinions are used must be selected. For example, there is no consensus on the best method for pooling opinions (Hammitt and Zhang, 2013; Bolger and Rowe, 2015a; Cooke, 2015; Winkler, 2015) or whether expert opinions should be pooled at all (Reaven, 1987b; Keith, 1996; Bolger and Rowe, 2015b; Morgan, 2015). Recommendations for combining expert opinion range from complex weighting schemes to simple averages (cf. Clemen and Winkler, 1999). Critiques of pooling have pointed out the difficulty of determining the "best" experts and have even argued that pooling is no better than randomly selecting one expert's judgement (Reaven, 1987b). It is widely acknowledged in academic literature that experts are subject to a variety of biases, but attempts to address these biases are less common in practice (Sutherland and Burgman, 2015). Furthermore, some forms of cognitive bias (e.g., insensitivity to sample size) are more amenable to "debiasing" than others (e.g., overconfidence), while most forms of motivational bias (e.g., confirmation bias and preference for positive outcomes) can be particularly obstinate (Montibeller and von Winterfeldt, 2015).

3.7 Accounting for Uncertainty

The treatment of uncertainty in a risk assessment is a fundamental issue. Uncertainty exists in all measurements, models, and assessments. Whether deliberately or accidentally, analysts make epistemic value choices on how to handle the various forms of uncertainty. Parameter uncertainty will be discussed first followed by model uncertainty. Cox (2012b) provides the most complete summary to date of the various approaches for dealing with extreme uncertainty in risk assessments.

3.7.1 Deterministic versus probabilistic

The simplest choice is to ignore uncertainty and use a deterministic model to perform the risk assessment. The fully deterministic approach has been popular in computer modeling (Cipra, 2000) but it limits the uses of a quantitative risk assessment. In cases where there are no defensible estimates of uncertainty, a qualitative assessment may be more appropriate. This can be a reasonable approach with relatively new threats, such as those found in cybersecurity (e.g., Coull and Kenneally, 2012). However, even a deterministic model can explore uncertainty by varying the model parameters and building a range of scenarios. This is commonly done in climate modeling. This at least gives the analyst a range of potential outcomes albeit with no associated likelihoods.

3.7.2 Objective versus subjective probabilities

The most popular technique in risk analysis is to use a probabilistic model where parameters are represented as distributions and the output is likewise a distribution. Using the popular definition of risk as a probabilistic expectation value, even the simplest risk estimate requires a hazard and a probability. When using probabilities, the analyst must decide if the probability will be based solely on data or if it will also include subjective probabilities (i.e., a

frequentist versus a Bayesian interpretation of probability) (cf. Hansson, 2010b). The subjectivist approach is increasingly preferred (Flage et al., 2014) and often necessary when we consider trans-science – questions that are technically answerable by research, but will not be answered because of practical limitations (Weinberg, 1972, 1981). For example, a calculated failure risk of 10^{-6} per year may not be empirically verifiable if the risk applies to something large and expensive like nuclear power plants.[49] These types of low-probability events are unlikely to be testable or known with much confidence. Given the difference in confidence one might place on a probability computed from empirical data versus expert opinion, differentiating between empirical and subjective interpretations of uncertainty may be important (Reaven, 1988; Doorn and Hansson, 2011)[50] – especially if these various sources of data are being combined within the same analysis (e.g., a fault-tree). In some models, second-order uncertainty (i.e., uncertainty regarding the uncertainty) is included (Hansson, 2008).

3.7.3 Hybrid probabilistic methods

A third approach is to distinguish between two types of uncertainty: aleatory (normal variability within a population or process) and epistemic (lack of knowledge or incertitude). [51] Some risk analysts account for aleatory uncertainty with standard probability distributions and epistemic uncertainty with alternate techniques. One alternative is to represent epistemic uncertainty with intervals that represent upper and lower bounds on possible values. This is done in probability bounds analysis (Ferson and Ginzburg, 1996; Ferson and Hajagos, 2004) or the

[49] With approximately 500 nuclear power plants worldwide, a 1 in a million per year risk should, on average, occur once in 2000 years.
[50] This comparison is a simplification in that all empirical data have some subjective components and expert opinions have some empirical basis.
[51] Epistemic uncertainty can take several forms such as: data uncertainty, lack of data, unknown model/theory, conflicting models, general ignorance, etc.

more generalized method of imprecise probabilities (Walley, 1991; Weichselberger, 2000). Incorporating alternatives to probabilistic risk is not common for two potential reasons. First, the mathematical techniques may be less familiar to analysts. Second, and more importantly, there is a tendency for decision-theoretic analyses to assume that simple probabilistic analysis is always possible. Hansson (2009b) uses the term "tuxedo fallacy" to describe the error of treating real-world uncertainty as known probabilities (e.g., those commonly found in casino games).

3.7.4 Non-probabilistic methods

A fourth approach is to use alternatives to probability theory or extensions of probability theory such as evidence theory, also known as Dempster–Shafer theory (Dempster, 1967; Shafer, 1976, 1990), or possibility theory (Dubois and Prade, 1988; Dubois, 2006). Likewise, related issues of vagueness or ambiguity can be addressed by fuzzy set theory (Zadeh, 1965; Unwin, 1986). Newer risk analytic methods with novel treatments of uncertainty include info-gap analysis (Ben-Haim, 2006) and confidence structures (Balch, 2012). There are even semi-quantitative approaches (e.g., Aven, 2008a, 2008b) that take into account the varying degrees of confidence we have in knowledge (Mosleh and Bier, 1996; cf. Renn and Graham, 2005; Flage and Aven, 2009).[52] These approaches generally make fewer epistemic assumptions. That is, they use a broader, more conservative conception of uncertainty which may be useful when data is sparse or unreliable.

[52] These methods can be used in risk assessments with alternative conceptions of risk. A traditional risk definition is hazard multiplied by probability. The alternative (C,U) risk perspective defines risk as consequence combined with uncertainty (Rosa, 1998; Aven and Renn, 2009; Aven, 2010).

3.7.5 Selecting an uncertainty representation

There is no current consensus on the best way to represent uncertainty or when to use one form over another (Flage et al., 2014). Furthermore, there is no consensus on the number of forms of uncertainty. Methods of representing uncertainty reflect an analyst's epistemological philosophy.[53] Bayesians (e.g., Apostolakis, 1990; Winkler, 1996) generally argue that there is only one type of uncertainty in that it is always a measure of belief irrespective of its source and that probability is the best way to express this uncertainty (O'Hagan and Oakley, 2004). At the other extreme, van Asselt and Rotmans (2002) propose an uncertainty typology that describes five sources of variability and seven types of incertitude, but without methods for treating each distinct form of uncertainty. Likewise, Walker et al. (2003) construct an uncertainty typology organized by location, nature, and level with multiple subtypes. The resulting uncertainty matrix has thirty-five unique forms of uncertainty. The utility of this level of detail is debated (Krayer von Krauss et al., 2006; Norton et al., 2006).

Distinguishing between incertitude and natural variability does have known advantages over a single treatment methodology. For example, when an analyst has no knowledge of the value of a parameter within a bounded range, it is common to use a uniform distribution (a frequently used "uninformed prior" in Bayesian inference). However, treating interval data as uniform distributions requires the equidistribution hypothesis (Bertrand and Goupil, 2000)[54]. Under this assumption, each possible value is considered to be equally likely and the interval is represented by the interval's midpoint which represents the mean and median of a uniform

[53] Tannert et al. (2007) categorize uncertainty across a spectrum that ranges from certainty to nescience (i.e., unresolvable ignorance). They also distinguish epistemological uncertainty from moral uncertainty which applies to decision-making.

[54] The ultimate source of this idea is Bernoulli's and Laplace's principle of insufficient reason and more recently the principle of indifference (Keynes, 1921).

distribution. While the uniform distribution approach is computationally expedient and easy to understand, it can also disguise uncertainty. One result of the Central Limit Theory is that any distribution with finite variance (e.g., a uniform distribution) will converge to a normal distribution approximation when repeatedly convolved (e.g., Grinstead and Snell, 1997). As shown in Figure 4, adding two uniform distributions results in a triangle distribution that would appear to have more certainty than the original uniform distributions. As the number of distributions increases, the resulting normal approximation does spread, but the assumed central tendency is maintained. By comparison, the addition of two intervals of [0, 1] yields the even wider interval [0, 2]. Repeating the process yields ever wider and more uncertain intervals without an artifact of precision.

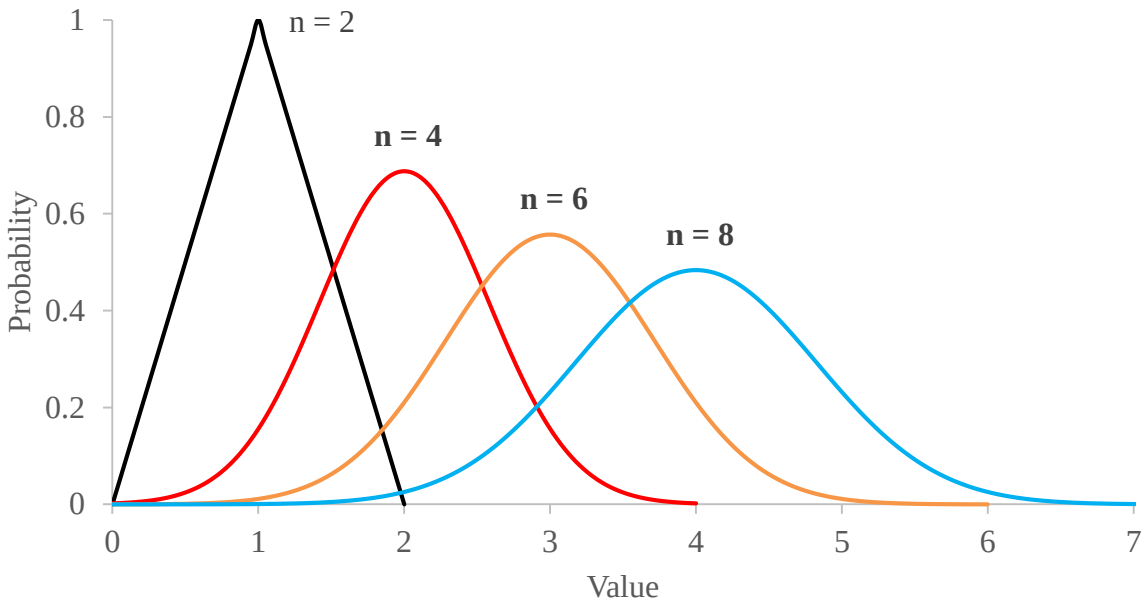

Figure 4: Convolution of *n* uniform distributions with bounds (0,1)

The practical implication of this phenomenon is that standard Bayesian and Monte Carlo techniques have an important limitation.[55] These approaches will yield valid results when the model inputs represent normally distributed natural variability. However, when the uncertainty due to lack of knowledge is large compared to the natural variability, the uncertainty in the model output will be understated.

In practice, uncertainty representation may be based on extraneous factors such as familiarity, academic tradition, or lack of knowledge of alternatives. Likewise, risk communication concerns (cf. Tucker and Ferson, 2008) might trump epistemic concerns. That is, an analyst might prefer mathematical simplicity over a slightly more informative, but less understandable method of treating uncertainty.

3.7.6 Treatment of variable dependence

Regardless of how uncertainty within a variable is treated, epistemic assumptions must also be made regarding the dependence among variables within the model unless the relationship is already well-known. A common occurrence in risk assessment is to ignore any potential dependence and to assume that all model variables are independent. However, this can cause the model to substantially misrepresent the potential range of outcomes. Ferson et al. (2004) provide a detailed discussion of inter-variable dependence in risk models including methods that allow an analyst to quantitatively account for limited knowledge about dependence between variables.

[55] Second-order Monte Carlo techniques are designed to addresses these concerns. Likewise, robust Bayesian or Bayesian sensitivity analysis (Berger, 1985; Berger et al., 1994) is a particularly useful technique in that is allows for the recognition of incertitude within Bayesian inference (Walley, 1991).

3.7.7 Model uncertainty

While parameter variability is often well-characterized in risk assessments, many forms of epistemic uncertainly are underestimated or ignored due to lack of appropriate methodology. Model uncertainty, also referred to as model form uncertainty or model structure uncertainty (Walker et al., 2003), is rarely acknowledged and the few methods for quantitatively addressing this form of uncertainty are not universally accepted (cf. Ferson, 2014). The analyst must make the epistemic value judgment of whether and how to account for model structure uncertainty.

In some cases, model structure may be the most important source of uncertainty. For example, Refsgaard et al. (2006) gives an example of five well-respected Danish consulting firms each generating a groundwater pollution risk model based on the same field data. Variations in methodological approach among the models resulted in conceptually unique models with unique structures and no common predictive capability.

Draper (1995) describes the general assumptions of a model, M, as $M = (S, \theta)$ where S represents the structural assumptions of the model and θ represents the parameter assumptions of the model. The uncertainty related to the parameter assumptions are often explicitly and quantitatively addressed in stochastic models as previously discussed. Meanwhile, structural uncertainty is rarely explicitly addressed. Often, a model structure is selected from a range of possibilities based on available data and then the modeler proceeds to account for parameter uncertainty while treating the model structure as a given (Draper, 1995). Uncertainty related to model structure is more difficult to characterize than parameter uncertainty primarily because it is not based on natural variability, but rather on a lack of knowledge that may not be amenable to probabilistic representation. The following reviews the few existing strategies for dealing with model structure uncertainty.

3.7.7.1 Lumped uncertainty

Where copious data exists, the modeler can use a split data set to first select the parameters during the model calibration phase and then evaluate the entire model during the validation[56] phase (cf. Klemes, 1986). Model structure uncertainty can be inferred from deviations found during the validation phase. The structural uncertainty can then be accounted for by either increasing parameter uncertainty until it is assumed to account for structural uncertainty (this occurs in inverse modeling methods) or by quantifying an explicit structural uncertainty term (Refsgaard et al., 2006). Using this approach has shortcomings. It relies on data that may itself be very uncertain or non-representative of the potential range of inputs. Likewise, separating model uncertainty from parameter uncertainty can be difficult (Vrugt et al., 2005). For models that are primarily predictive or have known non-stationary underlying processes (e.g., changes in precipitation patterns due to climate change), data often does not exist that can be used to directly assess structural uncertainty.

3.7.7.2 Sensitivity analysis with multiple models

One of the mathematically simplest methods of addressing model structure uncertainty is to create multiple models that address the range of possible model structures. Each model is evaluated separately and the output of all the models is summarized as a set of results. An example of this approach is climate modeling projections (e.g., Visser et al., 2000). This approach is simple to understand, but it can become arduous to perform a full analysis if there

[56] Validation is commonly defined as proving something to be true, but model validation is a much more modest process. The petroleum industry has historically used the term "history matching" instead of "validation" (Bredehoeft, 2003). This practice should be more widely used in other fields because it more accurately represents what is actually happening. Oreskes et al. (1994) use the term "evaluation" instead of "validation" for the same reason.

are many models. Likewise, summarizing the many results and guessing the likelihood of individual model outcomes can be difficult. Most importantly, there is a completeness problem because the analyst must describe the set of all possible models (Reaven, 1988).

3.7.7.3 Monte Carlo model averaging

A Monte Carlo probabilistic model (cf. Morgan et al., 1990) can account for multiple model forms by sampling each possible model. The sampling can even be weighted if some models are more probable. Expert elicitation is commonly used to estimate the probabilities of individual models although there are objections to averaging theoretically incompatible models (Reaven, 1988; Keith, 1996). This method also requires that all possible model structures be identified by the analyst. However, instead of evaluating each model independently, the output is a single probability distribution. Another concern is that averaged distributions will very likely underestimate worst-case tail risks which may be an issue for some risk assessments. If decision-makers are most concerned about catastrophic losses, then averaging models will obscure the range of modeled outcomes – some of which will predict that the extreme tail risks are more likely than the averaged model (cf. Ferson and Tucker, 2006).

3.7.7.4 Bayesian model averaging

Bayesian model averaging (cf. Hoeting et al., 1999) is similar to Monte Carlo averaging in that it combines all the identified potential model structures into an aggregated output. In this case, the model uncertainty, both parameter and structural, are evaluated together (Draper, 1995). Bayesian averaging has the same limitations as the Monte Carlo approach: completeness concerns, averaging incompatible theories, and underestimating tail risks.

3.7.7.5 Bounding analysis

Unlike averaging methods, a bounding approach does not combine the output of each model structure into a single output distribution. Using bounding analysis, the results for all the potential models are compared and an envelope is drawn around the entire set of results. The final product is a single bounded region that presumably will contain the correct model output. A benefit of this method is that all possible models structures need not be identified – only those that would yield the most extreme outcomes. While there is no guarantee that the analyst will be able to identify the extreme model structures, it is a lower epistemic hurdle to clear than identifying all possible models which should result in fewer surprises. Bounding analysis also avoids the issues of underestimating tail risks and averaging incompatible theories. It propagates rather than erases uncertainty (Ferson, 2014). Weaknesses include the inability to weight the credibility of individual model structures and the inability to distinguish likelihoods within the bounded region.

3.7.8 Unknown unknowns

Donald Rumsfeld (2002) brought the phrase "unknown unknowns" into common usage, but the concept has been recognized for a very long time. In Plato's *Meno*, Socrates proposes that one cannot inquire into a thing about which one is wholly ignorant. By definition, we are ignorant of unknown unknowns and therefore we almost universally exclude them from risk assessments.[57] Nonetheless, approaches for considering ignorance have been proposed. Jonas

[57] However, after events occurs, claims of events being unpredictable are often overused as a way to deflect blame for poor risk management (Paté-Cornell, 2012; Paté-Cornell and Cox, 2014).

(1984) suggests analysts expand their conception of possibilities through imaginative thinking and Attenberg et al. (2015) suggest expanding public discourse.

An easier approach is to merely increase uncertainty estimates. Taylor (2012) notes that unknown unknowns are often lumped into a single "catch-all probability" along with other unquantifiable factors.[58] While this type of uncertainty is unquantifiable, it is at least a display of Socratic wisdom to acknowledge that we are often uncertain of our own level of ignorance (cf. Elahi, 2011; Sahlin, 2012).

In summary, while there is uncertainty inherent in all knowledge, formal risk assessments tend to make specific, often quantitative, claims regarding the level of certainty in an analysis. Thus, special attention is warranted regarding the caveats placed in an assessment which reflect the analyst's value judgments regarding what is known and knowable. An analyst may believe that an assessment has captured all the salient points worthy of consideration to a degree of accuracy and precision that conclusively answers the question posed by the decision-maker. This analyst is likely to present the findings with few caveats. A more skeptical and humble analyst will qualify the assessment in hopes that readers do not over-interpret the results.

3.8 Comparing Risks

One reason to create a risk estimate is so it can be compared to other risks or risk management options. However, the comparison process is full of value assumptions. For example, one of the most common comparisons is to naturally occurring levels of a potential harm (e.g., background radiation levels) (Hansson, 2003). However, it is a value judgment to adopt this natural-level as the standard for comparison as there is often scant reason to assume that this level constitutes no-harm or an acceptable level of harm merely because it's natural or

[58] This is sometimes pejoratively termed a "fudge factor."

unavoidable. The widespread advocacy and use of sunscreen suggests that public health experts and the general public do not find all natural risks acceptable – especially when they can be avoided. Experiencing a level of harm by default does not imply that technologies that subject us to similar levels of risk are acceptable (Fischhoff et al., 1978; Slovic, 2000).

Along the same lines, public concerns regarding various technologies (e.g., synthetic food additives or genetically modified foods) are sometimes based on the "unnaturalness" of the technology (Viscusi and Hakes, 1998; Hansson, 2012). However, this assessment is based on a belief that naturally occurring substances are safer. This is a reasonable assumption in the sense that humans have evolved and co-existed with naturally occurring materials for a very long time. Compared to a synthetic material that may have existed for only a few years, this extensive experience is helpful for judging safety. However, given the vast range of naturally occurring substances, it's a naïve generalization to assume natural equates to safe. For example, arsenic is often naturally occurring in groundwater but is considerably more dangerous to human health than many anthropogenic risks. This basic assumption is even ensconced in US regulations; the FDA has far fewer requirements for botanical medicines (sold to the public as "nutritional supplements") than for synthetic drugs (Marcus and Grollman, 2002). Even equivalent harms caused by natural sources are perceived to be less scary than human-caused harms (Rudski et al., 2011).[59]

3.8.1 Incommensurability

One of the most basic assumptions in risk assessment is the belief that risks can be compared. Whether a risk assessment is part of a risk-benefit analysis, risk-cost-benefit analysis,

[59] The source of this distinction appears to be the "risk as feeling" model (Loewenstein et al., 2001) and affect heuristic (Slovic et al., 2007) which argue that perceptions of risk are very dependent upon an individual's positive or negative emotions associated with a risk.

or multi-criteria decision analysis (Edwards and von Winterfeldt, 1987), the basic assumption is that all options are comparable.[60] It may seem obvious that risks that are different in nature (e.g., health risks and risk of habitat loss) are difficult to compare due to the need for a controversial common unit of measurement (e.g., economic value) or equally controversial subjective ranking (Espinoza, 2009). However, even risks that appear to be of the same kind (e.g., all risks that can shorten a human life) can still be incommensurable due to important ethical distinctions. Hansson (2005) argues that public rejection of quantitative risk assessments in the past are due not to risk communication failures, but rather the failure of these formal assessments to account for ethical distinctions important to the public, such as justice and consent (e.g., Kuzma and Besley, 2008). Some distinctions frequently ignored in quantitative risk assessments that do not share normative equivalency include (Slovic, 1987; Gillette and Krier, 1990; Cranor, 2009; Espinoza, 2009):

- natural versus anthropogenic risks;
- detectable versus undetectable (without special instrumentation) risks;
- controllable versus uncontrollable risks;
- voluntary versus imposed risks;
- risks with benefits versus uncompensated risks;
- low incertitude versus high incertitude risks;[61]
- risks central to people's everyday lives versus uncommon risks;[62]

[60] The precautionary principle is a less comparative decision tool in that it only implies a comparison between the potential risk and doing nothing.
[61] This is due to ambiguity aversion – the preference for known risks over unknown risks (cf. Fox and Tversky, 1995) commonly phrased as, "Better the devil you know than the devil you don't."
[62] This may be one reasons that virologists often underestimate the biosafety risks of the pathogens they work with compared to scientists who have less working experience with

– future versus immediate risks; and

– equitable distribution of risks and benefits (in both space and time).[63]

In each pair above, the first risk type is generally more acceptable than the second risk type. In practice, risks often fit multiple categories and can be ranked accordingly. For example, common, self-controlled, voluntary risks, such as driving, generate the least public apprehension whereas uncommon, imposed risks without benefits, such as terrorism, inspire the most public concern. Ignoring these distinctions can lead to a simplification of risk that renders a risk assessment unusable for building public consensus. It is for this reason that risk governance guidelines stress the importance of providing social context for risk assessments (Renn and Graham, 2005). While the importance of these distinctions has been understood for some time in theory (Jasanoff, 1993), there remains limited evidence of this occurring in practice (Pohjola et al., 2012).

3.8.2 Risk ranking

If an assessment compares risk, a risk ranking method may be used. The ranking may be quantitative or qualitative and it can take many forms: letter grades, number grades, cumulative probability distributions, exceedance probability or F-N curves (cf. Evans and Verlander, 1997), color or word categories, etc. The ranking method selection is an epistemic value judgment with important communication implications (MacKenzie, 2014).

Some ranking methods are incomparable to other methods (Cox et al., 2005; Cox, 2008; Rozell, 2015). A simple example demonstrates that even seemingly comparable ranking systems

---

dangerous pathogens. This is one of many factors to be considered in the Chapter 6 analysis of the GOF research controversy.

[63] Similar justice issues arise when the exposed, the beneficiaries, and the decision-makers are non-overlapping groups (Hermansson and Hansson, 2007).

can give rise to mathematical inconsistencies. Let us rank the risk of two events, *A* and *B*, and define the risk, *R*, associated with each event as being the product of both the probability, *P*, of the event occurring and the cost, *C*, of the event, such that $R = PC$. If we have detailed information for the probability and cost of events *A* and *B*, such that $A = (0.4, \$400)$ and $B = (0.1, \$900)$, we would say that $R_A = \$160$ is riskier than $R_B = \$90$.

However, if we had less confidence in this information or were adding this assessment into a larger assessment with different parameters, we might decide to use qualitative methods. Using two equal-sized categories of low (L) and high (H), it would be reasonable to assign the probabilities as L = (0, 0.5) and H = (>0.5, 1). Likewise, costs could be categorized as L = ($0, $500), and H = (>$500, $1000). Using these qualitative labels, then R_A = (L, L) is less risky than R_B = (L, H). In this case, switching from quantitative to qualitative ranking and using a seemingly logical categorization caused the risk rankings to reverse. This issue is commonplace in qualitative risk ranking.

Most risk rankings are based on consequences and probabilities, but rankings involving ethical dimensions, such as source of risk (e.g., Gardoni and Murphy, 2014), have been proposed. This form of ranking is more inclusive, but also extremely subjective. For example, what is the most concerning: 1,000 deaths caused by heart disease, 100 smoking-related lung cancer deaths or one homicide? The question is almost meaningless when stripped of its context as can occur in a ranking exercise.

3.9 A Value Assumption Roadmap

Given the variety of value assumptions made during a typical risk assessment, an aid in the process is useful. Rather than organize the assumptions thematically, a chronological approach is used here so it can serve as a checklist for risk analysts. The list, detailed in Table 1,

is not exhaustive in its coverage of potential value judgments, but it is intended to cover the most common and contentious assumptions that, left unexamined and unaddressed, decrease the utility of a risk assessment.

Table 1: A process map of value judgments in risk assessment. Details for each question are discussed in the previous sections.

Step	Fundamental Value Questions
Selecting a topic	– How are hazards screened? – What heuristics are influencing choice?
Defining the assessment boundaries	– What is an appropriate time, space, and population? – Holistic or component analysis?
Choosing the form of the assessment	– What unit will be used? – What is a life worth? – What are deeply held values worth? – What discount rate should be used? – Qualitative or quantitative? – Which definition of risk? – Maximizing efficiency or resiliency? – Focus on preventing false positives or false negatives?
Model selection	– Accuracy or precaution? – Theoretical or empirical? – Simple or complex model?
Data selection	– How are data screened? – How are rare events treated? – How is expert opinion used?

Accounting for uncertainty	– Deterministic or probabilistic?
	– Objective or subjective probabilities?
	– How is incertitude addressed?
	  – Dependence between model variables?
	  – Model structure uncertainty?
	  – Deep uncertainty and unknown unknowns?

Comparing risks	– Can the risks be compared? – Qualitative ethical distinctions? – Risk ranking?

3.9.1 Example: risk analysis of farmed salmon

To see the utility of the map, it helps to apply it to a real risk assessment debate. A past risk assessment with a series of conflicting assessments demonstrates the importance of the value judgments in risk assessment. In this example, the original study (Hites et al., 2004) found high

concentrations of carcinogenic organochlorine contaminants in farm-raised salmon and suggested that the risk of consumption outweighed the benefits. The analysis prompted a series of strong response letters. One letter (Rembold, 2004) pointed out that, even using the conservatively protective EPA linear cancer risk model, the expected number of cancer cases from consuming farmed salmon was a fraction of the number of cardiovascular mortalities that would be avoided by eating salmon. Furthermore, the quality of the data was not the same; the cardiovascular benefits data was based on randomized clinical trials while the cancer risk data was based on less reliable observational studies and nonhuman dose-response models. The response by Hites et al. raised the possibility of other non-cancer harms (e.g., neurological) from contaminated fish consumption as well as a reminder that the beneficial omega-3 fatty acids found in salmon were available from other non-contaminated dietary sources.

A second letter (Tuomisto et al., 2004) also compared the cardiovascular benefits of salmon consumption to cancer risks and additionally included a value-of-information analysis to argue that any uncertainty in the risk and benefit data did not affect the assessment that salmon consumption had net health benefits. For this reason, the letter went so far as to state that any suggestion to limit salmon consumption was nonscientific because it ignored the data. Again, the response pointed out that fish were not the sole source of dietary cardiovascular benefits.

A third letter (Lund et al., 2004) questioned the EPA's linear dose-response model citing additional data that suggested there were no effects at low exposures. The response by Hites et al. argued that the Lund et al. study used a sample size too small to detect the estimated cancer risk. Furthermore, they pointed out that the potential neurological and cancer risk is larger for young people due to bioaccumulation while the cardiovascular effects primarily benefit older individuals which suggested the need to distinguish subpopulations. A fourth letter (Weaver,

2004) questioned the mechanism by which fish become contaminated in order to argue that farmed salmon are no riskier than wild salmon.

Looking at the various critiques and responses, value assumptions are made that correspond to each step of the risk assessment process. Boundary value assumptions are made regarding what can be counted as a risk (cancer and neurological impairment) or a benefit (cardiovascular health) and whether sensitive populations (e.g., children) should be considered. Likewise, there are debates regarding what is the appropriate risk model (the EPA linear model or a no-effects threshold model), what data are worthy of including in the assessment (e.g., observational studies, animal studies, and small sample size studies), and how important it is to account for uncertainty. Finally, there is an important value judgment regarding risk comparison: will people who potentially stop eating salmon substitute other foods rich in omega 3 fatty acids (e.g., can salmon be compared to walnuts)?

3.10 Discussion

The map of values ignores some of the more uncontroversial values inherent in risk assessment as well as broader value discussions within science. For example, the general conception of what constitutes quality science is itself value-based (e.g., Merton, 1942)[64] – and much debated (Kuhn, 1962; Feyerabend, 1970; Lakatos, 1970). However, it is not always clear what epistemic, ethical or aesthetic values are considered to be uncontroversial or for how long they will remain uncontested (Hansson and Aven, 2014). For example, maximizing happiness is a value common in contemporary economic analyses. However, this is not a universal goal; different eras and cultures have valued duty over self-interest (cf. Kant, 1785)

[64] "The ethos of science is that affectively toned complex of values and norms which is held to be binding on the man of science" (Merton, 1942).

The emphasis placed here on subjective values in risk assessment is intended to help risk analysts improve their work. It is important to dispel any notions that "subjective" is a derogatory term or that it is necessarily arbitrary. Subjective values can be deeply-held convictions with rational supporting arguments. Past neglect of value judgments, especially ethical considerations, has ignored the importance of emotional content in decision-making or treated it as adventitious (Roeser, 2011). Appreciation of these value-laden assumptions improves both the process and the final assessment.

3.10.1 Choice of ethical framework

The following is a brief review of some popular ethical frameworks proposed in moral philosophy as they pertain to risk assessment. The purpose is to point out that no normative theory is perfect and that the ethical underpinnings of any risk assessment can be a potential source of controversy. While an ethical framework must eventually be chosen as a basis for assessment, recognition of the weaknesses of any particular approach is necessary if an analyst is to address these weaknesses or at least openly communicate them.

Quantitative risk-benefit analysis uses utilitarian or consequentialist ethics (Hansson, 2005); that is, an action is good if the consequences are more beneficial than harmful. In the most simplistic form of consequentialism, the simplest way to incorporate probability is to judge the consequences in retrospect considering only what actually happened (Bergstrom, 1996). Using this simplistic approach, speeding down a busy road would only be unethical if someone was injured. However, this approach leads to irresponsible risk-taking behavior (i.e., actions that lead to unlikely, but easily avoidable harm) that is contrary to typical conceptions of moral behavior.

Instead, the common method for accounting for uncertainty in a utilitarian framework is to use expected utility. While this addresses the risk-taking issue, it creates an aggregation problem (Hansson, 2012). All the risks and benefits to all parties involved are combined and compared; the individual is ignored. Even if we can legitimately compare the risks and benefits of different individuals, it is not clear that we can say that the benefits to one individual compensates for the risks to another (Kelman, 1981).

Alternative frameworks to utilitarian ethics include: deontology (duty-based or rights-based ethics), virtue ethics (which focuses on moral character and moral wisdom), and social contract ethics (which focuses on process).[65] Using rights-based ethics resolves the aggregation problem of utilitarian ethics, however, it is still problematic because there are no recognized probabilistic limits to rights (Nozick, 1974; Hansson, 2004). For example, an innocent person's right to live is generally considered to be absolute. That is, various declarations of human rights (e.g., Article 3 of the 1948 UN Universal Declaration of Human Rights, "Everyone has the right to life, liberty and security of person.") do not make conditional statements about the right to life. Yet, we do not prohibit a wide range of activities (e.g., driving) that create a very small probability of killing an innocent person or even some activities that have an appreciable probability of killing an innocent person (e.g., targeted drone strikes against terrorists).

While virtue ethics avoids the issues associated with consequentialist and deontological ethics, it is perhaps the least practical framework for risk assessment because ethical criteria under conditions of uncertainty are left to the moral wisdom of the analyst. This is a rather vague

[65] One way to organize different ethical frameworks is by their emphasis on different parts of an ethical act (Koehn, 1995). Virtue ethicists focus on the character of the actor, deontologists focus on the act itself, social contract ethicists focus on the decision process, and utilitarians focus on the consequences.

solution (cf. Nussbaum, 1999) because we don't know how variable moral wisdom is among risk analysts nor do we know if differing conceptions of moral wisdom renders risk assessments incomparable.

Social contract ethics can have aggregation issues similar to utilitarian ethics (Hansson, 2004), but with consent rather than expected utility. Actual consent is too difficult to obtain in a large and interconnected society so we must instead use hypothetical consent.[66] However, this may involve forming a majority consensus that still ignores the individual. Rawls' (1971) social contract version of equity envisions a procedure where a system of rules is constructed before individuals know their place in the system. Ideally, rational self-interest then generates a maximally equitable system because decision makers would employ a "maximin" criterion which, in the face of uncertainty, selects the option that maximizes the benefits of the worst outcome (cf. Harsanyi, 1975; Rozell, 2014). However, as Harsanyi argued, the maximin criterion can also result in some rather unrealistic choices. For example, the criterion might lead one to avoid travel that only slightly increased the risk of accidental death, but would also substantially increase personal wealth and happiness (e.g., driving across country to take a dream job).

The contrast between the utilitarian and social contract approaches is similar to two competing economic policy decision goals: Kaldor-Hicks efficiency and Pareto optimality (cf. Stokey and Zeckhauser, 1978). The Kaldor-Hicks criterion requires maximizing net social

[66] Zandvoort (2009) argues that actual consent is necessary for an ethical society and that the only alternative is unlimited liability when consent is not granted. That is, those harmed by a technology should be able to seek full redress from the creators or users of the technology. Because this level of legal liability may be unbearable in cases of irreparable damages (e.g., species extinction), the implication is that certain technological activities are not economically feasible whenever actual consent is not possible.

benefit while the Pareto criterion prefers actions that improve at least one individual without making any other individual worse off.

The aggregation problem seen in utilitarian ethics and the Kaldor-Hicks criterion, which forms the basis of cost-benefit analysis, artificially obscures legitimate problems, such as the natural rationality of "not in my backyard" issues. When an unpleasant, but necessary, project is sited, the local residents are rationally self-interested to oppose the project because they comprehend the concentrated harms and dispersed benefits that will occur. Treating individuals as moral units can avoid this issue. For example, biomedical ethics uses a principle of beneficence or individual-centered risk assessment (Hansson, 2012). That is, the benefits to future patients cannot be considered in a risk-benefit analysis for a patient to participate in a drug trial – only personal benefits and risks.

3.10.2 Ignoring the normative

One potential reason for the common lack of normative considerations in formal risk assessment is the somewhat artificial and prejudiced distinction between risk assessment and risk perception (Coeckelbergh, 2009). This distinction is often paternalistically interpreted as risk assessment being the domain of experts where "real" risks are calculated and risk perception as the domain of the public where risks are imagined (e.g., Choi, 2013). This is not to suggest that the distinction is false. The training, experience, and considerable effort of risk experts is beneficial and informative to risk decisions. However, the formal methods in the field of risk analysis have traditionally idealized objective and quantitative analysis at the expense of difficult-to-quantify, but important risk characteristics. Objective, but narrow assessments are not clearly superior to subjective, but broad assessments.

Another reason may be the relative paucity of normative work in risk analysis. The inherent ethical value judgments that pervade risk analysis have led to claims that the field is a branch of ethics and that risk analysts need as much training in ethics as in economics or decision theory (MacLean, 2009). Given the proposed level of value judgments inherent in risk assessments, it seems reasonable that more philosophers should be involved in risk assessments – both epistemologists and ethicists. While the concept of moral experts has been rejected in the past, it is only if one defines a moral expert as someone who is an arbiter of what is right and wrong (cf. Gesang, 2010). Rather, moral experts exist if we regard them as individuals who act as knowledgeable guides familiar with the various moral concepts that frequently arise in decision-making (Singer, 1972; Nussbaum, 2002).

3.11 Conclusion

Risk assessment has become increasingly difficult for two reasons. First, for most of human history, the management of risk was simply trial and error (Wilson and Crouch, 2001). New technologies were adjusted or abandoned as evidence of harm was discovered. However, as the twentieth century unfolded, the power and scale of technology grew and it became evident that society needed to manage risks proactively. Second, science and technology also allowed us to recognize and measure new subtle hazards that required assessment. The result was a new multidisciplinary field with new and considerable challenges.

Where risks and benefits are clear and certain, formal risk assessment is unnecessary; it is intended for contentious issues where the correct policy decision is not obvious (Wilson and Crouch, 2001). However, it is in these very situations where the many inherent epistemological and ethical judgments that the analyst (maybe unknowingly) makes can influence the outcome of a risk assessment.

The purpose here is not to criticize formal risk analysis except, perhaps, in the way in which it is employed. If an analyst (or the funding source) believes that risk assessments will produce objective answers, there may be a subtle influence on all steps in the risk assessment process and a general understatement of subjectivity and incertitude. While the purpose of risk assessment is understood in theory, it may be less so in practice (cf. Apostolakis, 2004). This is peculiar considering that the contingent nature of risk analysis has been known for decades; "Risk assessment cannot be intensive nor selective in the way traditional science is. It must deal with questions as they arise without regard to their disciplinary assignment or to the quality and completeness of data that are obtainable or at hand" (Cumming, 1981). As methods and techniques within the field of risk analysis improve over time, it is necessary to remind ourselves frequently that risk assessment depends on science, but encompasses more than science.

Acknowledging risk in the decision-making process is always a wise endeavor and the various approaches developed in the field of risk analysis help guide the process of considering risk. Rather, a review of the value assumptions in risk assessment is used to emphasize that the process is inherently rife with assumptions that hinder unqualified claims of objectivity and neutrality. Narrow risk analyses of well-understood phenomena with ample data might be uncontested, but we must be realistic in our expectations of formal risk assessment to resolve public debates regarding controversial research, environmental contaminants, and emerging technologies. The various forms of value assumptions made in any risk assessment are a primary reason for the common inability of scientific studies to resolve disputes over the risks of new technologies (Sarewitz, 2004; Small et al., 2014). If astute stakeholders and policymakers intuitively understand the limitations of risk analysis, an assessment that is overly conclusive may be dismissed under the assumption that the analyst was deeply biased or naïve.

We still do not understand how risk analysis can be used to build consensus and reach decisions (Aven and Zio, 2014). However, the explication of uncertainty can aid rather than hinder public trust in formal risk analysis (Rabinovich and Morton, 2012). When assumptions are hidden or not obvious, policymakers may feel intuitively distrustful of an assessment without being able to pinpoint their misgivings. Ultimately, the reason to perform a risk assessment is to provide insightful or actionable information. Considering and clearly discussing the value-laden assumptions in a risk assessment improves trust in the provided information by allaying concerns of hidden agendas.

4 TECHNOLOGICAL RISK ATTITUDES

A common theme of policy debates over controversial research and technology is a call for risk-benefit analysis that is based on "sound" science. But which science studies have value or credibility is a contentious issue in its own right (Yamamoto, 2012; Sarewitz, 2015a). There is some consensus over what good science should look like in theory, but much less regarding what it looks like in practice or in individual cases (Small et al., 2014). Additionally, as argued in Chapter 2, there are many methods of valuing the benefits of research. While particular methods periodically rise and fall in popularity, none of them are universally accepted and often yield disparate results. Likewise, as argued in Chapter 3, risk assessment is a process inherently full of value-laden assumptions that can substantially affect the assessment outcome. The outcome of these two lines of argument is that risk-benefit analysis usually fails to resolve debates over controversial research.

Jasanoff (2007) critiques the modern attitude that difficult technological decisions are resolvable if only further research is conducted. She argues that not only is some uncertainty unavoidable, but often the most difficult parts of the decision process are not scientific, but rather ethical and political. Thus, delaying decisions for more fact-finding is often a form of wishful thinking. Calls to perform risk-benefit analysis for emerging technologies often fall into this category.

For example, as of early 2015, research on the health effects of e-cigarettes is sparse and inconclusive. Little is known of the various chemical hazards of the non-nicotine constituents of e-cigarette products (Henkler and Luch, 2015). Nonetheless, competing groups of public health scientists are taking strong positions on e-cigarette regulation based primarily on whether they see the technology as a replacement for or gateway to traditional tobacco use (Fairchild and Bayer, 2015). Further research should help resolve the question or at least clarify its dimensions.

However, regulatory decisions will ultimately be based on non-scientific issues over acceptable harm and the limits of freedom. Even reasonable individuals with the same public health goals and appreciation of science-based policy can arrive at very different conclusions. This leads to a fundamental question underlying the subjective dimensions of risk-benefit analysis. *Why* do reasonable individuals arrive at very different conclusions?

One possible answer comes from individual attitudes regarding risk and technology. Individuals tend first to evaluate novel information using mental shortcuts (heuristics) that are replaced by more systematic reasoned thinking as familiarity increases (Chaiken, 1980). The process of evaluating research and new technologies is no different (Scheufele, 2006). The heuristic-to-systematic thinking model also applies to interpreting risk assessments (Kahlor et al., 2003). New information can influence an individual's attitude, but the pre-existing attitude also influences how new information will be interpreted (cf. Eagly and Chaiken, 1993). That is, individuals with positive attitudes about technology will tend to expect more benefits from new technologies (Kim et al., 2014).

Many factors can influence attitudes towards technology. In this chapter, I first discuss the nature and potential origins of attitudes about technological risk that may account for why well-informed and reasonable people disagree on such matters. I then summarize the effects of these disagreements as a general "principle of controversial research."

4.1 Technological Optimism and Skepticism

It is natural to wonder what causes an individual to see a particular technology as primarily benign or dangerous. For example, the Haber-Bosch process converts atmospheric nitrogen to ammonia in the presence of natural gas, heat, and pressure. What causes one person to view this method of artificial nitrogen fixation as a dangerous innovation that has killed

countless people by providing crucial feedstock for the instrumentality of war, while another person seizes upon the countless people fed because the process provides crucial feedstock for the fertilizer industry?

Although attitudes regarding technological risk exist along a broad continuum, for our purposes it is helpful to define two general categories: technological optimism and technological skepticism. This simplification is appropriate because we are primarily concerned with public policy formulation and the binary outcome of whether an individual endorses or opposes a particular line of research. It has also been used by other academic and popular writers. For example, ecological economist Robert Costanza outlines four future scenarios that hinge on collective technological optimism or skepticism (Costanza, 1999, 2000). He further differentiates futures based on whether these technological attitudes are right or wrong. If the prevailing worldview of technological optimism is supported by future technological wonders that solve all our problems, Costanza describes the future as "Star Trek" and, if incorrect, "Mad Max". Likewise, a world following more technologically skeptical policies would result in "Ecotopia" if proved correct or needlessly oppressive "Big Government" if wrong. However, these scenarios don't address why individuals believe in the likelihood of various future scenarios. But, first, let us start by defining what we mean by a technological optimist or skeptic.

4.1.1 Technological optimism

Technological optimists believe in the liberating power of technology – modern medicine liberates us from disease and space exploration literally liberates us from this planet. This attitude is captured in the modernist[67] movement (cf. Reaven, 1989b) and is still a common

[67] Modernism, with its various incarnations in Western philosophy, architecture, art, literature, etc., is embodied in Ezra Pound's famous command, "Make it new!"

American attitude (e.g., Ridley, 2010; Tierney, 2010) with good reason. Over the past 160 years, life expectancy has increased linearly in the most developed nations (Oeppen and Vaupel, 2002) and trends suggest there may be no foreseeable limit to the average human lifespan (Cohen and Oppenheim, 2012). The technological optimist has a basic faith in technology and requires proof of harm to abandon a specific technology. An extreme version of technological optimism might be better labeled "technophilia" or the unquestioning embrace of new technology for its own sake.

This attitude is prevalent in Silicon Valley where, for example, there is often effusive praise for all things related to the internet and mobile communications technology. A common assumption is that social media has and will fundamentally restructure human communications to everyone's benefit. Many of these perceived benefits are as yet unimagined based on (the correct) assumption that technology is often repurposed by its users. Technological optimists see future repurposing of technology as empowering and creative rather than harmful or disenfranchising.

4.1.2 Technological skepticism

Technological skeptics reject the technology-as-panacea paradigm. This attitude is common in postmodernists (along with small enduring enclaves of pre-modernists, such as the Amish) and is linked to some of the failings of modern industrialization (e.g., nuclear accidents). In particular, within the environmental movement, there has been an ongoing critique of modern Western society that includes a general aversion to technology. Seminal examples include: Aldo Leopold's "land ethic" (1949), Rachel Carson's *Silent Spring* (1962), and Arne Naess's "deep ecology" (1973). However, the roots of technological skepticism date back at least to the early nineteenth century when the Industrial Revolution was underway in Great Britain. The Luddite

rebellion, a brief spasm of violence between 1811 and 1813, was a reaction to the social upheaval and suffering caused as the steam engine and power loom rapidly shifted the wool industry in central England from family cottage weavers to coal-powered mills run by a few adults and cheap child labor (Sale, 1995).[68] The full history and tradition of technological skepticism is too rich to cover here, but some notable works include writings of Freud (1930), Ellul (1964), Dreyfus (1972), Abbey (1975), Heidegger (1977), Mander (1978), Jonas (1984), Postman (1992), Rifkin (1995) and Jones (2006).

Technological skeptics are more likely to question and critique the work of scientists and engineers (e.g., Watts, 2014). An extreme version of skepticism could be labeled "technophobia" – a knee-jerk reaction to new technology. Technological skeptics are often pejoratively labeled "neo-Luddites", but the attitude is not associated only with those who eschew technology. The ranks of technological skeptics include engineers (e.g., Joy, 2000) who recognize that our present-day society is privileged and powerful primarily due to technology, yet warn that this technology hangs over our heads like the sword of Damocles.

4.1.3 Explaining the differences

If we accept that there are differences in technological risk attitudes that can be roughly categorized, we invite questions regarding the nature and measurement of these difference. The following is a summary of notable attempts, with varying degrees of success, to explain the differences in risk attitudes that are relevant to technology.

[68] It is not coincidental that Mary Shelley's *Frankenstein*, a seminal piece of technologically skeptical literature, was published in London five years later in 1818.

4.1.3.1 Cultural theory

The cultural theory of risk (Douglas and Wildavsky, 1982) proposes that risk attitudes originate from cultural influences. Individuals filter information through cultural world views that determine their perception of risk. The theory categorizes four risk ideologies: laissez-faire *individualists*, social justice *egalitarians*, utilitarian *hierarchists*, and apathetic *fatalists*. Individualists view human ingenuity as boundless and nature as robust which roughly corresponds to technological optimism. Conversely, egalitarians, who roughly correspond to technological skepticism, view nature as fragile and have a more precautionary view of technology. In the comparison between technological skeptics and optimists, it should be noted that there is a considerable middle ground in the general public consisting of ambivalent and indifferent individuals (Seidl et al., 2013). The last two groups in the cultural theory of risk fall into this category. The hierarchists do not have strong attitudes regarding technology, but they do value authority, expertise, and the status quo (van Asselt and Rotmans, 2002). Because formal analysis is conducted by "experts", the hierarchists should be the most influenced by the traditional economic analysis approach to risk assessment. Lastly, the fatalists correspond to those who have little faith in the predictability of nature or humanity or in our ability to learn from past mistakes. Because fatalists doubt their capacity for control and self-determination, they often opt out of the policy making process and see hubris in those who do not.

One implication of this classification is that many analysts might best be described as hierarchists whose analyses may tend to ignore the other cultural world-views. Open public debate and more democratic decision-making processes are particularly useful given such substantial differences in risk attitudes (Davidson, 2009).

These risk attitudes can make policy-making even more difficult. For example, individualists view government intervention in technology policy as a form of risk (i.e., policy makers may abruptly change policies) that can discourage investment (Chassot et al., 2014).

4.1.3.2 Psychometric measures

Despite its theoretical elegance, cultural theory has had less predictive success than psychometric approaches, such as risk-as-feeling concepts (Slovic et al., 2004), which still have limited explanatory power in empirical studies (Sjöberg, 2000). For example, cultural world view accounted for only 3% of the variance in surveys measuring the perception of risks and benefits of childhood vaccinations (Song, 2014). Psychometric studies have been found to be more useful for explaining why the public perception of risks has substantially deviated from calculated risks for particular hazards. Psychometric studies of risk perception also provide insight into attitudes about technology because attitudes about technology are dependent on risk perception. Sjöberg (2002) applies psychometric techniques to directly investigate attitudes about technology and finds that the acceptability of a technology is explained by its perceived utility and risk, and by whether the technology could be substituted with something else. Technologies that tamper with nature and may have unknown effects (e.g., genetically modified food) are perceived to be riskiest. Most importantly, even an expanded set of psychometric characteristics accounts for only about half of the variance, indicating that the basis of attitudes about technology is still not well understood. For example, in Yeosu, South Korea, which is home to a large petrochemical complex, a survey of 1446 residents regarding the perceived risk from the chemical industry found that less than 10% of the variance was explained by cultural theory or psychometric measures (Choi, 2013). In the study, demographic factors, such as education level or gender, were more strongly correlated.

4.1.3.3 Cultural cognition of risk

Cultural cognition theory (Kahan et al., 2006) is a hybrid of cultural theory and psychometric models which has been applied to risk perception topics ranging from vaccines (Kahan et al., 2010) to geoengineering (Kahan et al., 2015). Cultural cognition recognizes a set of psychological mechanisms that allow individuals to preferentially select evidence that comports with the values of the cultural group with which they identify (Kahan et al., 2011). For example, actions perceived to be culturally immoral are more likely to be seen as harmful even when given evidence to the contrary (Gutierrez and Giner-Sorolla, 2007). Likewise, evidence from experts that are perceived to share the same values are viewed as more credible. However, in general, cultural cognition has similar explanatory power to its parent theories. This hybridization may be the best of both worlds, but is still insufficient to adequately explain the source of risk attitudes.[69]

4.1.3.4 Other theories

Other theories embracing or eschewing technology have been proposed. One alternative conception of attitudes about technology comes from philosophers of technology. Framed in terms of trust, technological optimism and skepticism can be viewed as trusting or distrusting the reliability of technology as well as trusting or distrusting our ability to use technology appropriately (Ellul, 1964; Heidegger, 1977; Kiran and Verbeek, 2010).

Sociologist Daniel Fox ascribed technological optimism to fatigue with the political process and a misguided desire to advance seemingly intractable social problems: "The rejection

[69] One valuable outcome of cultural cognition is the clear recommendation to use identity affirmation in risk communication. That is, information is more likely to be accepted if it can be presented in such a way as to be consistent with cultural values (Kahan, 2010).

of politics among intellectuals often takes the subtler form of what I call technocratic solutionism. Experts who practice solutionism insist that problems have technical solutions even if they are the result of conflicts about ideas, values and interests" (Fox, 1995).[70] Understanding which problems require or are amenable to technological solutions is not obvious. For example, conventional wisdom has attributed famines to food production failures. However, (Nobel economist) Amartya Sen (1981) argued that famines in the past century occurred during times of adequate food production; the real culprits were hoarding and high prices brought about by bad governance and corruption.

None of the ideas presented here make for convincing explanations for attitudes about technological risk (cf. Renn and Benighaus, 2013). However, they do give us a general idea of the range of factors at play in the formation of these attitudes. Given our inability to explain the origins of these attitudes, it may be more helpful to shift our focus from why these attitudes exist to the question of how flexible and variable these technological risk attitudes are.

4.1.4 Variability in technology attitudes

How static are the attitudes of technological optimism and skepticism? This question is important if we are to ascertain whether an event or additional information, such as a risk-benefit analysis, is likely to have a substantial policy impact.[71]

First, it appears that cultural attitudes regarding technology change over time. The early and mid-twentieth century saw technological wonders culminating with a man walking on the moon. Subsequent generations have failed to witness such technological spectacles leading to a

[70] Clearly, he was a technological skeptic.
[71] Influencing attitudes about technology is not as easy as simply presenting new information. Even risk communication meant to reduce concerns can have the unintended effect of raising fears about related risks among skeptics (Nakayachi, 2013).

concern by some that society can no longer solve big technical challenges (Pontin, 2012) and now amuses itself with gadgetry. Nonetheless, the internet age has reignited technological optimism in certain segments of society.[72]

Trends in technological optimism and skepticism can also be traced through science fiction literature. The romantic era of early science fiction, which encompassed the second half of the nineteenth century, envisioned the utopian potential of technology,[73] while the first few decades of the twentieth century, dubbed science fiction's "radium age" (Glenn, 2012), saw the first signs of skepticism that technology was a solution to social ills. Aldous Huxley's *Brave New World*, 1932, is the epitome of the era. More recent trends in dystopian science fiction clearly represent an increase in technological skepticism (and perhaps contribute to it).[74]

While prevailing attitudes about technology have changed over time, they also vary geographically. In a comparison of U.S. and British attitudes regarding analog computing technology prior to World War II, Bowles (1996) argues that U.S. computing technology surpassed British efforts partially due to cultural differences regarding resistance to change and general enthusiasm for technological innovation. Likewise, emerging nations (e.g., South Africa) are considered more optimistic regarding the ability of technology to solve problems than developed nations with strong environmentalist movements (e.g., Germany) (Gruner, 2008).

---

[72] Ray Kurzweil, a famed futurist who predicts that artificial intelligence will exceed any human intelligence by 2030, was hired as Director of Engineering at Google in 2012. His predictions were considered unrealistic 20 years ago, but his track record of success has made him mainstream in Silicon Valley.

[73] For example, Jules Vernes' submarine Nautilus, in *20,000 Leagues Under the Sea*, is a technological marvel used to explore and catalogue nature, aid the oppressed, and oppose militarism.

[74] Some notable technology-induced apocalyptic novels include: Walter Miller's *A Canticle for Leibowitz*, 1959; Kurt Vonnegut's *Cat's Cradle*, 1963; Stephen King's *The Stand*, 1978; and Margaret Attwood's *Oryx and Crake*, 2003.

The variability of technology attitudes within an individual may be just as complex as temporal and geographical trends in society. Unlike the cultural theory of risk, the risk-as-feeling framework allows for variable attitudes about technology within a person. This agrees with our personal experiences where we may encounter individuals who are technological optimists in some fields of science (e.g., the potential for medical technology to improve people's lives) while being technological skeptic in others (e.g., blaming modern telecommunications for loss of privacy).

It also seems reasonable that dramatic personal experience could substantially change an individual's technological risk attitude. Medical technologies, such as deep brain stimulation and artificial joints have become commonplace and can substantially improve quality of life. This not only makes invasive technology more familiar, it can also greatly increase the technological enthusiasm of the person benefiting from the medical device. This appears to be the case for the technology theorist and transhumanism advocate Michael Chorost (2005, 2011) whose hearing was restored by a cochlear implant.

Conversely, technological optimists may evolve into skeptics when their jobs are made obsolete through technology. For example, an online educational system, created by educational technologists, could be used to increase access to education, but also to teach more students with fewer teachers. In this way, educators might innovate their way out of a job. Technology-induced obsolescence of this nature is common (e.g., travel agents, typists, mail carriers, video store and photo lab employees) and the trend is now reaching many of the skilled professions (cf. Rifkin, 1995; Rotman, 2013). That said, policy regarding controversial research is generally made by individuals who have not been substantially displaced by disruptive technology. Furthermore,

many decisions regarding controversial research pertain to low-probability, high-consequence risks that tend to seem more theoretical than real.

4.2 Principle of Controversial Research

So far, we have discussed the origins and flexibility of technological risk attitudes. Despite the lack of a comprehensive theory, we do know that these attitudes appear to be influenced by a variety of factors including: culture, feelings, and personal circumstances. We also know that they are malleable and do change over time at both the individual and societal level. So does this have any implications for how controversial research decisions are made? Based on the complexity of technological risk attitudes, can any general statements be made?

Let us start with the observation that very few lines of research have been banned or abandoned in the past for reasons not related to science or practicality. Many controversial lines of research and technologies, such as the practice of lobotomies, have been abandoned for lack of scientific value and better alternatives.[75] However, the list of scientifically valid research stopped for moral reasons is relatively short and consists primarily of weapons technology (e.g., biological and chemical weapons research (cf. Tucker, 2012) or weather modification for military purposes[76]). Likewise, the list of highly restricted lines of research (e.g., embryonic stem cell research in the U.S.) is also short. Despite substantial ethical concerns or general public opposition, a wide range of controversial research activities are currently legal in the U.S. including: human cloning, cosmetics testing on animals, genetically modifying food, and medical testing on chimpanzees.

---

[75] The inventors of the procedure were awarded the 1949 Nobel Prize in Physiology or Medicine which speaks to our inability to predict the long-term value of research.

[76] This is banned under the international 1978 Convention on the Prohibition of Military or Any Other Hostile Use of Environmental Modification Techniques, a.k.a., Environmental Modification Convention (ENMOD).

In general, the number of technologies and lines of research banned for moral or risk perception reasons is small enough to suggest a general principle of controversial research. Counterexamples, such as bioweapons research, serve as the few exceptions to the general rule.[77]

Principle of Controversial Research

No line of research or technology will be forbidden until it has been scientifically discredited, deemed impractical, or a viable alternative has been accepted. That is, research is rarely banned for moral reasons.

A precursor to this principle is the premise that if something can be done, someone will view it as a legitimate project and try it. This then leads to the principle which posits that we generally do not stop research just because it is ethically controversial. The result is that controversial research generally proceeds until something eliminates the controversy. The controversy can be eliminated in one of several ways: new information or extensive experience reduces the perceived danger, a shift in cultural acceptance decreases opposition to the research, or an alternative technology eliminates the need for the research.

A basis for this principle is a general attitude of inevitability attached to technology. The sense that technology controls us as much as we control technology is describe by social scientists as "technological determinism" (cf. Bimber, 1994).[78] Likewise, recent popular books on the philosophy of technology, such as *What Technology Wants* (Kelly, 2010) and *The Nature of Technology* (Arthur, 2009), describe the "evolution" of technology. Even after specifically

[77] Even bioweapons conforms to the rule to some degree because research continued in secret – on a massive scale in the Soviet Union – even after the research was banned (Leitenberg et al., 2012).

[78] The term is associated with debates regarding the degree to which technology steers social change and how much of it is good or bad.

defining evolution as the development of a system over time, the authors occasionally use the term in ways that connote the theory of biological evolution. By using multiple meanings of evolution, these technologists reveal a mental metaphor of technology as independent and alive and perhaps uncontrollable.

This biological paradigm also provides some insight into attitudes about technology. While technological skeptics might think of "living" technology as ominously virus-like and relentless, technological optimists would see it as benignly natural and inevitable. Although technological optimists embrace the concept of technological inevitability, even technological skeptics see an inexorable pull of technological progress; "we are free at the first step but slaves at the second and all further ones" (Jonas, 1984). The emphasis on describing the trajectory of technological progress also has implications for policy making. With an assumption of inevitability, technological discussions will tend towards predicting what will come next rather than discussing what *should* come next (i.e. technology ethics).[79]

The principle of controversial research appears to be particularly true for emerging technologies. Emerging technologies are noted for rapid change. When the time between basic research and usable technology is very short, the process of societal reaction, evaluation, and public policy formulation is slower, cannot keep pace, and lags behind. Likewise, because the research and technology are often nearly concurrent, available empirical evidence may be of limited value in estimating the future trajectory of the technology (Chameau et al., 2014). In the absence of convincing data or tractable theoretical frameworks, policy makers may revert to

[79] While there are academic communities working in specific subfields of technology ethics (e.g., biomedical ethics), the broader technology ethics community is surprisingly small considering the centrality of technology in modern society. Likewise, the primary focus on innovation literature is how to encourage innovation rather than normative discussions of where innovation should lead.

cautious permissiveness until more compelling evidence is available. That is, there is a pervasive fear of limiting technology development – especially in the face of marketplace competition and international military rivalry. Given the variety in technological risk attitudes, normative arguments rarely count as compelling reasons. The limitations of cultural theory in explaining the nature of technological risk attitudes suggests that these attitudes are distinct from political ideologies. However, the lack of shared values is a feature that opposing technological risk attitudes and opposing political ideologies have in common – appeals to values are largely unsuccessful in resolving debate. To further explore this idea, I next discuss the principle of controversial research in the context of three examples of emerging technologies.

4.3 The principle in action

4.3.1.1 Womb transplants

While organ donation is a well-established lifesaving procedure, a recent example of controversial research is the transplantation of a uterus from a post-menopausal woman to a woman of child-bearing age without a functional uterus. In October of 2014, the first of a dozen successful transplant cases resulted in the birth of a premature, but otherwise healthy child (Brännström et al., 2014). Several bioethicists have noted that this procedure is not medically necessary and is rather dangerous: it requires transplanting not only the womb, but much of the surrounding uterine vasculature during a lengthy surgery; both donor and recipient have a substantial risk for thrombosis and infection; and the recipient is at risk for all the standard concerns of organ transplant rejection (Farrell and Falcone, 2015). While the surgeons involved appeared to be deeply concerned for the safety and mental well-being of their patients who were desperate to experience pregnancy, there was also the unsavory appearance of a race for personal scientific glory to be the first to successfully perform this procedure (Orange, 2014).

The principle of controversial research suggests that this line of research will continue – especially now that it has proven successful. The United Kingdom's Health Research Authority has already granted ethical approval for ten transplants (Johnston, 2015).[80] However, given the expense and availability of less dangerous alternatives (e.g., surrogacy or adoption)[81], it is likely that this procedure will remain uncommon. The procedure is unlikely to be banned on account of ethical concerns, but it may be eventually abandoned due to its impracticality. The true end of the procedure will likely come when scientists are able to grow a new functional uterus from a woman's own stem cells – a technically (and perhaps ethically) superior alternative.

4.3.1.2 Embryo editing

While gene manipulation using recombinant DNA techniques is decades old technology, new techniques[82] have led to proposals to make genetic changes in embryos that are inheritable (i.e., germ line editing). This has raised concerns over potential multi-generational harm as well as a public backlash against "designer babies" (Carmen, 2014). In early 2015, the British parliament approved three-person *in vitro* fertilization (IVF) where a third person provides mitochondrial DNA to correct mitochondrial diseases (cf. Callaway, 2014; Wade, 2015). Some see this as a reasonable small first step, but caution that larger scale genome editing must await confirmation that it is safe (Baltimore et al., 2015). However, even mitochondrial editing is not without controversy. Research suggests that, rather than serving only as cellular power sources, mitochondria influence a range of cellular processes and are involved in multiple diseases

[80] The approved UK clinical trials will use deceased donors to minimize potential patient harm.
[81] It is interesting the note that Swedish surgeons are world leaders in uterine transplant research partially due to the ban on the use of surrogate mothers in Sweden which, ironically, is deemed unethical.
[82] Genome editing, and synthetic biology in general, has been revolutionized since the discovery of the editing technique CRISPR-Cas9 (Ledford, 2015; Marraffini, 2015). Improvements in CRISPR editing systems are progressing rapidly (Ran et al., 2015; Zetsche et al., 2015).

(Hamilton, 2015). Proponents of mitochondrial replacement argue that the long history of inter-racial marriages suggests that there are minimal concerns of mitochondrial mismatch for even the most genetically distant parents. Critics argue that the poorly understood role of mitochondria runs the risk of inadvertently introducing new genetic disease through mitochondrial replacement therapy.

Hesitancy on germ line editing has been compared to the initial concerns regarding IVF (Cyranoski, 2015). This comparison is apt. IVF was initially controversial, and while its safety is still not universally accepted (Servick, 2014), it is a widely used and now standard medical technology. Stronger critics argue that we may be crossing a more fundamental ethical line and that a safer alternative is to use screening, rather than editing, to ensure that parents do not pass genetic diseases to their offspring (Lanphier et al., 2015). In a call for a complete ban on this research, Robert Pollack (2015) represents the strongest critics; "This opening to germline modification is, simply put, the opening of a return to the agenda of eugenics: the positive selection of 'good' versions of the human genome and the weeding out of 'bad' versions, not just for the health of an individual, but for the future of the species."

Some researchers are concerned that germ line editing will cause the public to restrict all gene editing activities. However, history and the principle of controversial research suggests otherwise. There have been very few calls in the U.S. for moratoriums in the biological sciences: recombinant DNA research in 1974, human reproductive cloning in 1997, and influenza gain-of-function research in 2012 (Vogel, 2015). Of these three moratoria over the past 40 years, the recombinant DNA moratorium was lifted within a year (Berg and Singer, 1995). The human cloning moratorium is still unofficially in effect for federally funded research, however, human cloning was never banned despite legislative attempts in 1997, 2001, 2003, and 2005. The

moratorium was originally inspired by a spate of large animal cloning successes (e.g., Dolly the sheep). However, once there was proof of concept, scientific interest led elsewhere. More importantly, once it was found that defect rates were high in the cloned animals, commercial reasons to clone large animals disappeared.[83] Thus, the U.S. moratorium applies to a largely abandoned line of research. The third moratorium is ongoing and will be discussed in more detail in Chapter 6. In summary, as long as there are realizable benefits, research on and use of a new biotechnology will continue.

Current events appear to support this prediction. Chinese scientists have already attempted germ line editing with non-viable embryos obtained from fertility clinics in order to test the potential for curing a fatal genetic blood disorder (Liang et al., 2015). While the study had a low success rate with more unintended mutations than expected, others believe the technology can be quickly improved. Despite divisions, the majority of the scientific community appears to be cautiously supportive of this work (Hinxton Group, 2015; Kaiser and Normile, 2015; Martikainen and Pedersen, 2015; Miller, 2015; Savulescu et al., 2015). Other Chinese teams have already started similar work (Cyranoski and Reardon, 2015).

4.3.1.3 Autonomous Weapons

Perhaps the best counterexample to the principle of controversial research is in the field of autonomous weapons. In the last century, autonomous weapons essentially consisted of pressure-sensitive mines.[84] However, there now are a number of robots and drones equipped with

[83] One notable exception is a Korean laboratory that clones pet dogs. The success rate is low and the use of many donor and surrogate dogs has been called unethical (Stein, 2015). However, the price of $100,000 per dog has created a lucrative niche market that supports the principle of controversial research.

[84] The use of landmines was banned by most nations in the 1997 Ottawa Treaty (Anti-Personnel Mine Ban Convention).

lethal weapons in various stages of development that include fully autonomous operation (cf. Chameau et al., 2014). Proponents argue that autonomous weapons can minimize civilian casualties, while critics are concerned that these systems will lower the social cost of choosing war and will eventually be used universally by both police forces and terrorists (Russell, 2015). Led by computer scientist Noel Sharkey,[85] there has been considerable effort to ban autonomous weapons through the Convention on Certain Conventional Weapons (Stone, 2013).

Autonomous weapons systems may have the best chance of being banned on normative grounds for two reasons. First, a history of international weapons treaties (e.g., the 1968 Non-Proliferation Treaty, the 1972 Biological Weapons Convention, and the 1993 Chemical Weapons Convention) suggest that weapons technologies are unique – there is some consensus that their potential harm clearly outweighs their potential benefits.[86] Second, the 1997 Ottawa Treaty, which bans landmines, creates a specific precedent for banning (albeit simple) autonomous weapons.[87] Nonetheless, countries leading the rapid development of this technology, including the U.S., the United Kingdom, and Israel, have argued against any new restrictions (Russell, 2015).

[85] Sharkey chairs the nongovernmental International Committee for Robot Arms Control. Its mission statement includes the normative premise that "machines should not be allowed to make the decision to kill people" (Stone, 2013).

[86] Along the same lines, the reprocessing of nuclear fuel has been banned in the U.S. since 1976 due to concerns over nuclear proliferation. While the ban pertains to the use of a technology, it has had a chilling effect on related research (including breeder reactors) such that the functional outcome is the same. This suggests that fears of military/terrorist use of a technology are the most effective moral reason for banning a technology. However, it should be noted that calls to lift the reprocessing ban arise every few years and that reprocessing continues in other nuclear countries.

[87] A potential third reason is that every mention of autonomous weapons in the media at some point rather evocatively refers to them as "killer robots."

4.4 Conclusion

In this chapter, we discussed how, in the absence of convincing evidence, technological risk attitudes often guide decision making. Furthermore, because these attitudes are often independent of particular technologies, any single analysis of a controversial technology is unlikely to change risk attitudes and thereby resolve a policy debate. History suggests that there is a general underlying trend in research that can be summarized as a principle of controversial research: questionable research is rarely abandoned for normative reasons. However, when stakeholders attempt to resolve a research controversy with a formal risk analysis, they run into the many epistemic and ethical value assumptions inherent to risk-benefit analysis. The irony is that formal analysis is usually reserved for the most important and controversial technological risk decisions – the very situation in which it is least likely to provide conclusive "answers." In terms of management, this means that controversial research is likely to continue until it has been discredited or a better alternative has been found.

5 MANAGING CONTROVERSIAL RESEARCH

So how might the inherent subjectivities of risk-benefit assessment and the principle of controversial research guide the management of R&D controversies? This chapter will first briefly review some ways of managing research that have been proposed in the past and then evaluate their effectiveness in consideration of the previous discussions. Finally, I propose a series of recommendations for improving the assessment and management of controversial research.

5.1 Common risk management strategies

There are a few general criteria that can be used in risk management (Wilson and Crouch, 2001). First, one can attempt to maximize utility (i.e., seek maximum benefits for the minimum risk) which, ipso facto, lends itself to assessment by risk-benefit analysis. Second, one can attempt to solely minimize risk (zero risk being ideal) which leads to adopting a precautionary approach to technology management which can lead to technology bans. This is a common choice for extreme risks with deep incertitude. It is also a popular approach in Europe (e.g., Rabesandratana, 2015). Lastly, one can take a pragmatic technology-based risk minimization approach. This is often used in pollution control where regulations call for "as low as reasonably achievable" (ALARA) or "best available control technology" (BACT). This approach is often complemented by safety processes and training (e.g., CEN, 2012).

The first and third approaches are useful when data is plentiful. However, these approaches are of limited value for emerging technologies and controversial research. The precautionary approach is more easily applied to these situations, but the fear is that it stifles innovation.

Predicting the future is a popular pastime with generally limited success. Detecting and responding to early signs of risk remains a fundamental issue in risk analysis (Aven and Zio, 2014). More progress has been made regarding how to respond to risk events, but this is less helpful in situations where the consequences are potentially catastrophic. Additionally, even *post hoc* risk reduction is riddled with its own set of biases that prevent realistic reflection and improvement (Paté-Cornell and Cox, 2014). Furthermore, risk management is made more complex by the need to use different strategies for different risks (Renn and Klinke, 2004); science-based strategies (e.g., risk-benefit analysis and emergency management) are appropriate for risks that are well characterized, precautionary management is better for highly uncertain risks, and discursive strategies are best for risks that are well known, but the public consistently underestimates or overestimates.

5.1.1 Engaging multiple perspectives

A common prescription for high uncertainty risk management is to encourage substantial public involvement (Weinberg, 1972). Coeckelbergh (2009) posits that the public is critical to risk analysis because it serves as a source of "moral imagination" (Johnson, 1993) which lets us explore the consequences of our actions and imagine the plight of others. To this end, the U.S. Freedom of Information Act has helped open the process of public review of research (Jasanoff, 2003). However, like all processes, idealized public participation, which includes a well-informed, representative, and heterogeneous opinion, is difficult to achieve (Rossi, 1997). Likewise, increasing the number of stakeholders in a decision process can lead to increased confusion and general intransigence. Thus, public engagement, while widely considered a prerequisite to good risk management, is not guaranteed to be beneficial or sufficient itself.

Reaven (1987a) notes that citizen participation in technical policy decisions is often disappointing in practice because the participants start off overwhelmed by technical information and end up overwhelmed by their inability to judge conflicting expert assessments. Given the variety of epistemic and ethical value assumptions that are inherent to the risk assessment process, evaluating disputes in technological risk assessments is difficult enough for scientists, let alone the general public. One way to mitigate these issues is to provide "science literacy mini-courses" to public participants that discuss the types of scientific disputes and the contingent nature of science and scientific consensus (Reaven, 1987a).

As Reaven notes, one could also argue that the need for scientific literacy training extends to the scientists involved in research controversies. While most working scientists have formal training in methods and conduct of science, very few are educated as to the main schools of thoughts about its philosophical underpinnings. This may be why many scientists are ambivalent on questions about the nature of their work and why some scientists are even dismissive when these issues arise (e.g., Meyer, 2011). The roadmap of value assumptions detailed in Chapter 3 provides one potential framework for discussing the important assumptions in technological risk assessments with both scientists and the general public.

It is also important to engage the internal scientific community (Yarborough, 2014). That is, to ensure accountable science, dissenting scientists must have open channels for voicing concerns where they need not fear retribution. Likewise, scientists should be encouraged to reflect on their work rather than assuming that the existing regulatory process (e.g., institutional review boards) will catch any ethical or safety issues. However, as discussed in the second chapter, the attention of experts tends to be concentrated on their area of expertise. Scientists will have a natural tendency to permissively promote science while security experts will tend to

conservatively restrict perceived risks (Selgelid, 2009). Because the science research community is largely self-regulated (with notable exceptions in nuclear research), this tends to reinforce the general principle of controversial research. Although a robust dialogue within the scientific community is healthy, Sarewitz (2015b) argues that "risk is more a political and cultural phenomenon than it is a technical one." Thus, he cautions against letting unelected "privileged experts" (i.e., scientists, social scientists and ethicists) dictate the terms of technological risk debates.

The sphere of dialogue in risk management can also be expanded by consulting literature. The value of literature to improve empathy and moral imagination has long been argued (e.g., Nussbaum, 1990) and more recently supported by research (Bal and Veltkamp, 2013; Kidd and Castano, 2013). Philosopher Hans Jonas (1984)[88] called for "a science of hypothetical prediction" while still recognizing the impossibility of "preinventing" technology and the realization that the reach of modern technology often exceeds the reach of our of foresight. He argued that one of the best options is to cultivate an imaginative technological skepticism and suggested science fiction[89] as a valuable source of inspiration for exploring the possible and laying the groundwork for risk management. In the case of technological risk analysis, science fiction may provide fertile ground for moral imagination – that is, a way to play out various future technological scenarios and explore their ethical implications. Likewise, this mode of exploration may also suggest risk management options.

[88] Jonas (1984) suggested one of the earliest forms of the precautionary principle, "It is the rule, stated primitively, that the prophecy of doom is to be given greater heed than the prophecy of bliss."

[89] Experts who might take umbrage at the idea of consulting science fiction regarding future uses, misuses, and unintentional consequences of technology might be more at ease if the works were rebranded as "technological eschatology" or something equally pretentious.

5.1.2 Critique of traditional risk management

Despite our best efforts, history suggests that humans tend to be overconfident in their assessment of risk and good at rationalizing inadequate risk management after the fact (Paté-Cornell and Cox, 2014). Some of these excuses for poor risk management include: uncontrollable forces ("acts of God"), unimaginable events ("black swans"), rare confluences of events ("perfect storms"), lack of precedent, excusable ignorance, conformance to standard practice, someone else's responsibility, lack of legal liability, or even operator error. Although risk assessments are useful, unanticipated events with no historical precedent occur with such frequency that taking unrecoverable risks[90] should be avoided (Taleb et al., 2009). Aven and Krohn (2014) suggest adopting a "mindfulness" approach to risk management which includes the following characteristics: learning from past failures and anticipating future failure; avoiding simplification of complex systems; being sensitive to operations; fostering resilience, and deferring to expertise (not necessarily authority). This approach does not differ appreciably from a more traditional risk analytic approach (e.g., Wilson and Crouch, 2001) except in that it emphasizes the incompleteness of any formal risk assessment and promotes a general skepticism to quantitatively bounding the certainty of our knowledge.

A common theme in modern risk management is to increase the resiliency of a system in the face of deep uncertainty (Linkov et al., 2014). Physical systems can be made more resilient by increasing the redundancy or independence of functions. But how do we make a line of research or technology more resilient when the potential risk is in the outcome rather than the

[90] The idea of avoiding unrecoverable risks is common in our personal lives. The insurance industry depends on the natural preference to avoid extreme losses if given the opportunity. For example, most people do not calculate the risk of their house burning down before buying homeowners insurance (Taleb et al., 2009).

system itself? We could try to make society more resilient to the effects of research and emerging technologies, but it is not clear that this is possible. It is also a rather unsettling approach if the outcome may be catastrophic. A more acceptable and practical response to a potentially fatal technology is to attempt to reduce the risk rather than make society more resilient (e.g., by increasing population or by geophysical barriers such as underground bunkers or space colonies). However, in light of the essential inevitability of technological progress as proposed in the general principle of controversial research, it would seem that technological risk can only increase. What are we to do?

5.2 Inherent safety

Even if we accept the idea that controversial research will usually continue until it is discredited or something better comes along, we need not accept that research must be inevitable *and* unfettered. Between the extreme positions in risk management of outright banning technologies via the precautionary principle and complete permissiveness due to technophilia or fatalism, lie the traditional methods of risk management which fail to resolve controversy. This section describes the concept of "inherent safety" as a more satisfying, moderate risk management strategy.

The principles of inherent safety were first developed in the chemical industry (Kletz, 1978). If we define an accident as bad effects resulting from the combination of a material hazard and an instigating event (Srinivasan and Natarajan, 2012), then inherently safe design attempts to reduce the bad effects by reducing or eliminating the hazard. In contrast, conventional risk management generally focuses on reducing the *likelihood* of the instigating event that causes an accident through risk-reducing procedures added to an existing or near-complete process. The distinction between reducing the hazard versus reducing the likelihood of

the bad event has also been termed primary and secondary prevention, respectively (Hansson, 2010a). By selecting safer materials during the early design phase, overall risk is reduced far more than by merely adding on safety systems or safety procedures as an afterthought (Kletz, 1978). The basic principles of inherent safety consist of (1) minimizing the total amount of materials used, (2) substituting hazardous materials with alternatives that do not have the hazardous quality, (3) using materials under less hazardous conditions, and (4) simplifying processes to reduce errors (Kletz, 1985; Khan and Amyotte, 2003; Edwards, 2005). The formal inherent safety concept is frequently used in the chemical (cf. Khan and Amyotte, 2003; Zimmerman and Anastas, 2015) and nuclear engineering communities,[91] but has not been widely adopted by scientists and engineers in other areas (Srinivasan and Natarajan, 2012). Inherent safety has been described as common sense, but not common knowledge (Khan and Amyotte, 2003).

Inherent safety is one of many techniques within the field of safety engineering, but it is fundamentally different from the others in that they focus on secondary prevention (Hansson, 2010a). These other approaches include (Möller and Hansson, 2008): safety reserves, where reserve capacity is included in the design (e.g., safety margins and safety factors); fail-safe design, where a system is designed to fail in a way that minimizes damage (e.g., self-shutdown and safety barriers that isolate failure); and procedural safeguards (e.g., warning systems, safety audits, and safety training). These traditional approaches are applied to a variety of situations ranging from the mundane to existential threats (e.g., Jebari, 2014).

[91] In the nuclear engineering community, there is a different emphasis on the basic inherent safety principles due to the difficulty of reducing the hazardous material which is fundamental to the nuclear power process. Rather than reducing the material, the emphasis is on reducing hazardous conditions (e.g., designing systems that are incapable of reactor core meltdowns).

Another quality that sets the inherent safety concept apart from other risk management methods is its ability to address security concerns (Hansson, 2010a). For example, a secondary prevention measure, such as removing all ignition sources from proximity to an explosive agent, is of little security value; malevolent actors will bring their own ignition source. Likewise, terrorists are attracted to hazards that already instill public dread. Inherently safe design makes terrorism more difficult by removing or minimizing the exploitable hazard.

Because safety has traditionally been the concern of engineers at the production level, the R&D community often fails to consider these principles in the early stages of research when the most impact can be made (Edwards, 2005). However, inherent safety in research is sometimes recognized in hindsight. A Centers for Disease Control and Prevention report (CDC, 2014) summarizing an internal review of the June, 2014 exposure of laboratory workers to potentially viable *B. anthracis* at a CDC bioterrorism response lab, noted that a low-risk non-pathogenic strain could have been substituted in the experiment.[92] It is also interesting to note that in its list of responses, the report focuses on secondary prevention – primarily in the form of revised biosafety protocols and procedures. A reference to reducing the hazard (i.e., inherent safety) is made only within the fifth of eight recommendations. It appears that the philosophy of inherent safety has yet to be embraced in the biological sciences.

The findings of the report also reinforce two previous arguments. First, even at facilities with multiple safety barriers and extensive training, unanticipated errors do occur. This suggests that traditional risk management solutions will not allay public concerns regarding potentially

[92] The experiment was testing whether a new mass spectrometer could quickly detect *B. anthracis*. However, the instrument manufacturer stated that the instrument could not distinguish between virulent and avirulent strains of the same species. Thus, using a non-pathogenic strain would have yielded the same results.

dangerous research. Second, considering safety in the research design phase can often accomplish the same scientific goals while sidestepping controversy.

Edwards (2005) outlines several reasons why inherent safety has been slow to catch on in industry. Ironically, the main reason for the slow adoption is risk aversion. Businesses are hesitant to fix things that do not appear to be broken[93] and do not want to incur the potential risks of deviating from tradition and known practices. Costs may be incurred by redesigning processes or efficiencies may decrease. Speed-to-market and market share may be jeopardized by introducing new ideas.

These concerns also constitute potential roadblocks to using inherent safety principles in controversial research. By its nature, research is full of new ideas, but researchers often use well-established (and peer-reviewed) techniques in the race to new scientific findings and innovation. Thus, scientists may be just as hesitant to incur the extra expense and lost time needed to institute inherent safety principles. Researchers who are focused on work efficiency may view safeguards that are already in place as onerous and may see any additional requests as superfluous. However, public pressure and threats to funding may lead scientists engaged in controversial research to view inherent safety as a reasonable compromise.

Inherent safety, like all human endeavors, is limited by knowledge and creativity. That is, in order to invoke inherent safety, one must realize that a planned course of action is potentially unsafe. Likewise, inherently safe alternatives are not always obvious and may require considerable innovation. However, engaging multiple perspectives can help offset limitations of imagination. Ultimately, the primary value of inherent safety is that it provides a complementary

[93] This is a common attitude in risk management (and most human activities). Even systems with many "near misses" are considered to be safe until something too egregious to ignore occurs (Paté-Cornell and Cox, 2014).

philosophical approach to standard probabilistic risk analytic thinking which often treats the hazard as a given (Johnson, 2000).

An inherently safe design has a specific aim; to reduce a perceived risk. Yet, a design change can create a new unforeseen risk that did not previously exist. To avoid trading one potential risk for another, unrecognized risk, efforts to create more inherently safe research benefit from thorough risk exploration using the many tools of risk analysis.

5.3 Recommendations

In Chapter 2, I reviewed several approaches to valuing research and technology development activities. These approaches are interrelated, yet can yield disparate assessments. Likewise, in Chapter 3, I reviewed some of the many value assumptions inherent in risk assessments that hinder claims of objectivity and neutrality. These lines of argument detail why risk-benefit analysis is more appropriate for evaluating well-defined problems with copious data (e.g., many business and engineering decisions), but has limited utility for resolving debates over controversial research where uncertainty is high and data are sparse. Does this mean that risk-benefit analysis (RBA) has no place in assessing and managing controversial research? No, it means that, RBA, as commonly used, is ineffective as a primary decision criterion. However, by following a few heuristic recommendations, we can considerably improve the usefulness of RBA in assessing and managing controversial research.

Recommendation 1: Change expectations.

Expecting all RBA to yield clear-cut, consensus-building answers is unrealistic. The controversies surrounding some lines of research are based on fundamental disagreements over ethics and public risk that will not change due to the findings of a few assessments. Certain issues are more amenable to RBA than others for reasons already noted. For example, we would

expect that an RBA dealing with automotive safety would be far less controversial than one for euthanasia techniques.

Furthermore, any request to conduct an RBA will not be accompanied by a well-defined procedure because there is no consensus regarding what constitutes the proper accounting of benefits or what underlying value assumptions should be used to estimate risk. By focusing on the "correctness" of an assessment, stakeholders lose an opportunity to use the assessment productively as a risk exploration tool. The simple solution is to make sure that all stakeholders (i.e., policymakers, analysts, scientists, and the general public) understand that formal assessments are part of the decision process, not the entirety (or even the largest factor) of the decision process.

Recommendation 2: If a quantitative assessment is desired, ask narrow questions.

Although it may be inappropriate for stakeholders to expect definitive "answers," we must be realistic. The effort of generating a time-consuming and costly RBA carries an expectation that the results will translate into a decision. Another way to avoid disappointments, besides by moderating expectations, would be for assessments to address narrower questions that are more likely to be "answerable." Narrower questions are also more comprehensible and comparable.

As discussed in Chapter 3, there is a considerable tradeoff between broad, comprehensive assessments and narrow ones. However, if we follow the previous recommendation, a narrow assessment need not address all concerns, but rather becomes one more piece of evidence within the broader scientific and political debate. In this sense, an RBA can successfully generate an answer to part of the debate rather than fail to provide the answer to the entire debate.

Recommendation 3: Use broad uncertainty assumptions.

Even for RBAs that address narrow questions, it is important that the assessments have a broad conception of the uncertainty inherent in the work. This includes acknowledging that: not all uncertainty can be represented probabilistically without making important assumptions; dependencies among parameters are not always well characterized; and the correctness or completeness of the model is often uncertain. The value assumptions roadmap outlined in Chapter 3 can be used as a guide. Even in a highly quantitative assessment, broader notions of uncertainty can still generate informative results (e.g., Rozell and Reaven, 2012).

Recommendation 4: Use multiple methods.

When resources allow it, an RBA will benefit from using multiple techniques developed from risk analysis (e.g., using both a Monte Carlo simulation and probability bounds analysis). The work in Chapter 2 and Chapter 3 provide guidance on the range of methods available. The value of this approach is to give proponents of various statistical techniques (based on their epistemic preferences) a chance to feel confident about the results of the assessment. Likewise, individuals without any methodological preference will be comforted by the thoroughness of the analysis and the range of analytical frameworks.

Recommendation 5: Use the RBA to design better research.

In addition to any narrow quantitative assessments performed, a broad qualitative assessment should also be performed. Used in conjunction with the previous recommendations, the RBA should be useful for exploring risk. This can also lead to rethinking the purpose of the research and seeing if the same ends can be accomplished by inherently safer or more ethical means. By and large, the best opportunity to resolve a debate over controversial research is to apply the "inherent safety" design framework when possible.

5.3.1 Example: Gene drive risk assessment

The following example applies these recommendations to a recent risk assessment for controversial research. The research involves the proposal to eliminate some species of mosquito using a "gene drive." The general concept is that a species of mosquito can be driven to extinction by releasing genetically engineered males into a wild population. The engineered males are encoded with homing endonuclease genes (HEGs) – genetic elements that copy themselves into chromosomes in which they do not already appear (Burt and Koufopanou, 2004). In this case, the HEG generates a version of an enzyme, I-Ppol, (originally found in a slime mold) which has been engineered to only be expressed during sperm development and to cut and destroy X chromosomes (Galizi et al., 2014). The result is that male mosquitoes only pass on Y chromosomes and the offspring of these engineered males are entirely male. This male-dominant trait is passed on to subsequent generations and eventually the species is driven to extinction when there are no more females.

The benefit to humans of eliminating the primary vector for a multitude of diseases (e.g., malaria, yellow fever, and dengue fever) is obvious. According to WHO estimates, malaria alone results in approximately half a million deaths (mostly children) each year. However, any proposal to drive a species to extinction is inherently controversial. While many biologists believe that the ecological niche occupied by mosquitoes would be quickly filled by other insects (Fang, 2010), mosquitoes provide an ecological service as pollinators and are an important food source for many species (Poulin et al., 2010). The substantial biomass that mosquitoes comprise in some arctic and aquatic environments hints at the potential impact of their absence. Likewise, the concern that a gene drive mechanism might be transferred to another species and drive an unintended target to extinction is reason enough to proceed with great caution.

A report (Hayes et al., 2015) from the Australian Commonwealth Scientific and Industrial Research Organization (CSIRO) assessing the mosquito gene drive research exhibits some of the qualities of a good risk assessment as described in the previous recommendations. First, the assessment follows the second recommendation by limiting itself to a relatively narrow task - to "identify and quantify the ecological and human health risks associated with an accidental release of the I-PpoI mosquitoes from African insectaries in Burkina Faso, Mali and Kenya." The report does not attempt to identify or quantify all risks or benefits associated with gene drive technology and should be seen as a small piece of a much larger debate.

Likewise, the report assesses a more inherently safe version of the gene-drive technology. Early versions of the I-Ppol HEG inadvertently created fully sterile male mosquitoes (Windbichler et al., 2007). This creates the need to continually add genetically engineered males to the wild population in order to decrease total population. While this would be an economically unfeasible mosquito control mechanism in most regions where malaria is endemic, the self-limiting nature of this HEG variant is useful for early stage research.

The report also follows the fourth recommendation by using multiple methods to identify and quantify the potential hazards. A Boolean literature search was conducted to determine prior efforts to identify or quantify the risk of gene drives. Likewise, hierarchical holographic modeling (Haimes, 1981; Haimes et al., 2002) was used as a complementary approach to identify potential hazards. The basic idea of hierarchical holographic modeling is to create a detailed matrix of all the potential interactions of processes within a complex system. This matrix can then serve as a systematic guide for inductively investigating risk. Finally, fault tree analysis (Vesely et al., 1981) was used to deductively identify and quantify risks. Additionally, two methods were used for combining expert opinion within the fault tree (i.e., "Aggregate First

Then Convolute" (AFTC) and "Convolute First Then Aggregate"(CFTA) (Hayes et al., 2015, 40)).

The CSIRO report has mixed success at following the third recommendation, to use broad conceptions of uncertainty. The assessment does consider a range of risk estimates based on assumptions of both independence and strong positive dependence between model parameters. However, the assessment assumes the independence of expert opinion and does not consider interval uncertainty to express incertitude.

If one follows the fifth recommendation, the purpose of the RBA becomes risk exploration and communication. In this respect, the CSIRO report is less than ideal as a risk communication tool. First, the report is very technical and requires advanced knowledge of biology and risk analysis to comprehend. If the report is to be understood by non-technical policymakers or the general public, more introductory material is needed. Even the executive summary, where the most essential points of a report are distilled, is full of technically precise terms, such as "eukaryotes" and species "complex," which are jargon to non-biologists. This is unfortunate because the authors have unintentionally ceded control of risk communication to the technical policy advisors who will be needed to translate the material into plain English.

Likewise, the report may be giving an unintended false impression of precision by its combined use of Monte Carlo sampling of elicited probability density functions and linear pooling of expert opinion to create a quantitative assessment. It is understandable that the CSIRO report used expert opinion for model quantification due to the practical limitations of observing some of the modeled phenomena (i.e., it is "trans-science" (Weinberg, 1972)). However, unnecessary complexity of analysis is discouraged. As Casman et al. (1999) note, "a prescription that one's analytical formulation should grow in complexity and computational intensity as one

knows less and less about the problem, will not pass the laugh test in real-world policy circles." The CSIRO report authors have drawn a somewhat arbitrary line between qualitative and quantitative analysis; "Quantitative risk estimates in our opinion are essential to meet science quality criteria, [94] and we achieved this by using direct elicitation to fit subjective probability density functions to expert beliefs about the probability of the events in fault trees, and for a set of well established vectorial capacity parameters" (Hayes et al., 2015, 97). While the goal of quantitative assessment is laudable, the public often conflates the quantitative/qualitative distinction with the empirical/subjective distinction. Particular care is needed to properly communicate to users of the report that it is a quantitative risk assessment using subjective probabilities with difficult-to-estimate reliability.

5.4 Conclusion

This chapter reviewed standard suggestions for risk management of controversial research and their general shortcomings. Based on the previous chapters, five recommendations were made and an example was used to illustrate the recommendations and how they can improve RBA. Of the many risk management methods, the most universally acceptable way forward is to attempt to remove the characteristic of the research or technology that is causing the controversy while retaining the essential elements. In the next chapter, the ideas laid out so far will be applied to the example of controversial research presented in the first chapter – highly pathogenic avian influenza gain-of-function research.

[94] According to the authors, "A scientific risk assessment should be transparent, repeatable and wherever possible make predictions that are measurable and falsifiable."

6 THE GAIN-OF-FUNCTION RESEARCH DEBATE

In this final chapter, I revisit the influenza gain-of-function (GOF) research controversy. The history of the controversy and resulting moratorium was outlined in the first chapter. Here I apply the arguments presented in the previous chapters to this controversy with special attention given to the risk management proposals made by the scientific and science policymaking community to resolve the controversy. Finally, I summarize findings and draw some general conclusions that may be helpful to those attempting to manage controversial research.

6.1 Timeline

The moratorium instituted on October 17, 2014 by the U.S. Office of Science and Technology Policy called for "a robust and broad deliberative process" to "evaluate the risks and potential benefits of gain-of-function research with potential pandemic pathogens" (OSTP, 2014). The proposed multi-step process consists of a series of meetings by the NSABB and National Research Council (NRC) to draft recommendations on how a risk-benefit analysis (RBA) should be conducted followed by an evaluation of an independently performed RBA (Figure 5). The primary role of the NRC meetings is to provide additional feedback from the scientific community while the NSABB provides the final recommendations to the Secretary of HHS and Director of the NIH. The initial NSABB meeting was conducted on October 22, 2014 (NSABB, 2014b) followed by a conference call on November 25, 2014 (NSABB, 2014a). The subsequent NRC symposium occurred on December 15-16, 2014. A May 5, 2015 NSABB meeting was held to discuss and adopt the framework of the independent RBA. Another NSABB meeting to discuss ethical issues surrounding GOF research and receive an update on the ongoing RBA was held on September 28, 2015. As of this writing, the next NSABB meeting is scheduled for January 7 and 8, 2016.

A $1.1 million contract, administered by the NIH Office of Science Policy, was awarded to an independent Maryland biodefense consulting firm, Gryphon Scientific, on March 11, 2015 to conduct the RBA within 12 months (NIH, 2015). The original notice (NIH, 2014) requested that any RBA "be comprehensive, sound, and credible" and use "established, accepted methods in the field."

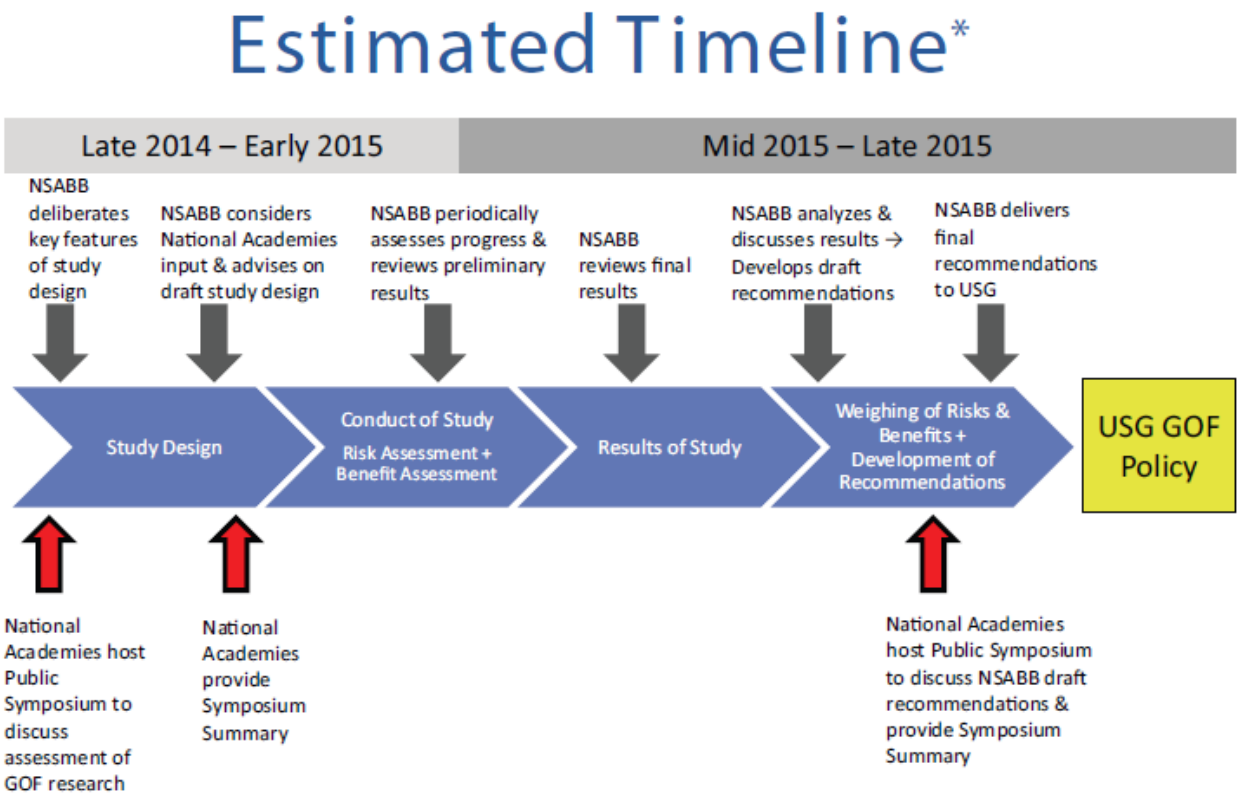

Figure 5: NIH proposed deliberative outline (from Sharples et al., 2015)

6.2 The Risk-Benefit Assessment

Evans (2013) summarizes the most commonly mentioned risks and benefits of influenza gain-of-function research. The primary risks are that: (1) a virulent engineered virus could be accidentally released causing a pandemic, and (2) an engineered virus could be intentionally released in an attempt to start a pandemic. The potential benefits include: (1) increased understanding of influenza virology, (2) improved influenza pandemic surveillance, and (3) more

effective and faster-to-market vaccines. There are also additional risks attached to the debate that are not specific to the research. That is, any restriction of research is accompanied by loss of academic freedom, chilling effects on related research, and the potential that the research will be ineffective.

An important distinction regarding the potential benefits of GOF research is that it is conditioned on factors outside the laboratory (Evans, 2013). That is, while the risk of accidental release is largely controlled by laboratory conditions, the beneficial use of any knowledge discovered depends on the public health system which varies widely among communities, regions, and nations. For example, one year into the 2009 influenza pandemic, there was still only enough vaccine for a quarter of the world's population (Stöhr, 2014). Given the uneven distribution of basic public health services in the world, any research benefits will be far more limited in extent than in an ideal world. The problem is further exacerbated by the frequent regression of public health services in regions experiencing war and failed governments (Hotez, 2015). In summary, the concern is that the potential risks are more immediate and wide-spread than the potential benefits.

While procedures for performing biosafety and biosecurity risk assessments for dangerous pathogen research have been previously formalized (e.g., Caskey et al., 2010), the NSABB (2012a) made a first attempt at RBA guidelines for the avian influenza controversy (Figure 2). However, the emphasis at the time was specifically focused on potential misuse and the decision whether to communicate or censor research results. The framework does not consider accidental laboratory release – a concern just starting to gain widespread attention at the time it was written (e.g., Berg, 2012). Thus, the NIH proposal for a comprehensive RBA for gain-of-function research has no obvious precedent to emulate.

COMMUNICATING DUAL USE RESEARCH OF CONCERN: RISK/BENEFIT ANALYSES

1 Begin — Risk Analysis

A *Are there reasonably anticipated risks to public health and safety from direct misapplication of this information*, i.e., is novel scientific information provided that could be intentionally misused to threaten public health or safety?

B *Are there reasonably anticipated risks to public health and safety from direct misapplication of this information*, i.e., does the information point out a vulnerability in public health and/or safety preparedness?

C Is it reasonably anticipated that this information could be directly misused to pose a threat to agriculture, plants, animals, the environment, or materiel?

D If a risk has been identified, in what timeframe (e.g., immediate, near future, years from now) might this information be used to pose a threat to public health and/or safety, agriculture, plants, animals, the environment, or materiel?

E If the information were to be broadly communicated "as is," what is the potential for public misunderstanding, that is, what might be the implications of such misunderstandings (e.g., psychological, social, health/dietary decisions, economic, commercial, etc.)? For sensationalism?

2 Pause to consider

In some very rare cases, the risks associated with misuse of information from dual use research of concern are so significant that no amount of potential benefits can outweigh the risks. In such cases, the decision would be *DO NOT COMMUNICATE.*

The conditions under which this could be the case is that the research yields sufficient information for bad actors to pose threats that:

- **Would cause substantial harm/severe impact**
- **Pose risk to large populations**
- **Require little or no additional information**
- **For which there are no countermeasures or only inadequate countermeasures in terms of efficacy or availability**
- **Require only readily available materials**
- **Require low levels of expertise or technology to execute**
- **Can be realized in the immediate or near future**

If this is <u>not</u> the case, then complete the risk/benefit analyses by resuming with steps 3A through 3D and step 4.

3 Resume — Benefit Analysis

A Are there potential benefits to public health and/or safety from application or utilization of this information?

B Are there potential benefits of the information for agriculture, plants, animals, the environment, or materiel (e.g., what potential solution does it offer to an identified problem or vulnerability)?

C Will this information be useful to the scientific community? If so, how?

D In what timeframe (e.g., immediate, near future, years from now) might this information be used to benefit science, public health, agriculture, plants, animals, the environment, or materiel?

4 Based on completed risk/benefit analyses and using best professional judgment, consider options and make a decision

Options

Communicate with specific conditions:

- **Content (as is or with additions and/or deletions)**
- **Timing (immediately, only after certain conditions are met, etc.)**
- **Distribution (broad, restricted, etc.)**

OR

Do not communicate

Figure 6: RBA guidelines (modified from NSABB, 2012a)

The NIH (2014) notice for an independent third-party[95] RBA included the following requested elements:

- updated literature on key pathogen characteristics needed for qualitative and quantitative assessments, such as dose-response;
- a comprehensive list of potential accidents and events;
- a detailed analysis of a sub-group of risk events;
- estimates of event frequencies based on available data and extrapolations;
- estimates of consequences of exposure events;
- a qualitative and quantitative risk assessment for selected pathogens and events that includes sensitivity and uncertainty analyses and addresses the following issues: biosafety, physical and personnel security, proliferation, information risk (e.g., publication risk), international relations, national security, and environmental risk; and
- a qualitative and quantitative benefits assessment that includes identifying less risky alternatives with similar benefits that addresses the following issues: the value of information of the research and improvements in surveillance, therapeutics, diagnostics, vaccines, international relations, and national security.

[95] An independent third-party analysis should not be confused with an impartial analysis. Analysts are human and may be influenced by the hiring organization, social pressures, and cultural norms. Likewise, there is a trend towards specialization among analysts as they are expected to quickly ascertain all the relevant information, key stakeholders, and alternatives (von Winterfeldt and Edwards, 2007). Analysts trained and embedded in the field of interest are more likely to have a strong personal position at the outset and in a small research community it is even more difficult to find someone both qualified and independent.

The many value judgments inherent in the RBA process have been previously described as well as the importance of not using RBA as a method to avoid the responsibility of making difficult decisions (see Chapter 3). This was well summarized by Baruch Fischhoff, an NRC symposium planning committee member, "Anybody who thinks that putting out a contract for a risk-benefit analysis will tell the country what to do on this topic is just deluding themselves" (Sharples et al., 2015).

Another difficulty of assessing and managing GOF research is the broad range of activities the term includes. The primary concern is research using potentially pandemic pathogens that may result in increased virulence, transmissibility, or pathogenicity. However, as discussed at the NRC symposium (Sharples et al., 2015), gain-of-function research is already widely used in virology for multiple beneficial and largely benign purposes. Some examples include: increasing vaccine yields (Schultz-Cherry et al., 2014), expanding genomic sequence surveillance databases (Davis et al., 2014), and creating animal models of human viruses to aid further research. Furthermore, naturally occurring gain-of-function mutations are common in research labs that work with RNA viruses. Thus, banning GOF research leading to increased virulence, transmissibility, or pathogenicity is difficult to accomplish in practice without casting a wide net. Because the current state of science is unable to predict what genomic changes will increase danger, it is not possible to foresee what experiments will result in new undesirable traits.

However, proponents of the moratorium point out the wording was specific enough that only 18 federally funded projects were affected and that public health surveillance and vaccine development activities were exempt (Lipsitch and Inglesby, 2014). Rather, critics of the

moratorium may be attempting to “muddy the waters” by widening the definition of what might be banned in hopes of weakening support for any restrictions.

This is not the only debate over terminology. It has been argued that the use of the term “pandemic” is an “apocalyptic rhetorical device” (Casadevall et al., 2014b, 2014d) that preempts any reasonable discussion of risks and benefits by appealing to our innate fear of rare, but catastrophic events. However, this assumes the risk of pandemic is rare despite considerable disagreement among informed scientists regarding the likelihood of such an event. Ironically, labeling the use of "pandemic" as rhetorical sophistry is itself a rhetorical trick because it dismisses a category of serious claims without due consideration of merit.

6.2.1 Quantitatively estimating the risks and benefits

While the NIH requests both qualitative and quantitative risk-benefit assessments if possible, the general expectation among the scientists involved is that a quantitative assessment should be performed (e.g., Lipsitch and Inglesby, 2014) under the assumption that numbers carry more credibility and may suggest solutions. Klotz and Sylvester (2012) performed a simple probabilistic risk assessment to argue that research with potential pandemic pathogens (PPPs) should be restricted. Starting with the assumption that the probability of release from a laboratory each year was 0.003 and that at least 42 labs worldwide were working with PPPs (in this case, highly pathogenic influenza, Middle East respiratory syndrome (MERS), or severe acute respiratory syndrome (SARS)), there is an 80 percent likelihood of a release every 13 years. An updated estimate (Klotz and Sylvester, 2014) finds the risk of a pandemic over ten years due to laboratories working on PPPs is between 5% and 27%.

Using BSL-3 lab infection data from Henkel et al. (2012), a simulation model for transmissibility, and historical pandemic data, Lipsitch and Inglesby (2014) estimate a

probability of creating a pandemic of between 0.01% and 0.1% per laboratory-year which would cause between 2 million and 1.4 billion fatalities. This yields an expected fatality rate of 2,000 to 1.4 million per BSL-3 laboratory year.[96] Alternatively, if using data from the National Institutes of Allergy and Infectious Diseases, the probability of pandemic would be between 0.05% and 0.6% per worker-year with a resulting expected fatality rate of between 10,000 and 10 million per laboratory worker. When Lipsitch presented these calculations at the NRC symposium, Ron Fouchier, the lead researcher for the original controversial H5N1 paper, responded, "I prefer no numbers rather than ridiculous numbers that make no sense"[97] (Sharples et al., 2015).

A subsequent risk estimate from Fouchier (2015b) starts from the same data (Henkel et al., 2012), but then argues that highly pathogenic viral experiments occur in special facilities (BSL-3+) and, using the Erasmus MC facility as an example, estimates that risks are much lower due to extra physical barrier biosafety measures, lab personnel vaccinations, and available anti-viral therapeutics. Thus, he estimates the risk of a laboratory acquired infection to be less than 1 x 10^{-7} per person-year. Taking into account that any infected lab worker would have already been vaccinated against a homologous H5N1 virus, would be taking anti-viral medication, and would be quarantined, Fouchier estimates a lab-induced pandemic would occur every 33 billion years.[98] Fouchier concludes with the observation that there have been no confirmed laboratory acquired influenza infections or releases in decades which suggests current measures are sufficient.

[96] As discussed in Chapter 3, expected utility (or disutility) can be an unhelpful way to express low probability-high consequence events, but that doesn't seem to stop anyone. The economists seem to have won this battle so far.

[97] One must acknowledge the boldness required for a virologist to make such a retort to a Harvard epidemiologist.

[98] This is more than twice the known age of the universe. Thus, the calculated risk for all practical purposes is zero.

A reply by Lipsitch and Inglesby (2015) questions Fouchier's claim that virology labs are safer than other BSL-3 labs. They also note that Fouchier's calculations incorrectly account for the uncertainty associated with zero observed events.[99] Furthermore, the assumption of zero events is unreasonable because viral laboratory-acquired infections have occurred in non-U.S. facilities (cf. Gaudioso et al., 2009) and underreporting is common. In separate comments, Klotz (2015) argues that Fouchier's calculations use the wrong method of calculating the elapsed time of escape for a laboratory-acquired-infection and that his estimate of less than one laboratory-acquired infection in one million years is too low. A reply by Fouchier (2015a) argues that Klotz does not provide "scientific justification" for higher estimates.

Within this debate among competing risk estimates, there appears to be a disagreement as to not only what constitutes the appropriate methodology, but also what constitutes evidence. For example, Fouchier (2015a) does not believe that recent laboratory errors (most notably at the CDC) constitute relevant data because: the errors did not result in laboratory-acquired infections, the pathogen was not an engineered avian influenza virus, and the work was not conducted specifically in a BSL-3+ laboratory. However, critics contend that these errors demonstrate the general failure of laboratory safety procedures that Fouchier's calculations depend upon. That is, estimates that rely on all systems functioning as designed with no safety measures being bypassed by unforeseen circumstances is an incredibly optimistic assumption unsupported by history.[100]

[99] Lipsitch and Inglesby reference the standard "rule of three" approach (e.g., Eypasch et al., 1995) where the 95% confidence interval is [0,3/*n*] when no events are recorded in *n* observations. However, this approach is not universally recommended (cf. Winkler et al., 2002; Quigley and Revie, 2011).

[100] See the section on "inherent safety" in Chapter 5.

A similar debate over evidence arose when Rozo and Gronvall (2015) proposed that the 1977 flu pandemic was caused by a vaccine trial or vaccine development accident rather than a research laboratory release; they concluded that, "it remains likely that to this date, there has been no real-world example of a laboratory accident that has led to a global epidemic." A critique by Furmanski (2015) argued that this nuanced position was largely irrelevant to the GOF debate. Since vaccine development is still a primary goal of GOF research, a vaccine mishap is no less worrisome than a research lab release – an epidemic could occur regardless of which lab made the fatal error.

This doubt is corroborated by a previous study that estimated a probability of between 5% and 15% that a laboratory escape event would not even be detected (Merler et al., 2013). There is also considerable doubt regarding the availability of relevant data for a realistic biosafety risk assessment. An investigative report of U.S. biolabs (Young and Penzenstadler, 2015) argues that a reliable biosafety assessment is not possible for many reasons, including: the federal government does not have a comprehensive list of labs working with PPPs; there is no consistent requirement for reporting laboratory accidents; incidences are generally underreported; and a 2002 U.S. bioterrorism law is frequently used to limit public disclosure of findings. Additionally, the exposé found isolated, but persistent reports of researchers that intentionally ignored biosafety rules and regulations.

An alternative view of government research oversight (Smith et al., 2015) suggests that the system is functional. A review of requests to the CDC's Division of Select Agents and Toxins (DSAT)[101] between 2006 and 2013 found that a total of 618 requests were made and that

[101] The DSAT was primarily created to approve or deny requests to conduct experiments that would create drug resistance in select agents that would make disease control difficult.

85% were determined not to be restricted experiments. The authors, who work at DSAT, interpreted these data as evidence that researchers were conscientious and practicing an abundance of caution. This was further evidenced by only four known cases during that time where restricted experiments were conducted without prior approval. It is difficult to determine where the DSAT program is an effective gatekeeper from these data alone. One can argue that the program is restrictive because only 34% of requests within the narrow purview of DSAT were approved. Conversely, in any given year, DSAT restricted less than a quarter of the requests received and, over the entire study period, about ten percent of requests were restricted. Without reviewing the actual risk assessments conducted (which are not public), it is hard to know if the program is effective, overly restrictive, or overly permissive.[102] The effectiveness of DSAT has become more relevant to the GOF debate since H5N1 avian influenza was proposed to be added to the select agents list (HHS, 2015). The important benefit of this move is that the select agents list applies to all research in the U.S. – not just labs with NIH funding. The major disadvantage is that it would also remove the risk assessment process from public scrutiny.

A review of the various assessments suggests that the most useful contribution of a single independent quantitative risk assessment may be to standardize the language of the debate. It is difficult enough to assess the quality of the data and validity of assumptions in each risk assessment. Comparison is made nearly impossible by inconsistent units (e.g., escape probability, risk per lab-year, and risk per worker-year) and different treatments of uncertainty (e.g., point estimates versus 95% confidence intervals). Using a single RBA as a starting point,

[102] Based on the arguments in this dissertation, the likelihood that these risk assessments would be controversial themselves is quite high.

hopefully the various stakeholders will at least be able to argue using the same mathematical and terminological framework.

Quantitative pandemic risk assessments are difficult because they are low probability-high consequence events. As Fineberg (2009) said, "The first lesson is to avoid over-confidence about scientific insights. Major flu pandemics arise on average only about three times every century, which means scientists can make relatively few direct observations in each lifetime and have a long time to think about each observation. That is a circumstance that is ripe for over-interpretation." While formal tools for assessing the pandemic risk of influenza viruses exist (e.g., Trock et al., 2012; Cox et al., 2014), their effectiveness is yet unproven and hampered by substantial data requirements.

6.2.2 Quantitative ≠ objective

Despite the NIH request for a comprehensive risk-benefit analysis, there is acknowledgement that this may not be possible. During the NRC symposium, both Baruch Fischhoff and Ronald Atlas discussed the difficulty of estimating benefits from the GOF research, or more generally, any basic research due to its unpredictable and serendipitous nature. Likewise, the time frame for reaping any rewards is also unknown.[103]

Further complicating the benefits analysis is the multiple ways in which evidence can be interpreted. For example, during the NRC symposium, it was widely acknowledged that genetic analysis of PPPs currently could not predict the resulting phenotype (cf. Russell et al., 2014). Critics of GOF research argued that this lack of prediction severely limited the benefits of this line research for any practical therapeutic purposes (e.g., vaccine design). However, proponents argued that this was the very reason that GOF research should continue. That is, the gap in

[103] See Chapter 2 for a detailed discussion of assessing the benefits of research.

understanding between genotype and expression of phenotype must be closed and GOF research is the best way to obtain this valuable information. Thus, an argument against the current practical value of the research is interpreted by others as an argument for the value of the research from the perspective of basic research. In this case, interpretation of benefit is a value judgment regarding what is most important: current practical utility or basic research.

The GOF controversy includes many other value-laden debates regarding risks, benefits, and assessment methodology. Proponents argue that GOF research has unique scientific value that should be considered a substantial benefit (Casadevall et al., 2014a). Meanwhile, critics argue that the scientific value may be no greater than safer alternatives which should be considered as an opportunity cost (Lipsitch, 2014). The debate also extends to disagreement regarding: the practical value of GOF experiments to policy makers (Casadevall et al., 2014b; Lipsitch, 2014); how we should count and compare the various ways of valuing research (e.g., intrinsic value versus instrumental value) (Casadevall et al., 2014c; Evans, 2014); and how publication criteria should compare public health risk to scientific merit (Dermody et al., 2014; Wain-Hobson, 2014).

Value judgments pervade the discussions. For example, Fouchier (2015a) argues that most of the biosafety and biosecurity concerns raised also apply to natural pathogen research which is currently not restricted; thus, GOF research should not be restricted. Temporarily ignoring the possibility that the consequences of an engineered virus release may be much greater, the underlying assumption is that natural pathogen research is relatively safe. However, with the recent spate of laboratory accidents, there is renewed concern that even research on naturally occurring pathogens may require further restriction.

Considerable disagreement even exists regarding ancillary effects, such as the impact of the various moratoria, "pauses," and potential regulation on the decisions of young scientists to work in virology (Casadevall and Imperiale, 2014a; Culp, 2014; Imperiale and Casadevall, 2014). The overall purpose of summarizing and critiquing some of the arguments within the debate is to emphasize the many epistemic and ethical value judgments inherent to risk-benefit analysis and to provide evidence for prior claims that a consensus-building assessment is unlikely (Casadevall and Imperiale, 2014b).

6.2.3 Ebola: lessons from a current epidemic

It may be useful to look at a recent epidemic, the 2014 Ebola outbreak in West Africa, to help us understand the current state of pandemic risk assessment. Adalja (2014) provides some background information regarding the Ebola virus, which has caused a dozen outbreaks in Africa since it was first recognized in 1976. Ebola virus disease (EVD) is a hemorrhagic fever with mortality rates of the five known strains ranging from the Reston strain, which does not cause disease in humans, to the Zaire strain which kills 90% of those infected. Much like avian influenza, EVD is a zoonotic disease, but in this case the primary natural reservoir is bats rather than birds. In addition to the much higher mortality rate, Ebola is much less transmissible than influenza. These combined factors tend to limit outbreaks and the pandemic potential of EVD is much lower than influenza.

Given its more modest growth characteristics, presumably an EVD outbreak should be easier to model than an influenza outbreak. However, multiple epidemiological models have overestimated the severity of the most recent EVD outbreak in West Africa (Butler, 2014b). A summary of predictions made in the latter half of 2014 is shown in Table 2.

Table 2: Accuracy of EVD predictions made in late 2014

Source	Model	Timeframe	Prediction	Actual (WHO data)
Gomes et al. (2014)	EVD cases in Liberia	10 weeks	9,400 - 47,000	< 7,000
Pandey et al. (2014)	daily cases in Liberia	15 weeks	249 - 545	< 50
Lewnard et al. (2014)	reported cases in Liberia	12 weeks	20,471 - 94,143	7,819
WHO Ebola Response Team (2014)	cumulative number of confirmed and probable cases	6 weeks	5,740 in Guinea 9,890 in Liberia 5,000 in Sierra Leone	1,820 in Guinea 4,240 in Liberia 4,602 in Sierra Leone

The general inaccuracy of these short-range predictions (i.e., all less than 4 months) suggests that the progression of an epidemic in developing countries is still difficult to model. Although one could argue that localized EVD outbreaks are dissimilar to influenza pandemics, there are still valuable lessons. Like EVD, influenza pandemics move quickly and the early stages are critical to assumptions regarding public health interventions. Likewise, failure to understand all the transmission routes can make the end stages of a pandemic equally difficult to predict (Enserink, 2015). There is no reason to believe that the progression of a highly pathogenic influenza epidemic would be any easier to model. Furthermore, the translation of Ebola transmission risk studies performed in developing nations to developed nations presents a considerable challenge (Haas, 2015). The reverse is also true. Thus, any influenza pandemic risk assessment would need to account for local and regional conditions that affect an outbreak's progression; averaged data tell a very incomplete story.

Any quantitative assessments using these types of estimates in either the costs or benefits calculations are contingent upon debatable assumptions. That is, the stated confidence intervals should not inspire confidence. Estimates of costs are just as uncertain. For example, the WHO

estimated that the cost of controlling the West African EVD outbreak was $4.8 million on April 10, 2014; $100 million on July 31, 2014; $490 million on August 28, 2014; and $988 million on September 16, 2014 (Butler and Morello, 2014). That the cost estimates increased by a factor of 200 within six months demonstrates how quickly costs can change and how difficult it can be to estimate the cost of a potential outbreak when it is not known at what stage the outbreak will be discovered or contained. The economic impact of a potential pandemic would also be expected to change rapidly and would need to include a range of extraneous factors that would only increase the uncertainty. Despite these daunting uncertainties, quantitative influenza pandemic risks models have been created, for use in the insurance industry (e.g., AIR Worldwide, 2013), that attempt to characterize uncertainty using complex proprietary models and thousands of stochastic simulations. However, a realistic accounting of uncertainty in a risk assessment should show that our incertitude is so large as to hinder decisive quantitative decision making.[104]

6.3 Underlying Risk Frameworks

Vogel (2013) attributes discrepancies in biotechnology risk estimates to two competing narratives on technology development: biotech *revolution* and biotech *evolution*. The revolution narrative is a dystopian form of technological determinism (cf. Smith and Marx, 1994) that predominates the biosecurity community.[105] The evolution narrative is based on the sociotechnical model of technology development (cf. Jasanoff, 2004) and comes from historical science and technology studies where the predominant view is that biotechnology is built on

[104] A secondary relevant policy issue arising from the Ebola outbreak is the noticeable difference in how scientists and public officials estimated risk. Because the transmissibility of Ebola is relatively low, medical experts recommended against the quarantine of U.S. physicians returning from West Africa. However, the cost and effort of tracking and monitoring potential exposures led public officials to lean towards precautionary quarantines.

[105] This view of an unstoppable and rapidly changing biotechnology is common in science fiction literature —some of which has directly influenced U.S. biosecurity policy (cf. Wright, 2006).

slow and incremental innovation (e.g., Nightingale and Martin, 2004). In this case, the revolution narrative roughly equates to the skeptical technological risk attitude as discussed in the previous chapter, whereas the biotech evolution narrative is a more benign and optimistic technological risk attitude. The biotech evolution narrative also forms the basis for incrementalism, an argument frequently used by proponents of research freedom (e.g., Morens et al., 2012). The argument is that each published paper is usually a small addition to the corpus of a particular field. Thus, if previous papers were not restricted, then why limit the next one. This is why regulating entire lines of research rather than individual projects is a more effective approach (Rappert, 2014).

At the December, 2014 NRC symposium, Gregory Koblentz discussed the propensity of risk attitudes to dominate risk assessments in the absence of good data (e.g., risks with few past events). In the context of biosecurity risks, Koblentz described three views: optimists, pessimists, and pragmatists. The optimists (e.g., Trevan, 2012) believe bioterrorism risk is exaggerated because: very few past terrorist attacks used bioweapons, terrorists tend to use more readily available weapons, and the technical obstacles are significant. Pessimists believe bioterrorism risk is understated because: the few past bioweapons examples show terrorist are innovative, terrorist acts have become increasingly lethal over time, terrorist ideologies embrace mass casualties, and technical obstacles are decreasing with time. Koblentz argued that these assumptions need to be clearly stated in any risk assessment.

The parallels of these opposing views to the concepts of technological optimism and skepticism, presented in Chapter 4, are clear. While I would agree with Koblentz that acknowledging technological risk attitudes in an assessment is ideal, it may be difficult in practice because, as discussed in Chapter 3, the influence of value judgments is felt throughout

the risk assessment process. That aside, the important point here is that these opposing views interpret the same available data to reach conclusions that can result in very different and controversial risk management strategies.

6.4 The Deliberative Process

Despite a desire for a "broad deliberative process," the NSABB and NRC evaluations are largely led by experts in virology and public health. It is unclear how strong a voice security experts, social scientists, and the general public have in the process. Furthermore, some of the most vocal experts have a vested personal interest in the outcome of the assessment. This has led to some confusing science communication. For example, when first presenting their findings, Fouchier and others made bold claims that influenza transmissibility in ferrets and humans was nearly equivalent. However, after the public backlash and subsequent threats to the viability of the research, Fouchier made much weaker claims regarding the use of ferrets as surrogates for humans. Lipsitch and Inglesby (2015) noted this confusion in the Fouchier (2015b) risk assessment. Similar confusing statements were made regarding the transmissibility and lethality of the engineered virus (Kahn, 2012). Sandman (2012) made the same observation but framed it in terms of risk miscommunication. That is, the initial Fouchier announcement played up the danger of the H5N1 research in order to attract attention from his peers and potential future funding. However, he made a miscalculation in that the public heard the same message, panicked, and thereby threatened all future funding for this line of research. Subsequent media interactions by Fouchier attempted to downplay the danger of the work. This is a good example of the balancing act that scientists performing controversial research must perform. The scientist

must appear to be cutting-edge while at the same time appearing safely incremental.[106] This tightrope walk only becomes more difficult as modern communications technology makes it harder to separate messages intended for the internal science community from the external public community.

The concerns of bias and narrow interests has led to calls for more active public engagement in the deliberative process; a reasonable request considering that most of the world's population would potentially be affected by both the risks and benefits of GOF research (Fineberg, 2015; Schoch-Spana, 2015).[107] That is, the public should have more access to the process than simply through public comment periods. The NSABB has acknowledged the need for public input (NSABB, 2015), but the deliberations are still unintentionally obscure and passive. A letter to the NSABB Chair (Roberts and Relman, 2015)[108] voices concerns regarding the deliberative process: that it does not include enough risk assessment experts, is not sufficiently international, does not have enough public comment opportunities, is generally opaque and moving too fast, and has an inherent conflict of interest by being funded by the NIH. Likewise, the GOF debate has been critically compared to deliberative processes associated with genome editing research (Lipsitch et al., 2015). In the latter case, researchers have proactively

[106] Fouchier has continued his risk miscommunications; his risk assessment of the frequency of a lab-induced pandemic was once every 33 billion years (Fouchier, 2015b). The number invites incredulity.

[107] There are further distinctions in that, while risks are broadly equivalent across the world, immediate benefits accrue primarily to individuals in developed nations with access to modern health care (Quinn et al., 2011). During the NRC symposium, Laurie Garrett of the Council on Foreign Relations noted this disparity, "We are having a very American conversation that excludes the rest of the planet."

[108] The letter was signed by Richard Roberts, who won the 1993 Nobel Prize in Physiology or Medicine for work on gene splicing, and David Relman, a professor of microbiology and biosecurity specialist at Stanford who was also a former NSABB board member.

called for discussions of ethics and safety in order to gain and maintain public support for controversial but potentially widely beneficial biotechnology.

Science policy making is an ethical process and therefore should not be left to the scientific community. "Scientists may approach their research with the best of intentions and the good of society in mind, but it would be a mistake to assume that they could know what is best for the world – and a tragedy to foist that burden on them" (Evans, 2013). As discussed in the previous chapter, engaging multiple perspectives increases the likelihood that important considerations in the risk assessment and viable options for risk management will not be missed. However, given the potential global reach of GOF research, truly inclusive public engagement at this scale is rarely attempted. The few previous attempts at international-scale public participation have failed to demonstrate any impact that justified the considerable effort, expense, and time involved (Rask, 2013). Some middle ground is necessary.

Another possibility is that the deliberative process is harmed by the nature of the institutions in which it takes place. That is, the NSABB may be the wrong model for appropriate technical oversight of controversial life science research. In a proposal to improve the objectivity of science and technology assessments for policy decisions, Kantrowitz (1967) suggested an "institution for scientific judgment" where science experts would argue their position and a disinterested scientist from another field would act as a judge or mediator. This would separate the judge and advocate positions as is common in the U.S. judicial system. This "science court" proposal is an attempt to be proactive considering research controversies will eventual find their way to the legal system if not adequately addressed elsewhere.[109]

[109] A deeper exploration of the concept of the science court can be found in the proceedings of a symposium held in Concord, NH, Oct. 6-7, 1994 (Field Jr., 1994).

6.5 Inherent Safety

Fischhoff suggested that, rather than use the RBA to only inform the eventual policy decision, it should instead be used to improve research design (Sharples et al., 2015). This implicitly includes inherent safety considerations. While this idea seems commonsense, it is a departure from most previous work on biosafety and biosecurity (e.g., Gaudioso et al., 2009; Caskey et al., 2010) which solely focuses on improving risk management through formalized processes and training.[110] The continued emphasis on these methods is unfortunate given the poor record of implementing such measures. For example, the European Committee for Standardization's CWA 15793 framework for lab biosafety was adopted in 2008, but by 2013 only a third of the European Biosafety Association's 118 members were using the framework and 15% were unaware of its existence (Butler, 2014a). Likewise, there has been insufficient effort to integrate biosecurity considerations into the culture of the life sciences. An informal survey of 220 graduate students and postdoctoral fellows at top NIH-funded U.S. institutions found 80% of the respondents had bioethics or biosafety training while only 10% had biosecurity training (Kahn, 2012).[111]

Lipsitch and Galvani (2014) make the same argument for improving research design when studying the potential human adaptation of avian influenza viruses, but in the context of

[110] For example, one biosafety practices report states, "biosafety best practices and international guidance span a wide variety of biosafety risk mitigation measures, which can be categorized as engineering controls, procedural and administrative controls, and the use of personal protective equipment" (Caskey et al., 2010).

[111] Some theories for the dearth of biosecurity training include: a lack of perceived relevance (Rappert, 2014), a lack of familiarity, and a general reluctance to acknowledge that the research community might contain bad actors or that research could be used to cause harm (Kahn, 2012).

responsible research principles. They argue that most gain-of-function PPP experiments are not ethically justifiable because they do not meet the criterion of yielding humanitarian benefits not attainable by safer alternatives.[112] Specifically, they contend that two main goals of GOF experiments, guiding vaccine[113] development and interpretation of surveillance data, are better achieved by other methods. Lipsitch and Galvani present an alternative list of safer approaches which includes molecular dynamical modeling, *in vitro* experiments that involve single proteins or inactive viral components, and genetic sequence comparison studies (cf. Russell et al., 2014). They also note that the ultimate goal of reducing influenza pandemics is more safely, and potentially more successfully, pursued by research related to universal influenza vaccines, broad-spectrum antiviral drugs (cf. Bekerman and Einav, 2015), and rapid vaccine manufacturing. Although Lipsitch and Galvani do not use the term, these alternatives all fall within the general concept of inherently safe design.[114]

The calls for inherently safe design appear to have yielded some consensus from the opposing camps. One sign during the NRC symposium was Yoshihiro Kawaoka, a principal investigator of one of the two original studies that started the debate, endorsing the idea that some – but not all[115] – research could be conducted with alternative techniques such as: loss-of-function studies, using less pathogenic viruses, and phenotypic (i.e., observable traits) analyses

[112] This is a principle taken from the 1947 Nuremberg Code – a foundational bioethics document that arose from revelations of Nazi medical research.
[113] It should be noted that vaccines are not inherently safe so including them within the goals of inherent safety is problematic. For example, one vaccine used in Europe during the 2009 H1N1 influenza pandemic has been associated with over 1300 cases of autoimmune-induced narcolepsy (Wekerle, 2015).
[114] While they do not mention costs, it is also useful to note that these inherently safe research alternatives are probably more cost effective in that all the additional physical, procedural, and administrative layers of biosafety and biosecurity risk management are reduced.
[115] His own research was, of course, unalterable.

(Sharples et al., 2015). Proponents of inherently safe PPP research have been buoyed by recent successes. For example, Langlois et al. (2013) have shown that species-specific microRNA targeting can be used to conduct relevant animal model PPP research that still poses low risk to humans. As Michael Imperiale stated, "you can develop safer approaches to do these types of experiments; it just needs a little bit of imagination on the part of researchers" (Sharples et al., 2015).

The suggestions to incorporate inherent safety have had an impact. While the initial HHS guidelines (Patterson et al., 2013) did include the concept of inherent safety in its third criterion,[116] the discussion among scientists and regulators was focused on a risk-benefit analysis. However, after critics started discussing legitimate alternatives, the most recent literature from the NSABB references "alternative methods that may be employed to yield similar scientific insights or benefits, while reducing potential risks" (NSABB, 2015).

Because the application of inherent safety is relatively new to the biosciences, there are some special considerations. For example, the principles of inherent safety are more restrictive than when applied to the fields of chemical and nuclear engineering. The important difference is the ability of hazardous materials in the life sciences to replicate. Thus, inherent safety requires that a hazardous organism be eliminated, rendered unable to reproduce, or have survival limitations imposed. This last approach is being pursued in synthetic biology; genetically modified organisms have been created that require a synthetic amino acid for metabolic processes (Mandell et al., 2015). The organism can only survive in the presence of the anthropogenic metabolite and has shown no ability to mutate around this restriction.

[116] "There are no feasible alternative methods to address the same scientific question in a manner that poses less risk than does the proposed approach" (Patterson et al., 2013).

As with all risk management, inherently safe design cannot be implemented unless a risk is perceived. For example, it was recently discovered that certain types of attenuated viruses have the potential to recombine into more virulent forms when the same livestock population is vaccinated with two different forms of the same vaccine (Lee et al., 2012). Traditional biosafety risk management would suggest that new policies and processes are needed to prevent double vaccinations. An inherent safety approach would focus on reformulating the vaccine (e.g., using an inactivated rather than attenuated virus) so as to remove the potential hazard altogether.

6.6 Potential Outcome

It would appear that the various accounts of laboratory biosafety lapses at the CDC that made the news in the summer of 2014 were the key events instigating the moratorium on gain-of-function research on select PPPs. These breakdowns in safety at a well-respected institution raised alarm among the public and policymakers. The standard assurances by scientists that research was being conducted safely began to lack credibility. The result was a precautionary backlash that embodies what Jasanoff (2003) has labeled "popular technological assessment." Fischhoff argues that such public responses often appear to ignore scientific evidence. However, they are not simply due to scientific illiteracy (cf. Reaven, 1987a) or incompetence,[117] but rather a reaction to sophisticated analyses that lack credibility; invoking the precautionary principle is a common response (Sharples et al., 2015). The precautionary principle has already been

[117] In Science, Technology and Society literature, this is referred to as the "deficit model" where any lack of public support is assumed to be based on ignorance and the solution is more explanatory lectures (Mooney, 2010). Meanwhile, surveys suggest that educational attainment is not strongly correlated with support for controversial science policy questions (Funk et al., 2015). While more widely accepted today, the first discussions of the underlying hubris of scientists (Feyerabend, 1978) were treated as blasphemous.

suggested as a reasonable response to most PPP research by biosecurity experts (Klotz and Sylvester, 2012).

At the NRC symposium, Alta Charo framed the GOF debate in terms of burden of proof. Whereas the government traditionally has the burden of proof to show that research may be harmful, after the declared moratorium on certain GOF research, the burden of proof now falls to the scientists to show that their work does not fall within the moratorium or that the benefits of their work exceed the risks. Harvey Fineberg, the Chair of the committee that organized the symposium, agreed with this idea noting that the NSABB had the burden of proof to halt the research when the controversy first arose in 2011, but that the burden of proof had shifted by 2014.

The principle of controversial research, as proposed in the previous chapter, suggests that the research will continue. Because the research was not banned outright, a cynical interpretation of the situation is that the pause was politically necessary for public attention to turn elsewhere and that the risk-benefit analysis will be used as a fig leaf for the continuation of the work. However, just as plausible, is that the research will be abandoned because inherently safe methods will have been found to be just as effective and less problematic for the scientists involved. Indeed, all the collected "public"[118] comments from the November 2014 NSABB meeting (NSABB, 2014a) opposed the research with general agreement that the potential benefits were minimal and the risks were substantial. It remains to be seen whether the underlying culture of technological inevitability that is prevalent in the U.S. will overcome the moratorium. Nonetheless, it is a substantial shift from only a decade ago when both the NRC and

[118] The comments were public only in the sense that they came from outside of the NSABB. The commenters consisted of six public health specialists, a virologist, and a policy analyst.

NSABB broadly supported self-regulation in the biosciences (Resnik, 2010). The situation is further complicated by a 2015 U.S. avian influenza outbreak that was the largest in decades (Hvistendahl, 2015). While the outbreak only affected poultry, the geographic immediacy and substantial economic losses may refocus public attention on this field of research.

In either case, it is important to note that this moratorium only applies to research funded by the U.S. government. An optimistic assessment is that any recommendations or regulations will carry as much weight as the Recombinant DNA Advisory Committee which issues guidelines that are mandatory only for NIH-funded research, but have become widely accepted (Resnik, 2010). While the influence of federal funding within the global research community is still substantial, it is shrinking and it is not inevitable that even closely aligned communities, such as the European Union, will adopt U.S. policies.[119] The international biotechnology regulatory landscape is inconsistent (Araki and Ishii, 2014). Some countries legislatively ban controversial biotechnology activities, some provide non-binding guidelines, while others have little or no restrictions at all.

Other contemporary trends make biotechnology regulation even more difficult. Given the increasingly competitive nature of NIH funding and the decreasing cost of advanced labs, the proportion of biomedical scientists in the U.S. working independently of federal funding is growing. While gain-of-function research on PPPs is conducted solely in well-funded labs with strict biosafety measures, the issues that underlie this debate are similar to controversies associated with synthetic biology – a field that has a much more do-it-yourself independent lab culture that includes entrepreneurs trained in non-biological fields.

[119] This was seen in the WHO's and NSABB's conflicting assessments for the original H5N1 papers (see Chapter 1).

6.7 Lessons on Managing Controversial Research

Lessons learned in a particular case study can be quite useful when presented to a larger audience. The research community is not monolithic; insight from past events can be quickly forgotten and best practices in one field are often ignored by other fields that face similar problems but have little interaction.

General rules for similar relevant activities have been proposed before. Sutherland et al. (2013) propose twenty "tips" for interpreting scientific claims. While the list was intended to help non-scientists have meaningful debates with scientists (cf. Reaven, 1987a),[120] the advice is also useful to policymakers assessing controversial research. The advice ranges from statistical (e.g., "correlation does not imply causation" or "regression to the mean can mislead") to sociological (e.g., "scientists are human"[121]). Another useful list containing ten principles for providing science advice to government (Gluckman, 2014) is also based on the realities of science policy, "These issues are urgent and of high public and political concern; the people involved hold strong positions based on their values, and the science is complex, incomplete and uncertain. Diverse meanings and understandings of risks and trade-offs dominate."

The primary point of this study is to argue that risk-benefit analysis is an ineffective tool for assessing and managing controversial research. A list of heuristic recommendations were made in the previous chapter that can improve the use of RBA by policymakers. The following are a few additional lessons to be gleaned from the H5N1 debate.

[120] The list also serves as a helpful reminder to practicing scientists.
[121] This point acknowledge "that scientific judgment itself is value-laden, and that bias and context are integral to how data are collected and interpreted" (Sutherland et al., 2013).

6.7.1 Post hoc management is difficult

Once funding has been obtained and work has begun, researchers have more than a purely intellectual position regarding a line of research. Funding creates an immediate personal financial interest as well as a long-term career impact.[122] More importantly, an investment of time and effort creates an emotional attachment; researchers will tend to defend their work for reasons that may have more to do with personal reputation and feelings of ownership than with the merits of the work. Expecting researchers to self-govern their work is asking too much of even the most well-intentioned scientists.

Furthermore, post hoc management options are limited. For controversial research, censorship is commonly proposed. However, the ubiquity of modern communications technology makes effective censoring increasingly difficult. Moreover, past censorship, mostly in the form of classification for national security purposes, has an unpleasant history of shielding scientists from essential public oversight (cf. Evans, 2013).[123] Also, in the U.S., restricting the dissemination of research can run afoul of the First Amendment under the unconstitutional conditions doctrine (Kraemer and Gostin, 2012) which prohibits a governmental benefit (e.g., research funding) from being conditioned on giving up a protected Constitutional right (e.g., free speech). These issues do not arise when considering proactive regulation at the funding stage.

[122] As Upton Sinclair said, "It is difficult to get a man to understand something when his job depends on not understanding it."
[123] Classification may be necessary in some cases to avoid Freedom of Information Act disclosures (Kraemer and Gostin, 2012).

6.7.2 Research controversies are predictably unpredictable

Research controversies are like snowflakes – unique in detail, but identical from a distant perspective. The principle of controversial research presented in Chapter 4 outlines a general trend in research controversies. With a few exceptions, a line of research continues unless it has been found to be unproductive or inferior to an alternative. As the H5N1 debate shows, research that is controversial due to safety concerns may give way to more inherently safe designs. In cases where the controversy is purely ethical, the research may never be abandoned unless pragmatic justifications can be argued.

So, while there are general trends in the evolution of research controversies, the outcomes are contingent upon the specific technical details, utility, and viable alternatives. Thus, the principle of controversial research has limited predictive power. However, predictive power is not the only valuable characteristic of a theory.[124] Rather, the principle is a simple framework for interpreting a situation that frequently arises in science policy and technology assessment.

6.8 Relevance

The avian flu gain-of-function research was used because it is a prime example of low-embodiment, low-governability science and technology. However, other examples abound (Brumfiel, 2012; Baum and Wilson, 2013). Cybersecurity research has even lower embodiment and represents a potential threat that is almost exclusively information-driven (a basic computer and internet access are the only physical requirements). Given that critical infrastructure (e.g., energy, telecommunications, and transportation systems) are at risk of disruption and potential

[124] For example, game theory, originally lauded for describing economic systems, has lost some of its luster in recent years due to its lack of predictive power (Rubinstein, 2006). Nonetheless, it still remains a valuable framework for analyzing certain problems and has generated useful niche applications, such as evolutionary game theory.

economic and national security losses increase as more systems become automated and interconnected, it is no surprise that the field of cybersecurity is rapidly growing in the twenty-first century (Nelson, 2014).[125]

Even more closely related to the influenza gain-of-function research is the emerging field of synthetic biology where detailed genetic-level bioengineering is already being used to make materials ranging from biofuels to food additives (Hayden, 2014). Applications have been primarily restricted to bacteria and fungi, but widespread applications to plants and animals are expected (Medford and Prasad, 2014). The field raises issues regarding biotechnologies such as "gene drives" which can potentially alter species (Gurwitz, 2014; Oye et al., 2014; Bohannon, 2015; Lunshof, 2015), the creation of entirely synthetic genomes (Gibson and Venter, 2014; Pennisi, 2014), or even the creation of new forms of life (Malyshev et al., 2014; Thyer and Ellefson, 2014; Woyke and Rubin, 2014). Like gain-of-function research with PPPs, expert opinions on synthetic biology vary widely (cf. Specter, 2009). Proponents claim that increasingly detailed modification of genetic sequences will result in decreasing unintended consequences (ter Meulen, 2014), while critics fear unpredictable emergent properties and scientific hubris (Cho and Relman, 2010). With the potential for substantial harm and benefits to humans and ecosystems, similar questions have been raised regarding whether the current regulatory systems and avenues of risk assessment are sufficient (Lowrie and Tait, 2010; Kelle, 2013; Carter et al., 2014; Drinkwater et al., 2014; Breitling et al., 2015; Kuiken, 2015). Recent work on bioengineering yeast to convert sugar into morphine (DeLoache et al., 2015; Fossati et al., 2015)

[125] Cyberwarfare is also full of unique ethical situations. For example, the malicious code, Stuxnet, which is believed to have destroyed approximately 1000 Iranian uranium enrichment centrifuges, has been argued to be the first highly targeted ethical war weapon (Singer and Friedman, 2014).

underlines just how socially disruptive this technology can be and the regulatory challenges that lie ahead (Ehrenberg, 2015; Oye et al., 2015)

This in no way minimizes the continued importance of "traditional" material-centric dual-use policy. The old threats have not disappeared, but they must now vie for our attention among a growing crowd of technological risks. The discussion of the role of RBA also applies to technologies with high material embodiment. For example, a relatively new uranium enrichment technique using lasers lauded by the commercial nuclear power industry has created considerable alarm in the nuclear non-proliferation community (Slakey and Cohen, 2010; Weinberger, 2012). When the US Nuclear Regulatory Commission licensed this technology in 2012, opponents criticized the commission's narrow conception of risks – considering only the physical security of an enrichment facility while ignoring the broader societal implications of encouraging the technology (Price, 2012). Likewise, entire new fields of controversial material-centric research are emerging. For example, the current debate over using geoengineering to mitigate climate-change has raised concerns regarding unintended consequences and moral hazard [126] (i.e., using geoengineering as an excuse to delay implementing more permanent and difficult mitigation strategies). Ultimately, no field of study is immune from controversy. One of the more controversial papers in the past decade was a mathematical model created by two economists that described the risk of intentional food supply poisoning (Wein and Liu, 2005); critics felt the paper served as a warning, but also as a roadmap for miscreants.

[126] Another example of a scientific moral hazard is the proposal to bring back extinct species through modern biotechnology. Unintended consequences aside, it is dangerous to encourage the belief that poor stewardship of global biodiversity has a simple technological fix (Minteer, 2014).

6.9 Future Work

This work is situated within a larger group of questions. When controversial research or technology becomes publicly known, questions are invariably asked regarding who funded the work, why it was allowed, etc. In light of the examples of controversial R&D programs previously discussed, an important general question arises: In a free and open society, how, if at all, should society control science research and technology development? Adequately addressing this question leads to a series of secondary issues that I have partially addressed: who should make the decisions (as many stakeholders as practicable – while acknowledging the limitations of public participation); what form control should take (primary controls are better than secondary); and, at what stage should controls be established (as early as possible). However, there are many questions that still need to be explored.

For example, it is unclear whether restricted communication is a useful technique for reducing technological risk. Open debate engages multiple perspectives and improves the likelihood that unintended consequences will be discovered. Secrecy and classification reduce open debate which works against risk discovery. While secrecy reduces misappropriation, it may also exacerbate the larger risks of accidents and unintended consequences. Openness can also reduce ethical issues when researchers are less inclined to be associated with R&D that is publicly controversial. Conversely, once research is published, other scientists may attempt to replicate the work or use the results to perform similar research thereby exacerbating the situation. If the work is potentially dangerous, publication may be encouraging other laboratories to also perform dangerous work. It would seem that advances in science, yielding increasingly powerful technology, challenge our ability to maintain a free society.

6.10 Conclusion

There is considerable disagreement among scientists as to whether gain-of-function research on potential pandemic pathogens has any value for current public health preparedness. However, there is some consensus that, as science, this class of research does help further our knowledge regarding how viral genetic variation affects phenotypic characteristics relevant to public health which may eventually have practical applications. The larger question is whether the risk of pandemic caused by intentional or accidental release outweighs these potential benefits. This question may be less relevant if it is determined that alternative techniques can achieve the same goals while avoiding the risks. That is, inherently safe design may be the best resolution of the H5N1 controversy, in which case the biosciences need to revamp their biosafety/biosecurity culture to adopt these principles to avoid frequently repeating this unpleasant process.

More generally, the many ways in which research benefits can be valued and the many value judgments inherent in assessing risk suggest that definitive quantitative risk-benefit analysis is not possible. This realization does not devalue RBA; it is still useful as an indispensable risk exploration and communication tool for engaging experts and the public in a conversation about risk-benefit tradeoffs. However, if calls for RBA become a knee-jerk response to what are essentially quantitatively intractable technological risk problems, everyone will be disappointed. Likewise, RBA should not be used as a stalling tactic to distract concerned citizens. RBA works best when expectations are realistic. When data are plentiful and there are no moral or cultural differences among the stakeholders relevant to the issue, RBA can generate answers or avenues to follow. However, for emerging technologies and controversial research where data are sparse and uncertainty is large, putting a number on an unavoidably subjective

quantity just yields an unavoidably subjective number that now has the suspicious appearance of objectivity. RBA works best when viewed as a tool for risk exploration and communication.

7 REFERENCES

Abaimov, S. G.; Turcotte, D. L.; Shcherbakov, R.; Rundle, J. B. Recurrence and Interoccurrence Behavior of Self-Organized Complex Phenomena. *Nonlinear Process. Geophys.* **2007**, *14* (4), 455–464.

Abbey, E. *The Monkey Wrench Gang*; Lippincott Williams & Wilkins, 1975.

Adalja, A. A. Ebola in West Africa: A Familiar Pattern? *Biosecurity and bioterrorism* **2014**, *12* (4), 161–162.

Adams, J. D. Fundamental Stocks of Knowledge and Productivity Growth. *J. Polit. Econ.* **1990**, *98* (4), 673–702.

Adams, J. D.; Sveikauskas, L. *Academic Science, Industrial R & D, and the Growth of Inputs*; Bureau of the Census, 1993; Vol. 6300.

Adams, W. M. The Value of Valuing Nature. *Science* **2014**, *346* (6209), 549–551.

Adleman, L. Resurrecting Smallpox? Easier Than You Think. *New York Times*. New York October 16, 2014, p A31.

Agan, T. Why Innovators Get Better With Age. *New York Times*. New York March 31, 2013, p BU8.

AIR Worldwide. The AIR Pandemic Flu Model https://www.air-worldwide.com/publications/.../air-pandemic-flu-model/ (accessed Sep 30, 2015).

Aizcorbe, B. A. M.; Moylan, C. E.; Robbins, C. A. Toward Better Measurement of Innovation and Intangibles. *Surv. Curr. Bus.* **2009**, *89* (1), 10–23.

Althaus, C. E. A Disciplinary Perspective on the Epistemological Status of Risk. *Risk Anal.* **2005**, *25* (3), 567–588.

Andretta, M. Some Considerations on the Definition of Risk Based on Concepts of Systems Theory and Probability. *Risk Anal.* **2014**, *34* (7), 1184–1195.

Apostolakis, G. E. The Concept of Probability in Safety Assessments of Technological Systems. *Science* **1990**, *250* (4986), 1359–1364.

Apostolakis, G. E. How Useful Is Quantitative Risk Assessment? *Risk Anal.* **2004**, *24* (3), 515–520.

Araki, M.; Ishii, T. International Regulatory Landscape and Integration of Corrective Genome Editing into in Vitro Fertilization. *Reprod. Bol. Endocrinol.* **2014**, *12* (108), 1–12.

Arita, I.; Francis, D. Is It Time to Destroy the Smallpox Virus? *Science* **2014**, *345* (6200), 1010.

Armstrong, J. S. Selecting Forecasting Methods. In *Principles of Forecasting: A Handbook for Researchers and Practitioners*; Armstrong, J. S., Ed.; Kluwer Academic Publishers: Norwell, MA, 2001; pp 365–386.

Arns, M. Open Access Is Tiring out Peer Reviewers. *Nature* **2014**, *515* (7528), 2014.

Arthur, W. B. *The Nature of Technology: What It Is and How It Evolves*; Simon & Schuster/Free Press: New York, 2009.

Assaad, F. A.; Bres, P.; Chi-Ming, C.; Dowdle, W. R.; Kendal, A. P.; Kilbourne, E. D.; Kit, L. S.; Milouchine, V.; Nejmi, S.; Noble, G.; et al. A Revision of the System of Nomenclature for Influenza Viruses: A WHO Memorandum. *Bull. World Health Organ.* **1980**, *58* (4), 585–591.

Van Asselt, M. B. A.; Rotmans, J. Uncertainty in Integrated Assessment Modelling. *Clim. Change* **2002**, *54* (1-2), 75–105.

Atkinson, P. M. Assess the Real Cost of Research Assessment. *Nature* **2014**, *516* (7530), 145.

Attenberg, J.; Ipeirotis, P.; Provost, F. Beat the Machine: Challenging Humans to Find a Predictive Model's "Unknown Unknowns." *J. Data Inf. Qual.* **2015**, *6* (1), 1.

Augustine, N. R.; Lane, N. *Restoring the Foundation: The Vital Role of Research in Preserving the American Dream*; AAAS: Cambridge, MA, 2014.

Aven, T. A Semi-Quantitative Approach to Risk Analysis, as an Alternative to QRAs. *Reliab. Eng. Syst. Saf.* **2008a**, *93* (6), 790–797.

Aven, T. *Risk Analysis: Assessing Uncertainties Beyond Expected Values and Probabilities*; John Wiley & Sons, 2008b.

Aven, T. On How to Define, Understand and Describe Risk. *Reliab. Eng. Syst. Saf.* **2010**, *95* (6), 623–631.

Aven, T. Foundational Issues in Risk Assessment and Risk Management. *Risk Anal.* **2012**, *32* (10), 1647–1656.

Aven, T.; Renn, O. On Risk Defined as an Event Where the Outcome Is Uncertain. *J. Risk Res.* **2009**, *12* (1), 1–11.

Aven, T.; Renn, O. *Risk Management and Governance: Concepts, Guidelines and Applications*; Springer Science & Business Media, 2010.

Aven, T.; Krohn, B. S. A New Perspective on How to Understand, Assess and Manage Risk and the Unforeseen. *Reliab. Eng. Syst. Saf.* **2014**, *121*, 1–10.

Aven, T.; Zio, E. Foundational Issues in Risk Assessment and Risk Management. *Risk Anal.* **2014**, *34* (7), 1164–1172.

Bal, P. M.; Veltkamp, M. How Does Fiction Reading Influence Empathy? An Experimental Investigation on the Role of Emotional Transportation. *PLoS One* **2013**, *8* (1), e55341.

Balachandra, R.; Friar, J. H. Factors for Success in R&D Projects and New Product Innovation: A Contextual Framework. *IEEE Trans. Eng. Manag.* **1997**, *44* (3), 276–287.

Balch, M. S. Mathematical Foundations for a Theory of Confidence Structures. *Int. J. Approx. Reason.* **2012**, *53* (7), 1003–1019.

Baltimore, B. D.; Berg, P.; Botchan, M.; Carroll, D.; Charo, R. A.; Church, G.; Corn, J. E.; Daley, G. Q.; Doudna, J. A.; Fenner, M.; et al. A Prudent Path Forward for Genomic Engineering and Germline Gene Modification. *Science* **2015**, science.aab1028 – .

Ban, T. A. The Role of Serendipity in Drug Discovery. *Dialogues Clin. Neurosci.* **2006**, *8* (3), 335–344.

Baum, A.; Fleming, R.; Davidson, L. M. Natural Disaster and Technological Catastrophe. *Environ. Behav.* **1983**, *15* (3), 333–354.

Baum, S. D. The Great Downside Dilemma for Risky Emerging Technologies. *Phys. Scr.* **2014**, *89* (12), 128004.

Baum, S. D.; Wilson, G. S. The Ethics of Global Catastrophic Risk from Dual-Use Bioengineering. *Ethics Biol. Eng. Med.* **2013**, *4* (1), 59–72.

Begley, S.; Steenhuysen, J. How to fix U.S. biosecurity leaks? Close some labs http://www.reuters.com/article/2014/07/14/us-usa-anthrax-labs-analysis-idUSKBN0FJ0BC20140714.

Bekerman, E.; Einav, S. Combating Emerging Viral Threats. *Science* **2015**, *348* (6232), 282–283.

Ben-Haim, Y. *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, Second Edi.; Academic Press: Oxford, 2006.

Ben-Haim, Y. Doing Our Best: Optimization and the Management of Risk. *Risk Anal.* **2012**, *32* (8), 1326–1332.

Berg, P. The Dual-Use Conundrum. *Science* **2012**, *337* (6100), 1273.

Berg, P.; Singer, M. F. The Recombinant DNA Controversy: Twenty Years Later. *Proc. Natl. Acad. Sci.* **1995**, *92* (20), 9011–9013.

Berger, E. Is NASA about Jobs, or Actually Accomplishing Something? *Houston Chronicle*. Houston May 17, 2013.

Berger, J. O. *Statistical Decision Theory and Bayesian Analysis*, Second Edi.; Springer Science & Business Media: New York, 1985.

Berger, J. O.; Moreno, E.; Pericchi, L. R.; Bayarri, M. J.; Bernardo, J. M.; Cano, J. A.; De la Horra, J.; Martín, J.; Ríos-Insúa, D.; Betrò, B.; et al. An Overview of Robust Bayesian Analysis. *Test* **1994**, *3* (1), 5–124.

Bergstrom, L. Reflections on Consequentialism. *Theoria* **1996**, *62* (1-2), 74–94.

Bernal, J. D. *The Social Function of Science*; George Routledge & Sons Ltd.: London, 1939.

Bertrand, P.; Goupil, F. Descriptive Statistics for Symbolic Data. In *Analysis of Symbolic Data*; Bock, H.-H., Diday, E., Eds.; Springer-Verlag Berlin Heidelberg: Berlin, 2000; pp 106–124.

Bessen, J. The Value of U.S. Patents by Owner and Patent Characteristics. *Res. Policy* **2008**, *37* (5), 932–945.

Bhattacharjee, Y. U . S . Panel Calls for Extra Review of Dual-Use Research. *Science* **2006**, *313* (5785), 284.

Bijak, J. Dealing with Uncertainty in International Migration Predictions: From Probabilistic Forecasting to Decision Analysis. *Int. Migr.* **2010**, No. April, 1–6.

Bimber, B. Three Faces of Technological Determinism. In *Does Technology Drive History: The Dilemma of Technological Determinism*; Smith, M. R., Marx, L., Eds.; MIT Press: Cambridge, MA, 1994; pp 80–100.

Biosafety. [Editorial] Biosafety in the Balance. *Nature* **2014**, *510* (7506), 443.

Blackwell, D. Comparison of Experiments. *Second Berkeley Symp. Math. Stat. Probab.* **1951**, *1*, 93–102.

Boddie, C.; Sell, T. K.; Watson, M. Federal Funding for Health Security in FY2015. *Biosecurity and bioterrorism* **2014**, *12* (4), 163–177.

Bogen, K. T. Does EPA Underestimate Cancer Risks by Ignoring Susceptibility Differences? *Risk Anal.* **2014a**, *34* (10), 1780–1784.

Bogen, K. T. Unveiling Variability and Uncertainty for Better Science and Decisions on Cancer Risks from Environmental Chemicals. *Risk Anal.* **2014b**, *34* (10), 1795–1806.

Bohannon, J. Biologists Devise Invasion Plan for Mutations. *Science* **2015**, *347* (6228), 1300–1300.

Bolger, F.; Rowe, G. The Aggregation of Expert Judgment: Do Good Things Come to Those Who Weight? *Risk Anal.* **2015a**, *35* (1), 5–11.

Bolger, F.; Rowe, G. There Is Data, and Then There Is Data: Only Experimental Evidence Will Determine the Utility of Differential Weighting of Expert Judgment. *Risk Anal.* **2015b**, *35* (1), 21–26.

Bondi, H. The Lure of Completeness. In *The Encyclopaedia of Ignorance: Everything you ever wanted to know about the unknown*; Duncan, R., Weston-Smith, M. M., Eds.; Pocket Books, 1978.

Bornmann, L. Measuring the Societal Impact of Research. *EMBO Rep.* **2012**, *13* (8), 673–676.

Boulton, G. *Science as an Open Enterprise*; The Royal Society: London, 2012.

Bowles, M. D. U.S. Technological Enthusiasm and British Technological Skepticism in the Age of the Analog Brain. *IEEE Ann. Hist. Comput.* **1996**, *18* (4), 5–15.

Boyd, I. A Standard for Policy-Relevant Science. *Nature* **2013**, *501* (7466), 159–160.

Bozeman, B. *Public Values and Public Interest: Counterbalancing Economic Individualism*; Georgetown University Press: Washington, DC, 2007.

Brännström, M.; Johannesson, L.; Bokström, H.; Kvarnström, N.; Mölne, J.; Dahm-Kähler, P.; Enskog, A.; Milenkovic, M.; Ekberg, J.; Diaz-Garcia, C.; et al. Livebirth after Uterus Transplantation. *Lancet* **2014**, *385* (9968), 607–616.

Brauer, F.; Feng, Z.; Castillo-Chavez, C. Discrete Epidemic Models. *Math. Biosci. Eng.* **2010**, *7* (1), 1–15.

Bredehoeft, J. From Models to Performance Assessment: The Conceptualization Problem. *Ground Water* **2003**, *41* (5), 571–577.

Breitling, R.; Takano, E.; Gardner, T. S. Judging Synthetic Biology Risks. *Science* **2015**, *347* (6218), 107.

Broad, W. J. Vindication for Entrepreneurs Watching Sky: Yes, It Can Fall. *New York Times*. New York February 17, 2013, p A1.

Brown, G. E. Past and Prologue: Why I Am Optimistic About the Future. In *AAAS Science and Technology Policy Yearbook*; AAAS: Washington, DC, 1999.

Van der Bruggen, K. Possibilities, Intentions and Threats: Dual Use in the Life Sciences Reconsidered. *Sci. Eng. Ethics* **2012**, *18* (4), 741–756.

Brumfiel, G. Good Science, Bad Science. *Nature* **2012**, *484* (7395), 432–434.

Brun, W. Cognitive Components in Risk Perception: Natural versus Manmade Risks. *J. Behav. Decis. Mak.* **1992**, *5* (2), 117–132.

Burt, A.; Koufopanou, V. Homing Endonuclease Genes: The Rise and Fall and Rise Again of a Selfish Element. *Curr. Opin. Genet. Dev.* **2004**, *14* (6), 609–615.

Bush, V. Science - the Endless Frontier. *Trans. Kansas Acad. Sci.* **1945**, *48* (3), 231–264.

Butler, D. Pathogen-Research Laws Queried. *Nature* **2013**, *503* (7474), 19.

Butler, D. Biosafety Controls Come under Fire. *Nature* **2014a**, *511* (7511), 515–516.

Butler, D. Models Overestimate Ebola Cases. *Nature* **2014b**, *515* (7525), 18.

Butler, D.; Morello, L. Ebola by the Numbers: The Size, Spread and Cost of an Outbreak. *Nature* **2014**, *514* (7522), 284–285.

Butler, D.; Maher, B. Risks of Flu Work Underrated. *Nature* **2014**, *511* (7507), 13–14.

Cairns-Smith, A. G. Synthetic Life for Industry. In *The Encyclopaedia of Ignorance: Everything you ever wanted to know about the unknown*; Duncan, R., Weston-Smith, M., Eds.; Pocket Books: New York, NY, 1978; pp 405–410.

Calabrese, E. J.; Baldwin, L. A. Toxicology Rethinks Its Central Belief. *Nature* **2003**, *421* (6924), 691–692.

Callaway, E. Reproductive Medicine: The Power of Three. *Nature* **2014**, *509* (7501), 414–417.

Calow, P. Environmental Risk Assessors as Honest Brokers or Stealth Advocates. *Risk Anal.* **2014**, *34* (11), 1972–1977.

Camerer, C.; Loewenstein, G.; Weber, M. The Curse of Knowledge in Economic Settings: An Experimental Analysis. *J. Polit. Econ.* **1989**, *97* (5), 1232–1254.

Carmen, I. H. Leadership Void for "designer Baby" Ethics. *Science* **2014**, *345* (6193), 149.

Carson, R. *Silent Spring*; Houghton Mifflin: Boston, 1962.

Carter, S. R.; Rodemeyer, M.; Garfinkel, M. S.; Friedman, R. M. *Synthetic Biology and the U.S. Biotechnology Regulatory System: Challenges and Options*; J. Craig Venter Institute: Rockville, MD, 2014.

Casadevall, A.; Imperiale, M. J. Reply to "A Brain Drain due to Increased Regulation of Influenza Virus Research Is Highly Speculative." *MBio* **2014a**, *5* (5), e01860–14.

Casadevall, A.; Imperiale, M. J. Risks and Benefits of Gain-of-Function Experiments with Pathogens of Pandemic Potential, Such as Influenza Virus: A Call for a Science-Based Discussion. *MBio* **2014b**, *5* (4), e01730–14.

Casadevall, A.; Howard, D.; Imperiale, M. J. An Epistemological Perspective on the Value of Gain-of-Function Experiments Involving Pathogens with Pandemic Potential. *MBio* **2014a**, *5* (5), e01875–14.

Casadevall, A.; Howard, D.; Imperiale, J. Reply to "Can Limited Scientific Value of Potential Pandemic Pathogen Experiments Justify the Risks ?" *MBio* **2014b**, *5* (5), e02053–14.

Casadevall, A.; Howard, D.; Imperiale, J. Reply to "Valuing Knowledge: A Reply to the Epistemological Perspective on the Value of Gain-of-Function Experiments ." *MBio* **2014c**, *5* (5), e02054–14.

Casadevall, A.; Howard, D.; Imperiale, M. J. The Apocalypse as a Rhetorical Device in the Influenza Virus Gain-of- Function Debate. *MBio* **2014d**, *5* (5), e02062–14.

Caskey, S.; Gaudioso, J.; Salerno, R.; Wagener, S.; Shigematsu, M.; Risi, G.; Kozlovac, J.; Halkjær-Knudsen, V.; Prat, E. *Biosafety Risk Assessment Methodology*; Sandia National Laboratories: Albuquerque, NM, 2010.

Casman, E. A.; Morgan, M. G.; Dowlatabadi, H. Mixed Levels of Uncertainty in Complex Policy Models. *Risk Anal.* **1999**, *19* (1), 33–42.

CDC. *Report on the Potential Exposure to Anthrax*; Centers for Disease Control and Prevention: Atlanta, GA, 2014.

Cello, J.; Paul, A. V; Wimmer, E. Chemical Synthesis of Poliovirus cDNA: Generation of Infectious Virus in the Absence of Natural Template. *Science* **2002**, *297* (5583), 1016–1018.

CEN. *Laboratory Biorisk Management*; European Committee for Standarization: Brussels, 2012.

Chaiken, S. Heuristic versus Systematic Information Processing and the Use of Source versus Message Cues in Persuasion. *J. Pers. Soc. Psychol.* **1980**, *39* (5), 752–766.

Chameau, J.-L.; Ballhaus, W. F.; Lin, H. S. *Emerging and Readily Available Technologies and National Security-A Framework for Addressing Ethical, Legal, and Societal Issues*; National Academies Press: Washington, DC, 2014.

Chassot, S.; Hampl, N.; Wüstenhagen, R. When Energy Policy Meets Free-Market Capitalists: The Moderating Influence of Worldviews on Risk Perception and Renewable Energy Investment Decisions. *Energy Res. Soc. Sci.* **2014**, *3*, 143–151.

Cho, M. K.; Relman, D. A. Genetic Technologies. Synthetic "Life," Ethics, National Security, and Public Discourse. *Science* **2010**, *329* (5987), 38–39.

Choi, S. Public Perception and Acceptability of Technological Risk : Policy Implications for Governance. *J. Converg. Inf. Technol.* **2013**, *8* (13), 605–615.

Chorost, M. *Rebuilt: How Becoming Part Computer Made Me More Human*; Houghton Mifflin: New York, 2005.

Chorost, M. *World Wide Mind: The Coming Integration of Humanity, Machines, and the Internet*; Free Press: New York, 2011.

Chubin, D. E. The Editor's Dilemma. *Science* **2013**, *342* (6155), 193.

Cipra, B. Revealing Uncertainties in Computer Models. *Science* **2000**, *287* (5455), 960–961.

Cirković, M. M. Small Theories and Large Risks--Is Risk Analysis Relevant for Epistemology? *Risk Anal.* **2012**, *32* (11), 1994–2004.

Clark, M. Thanks JFK : States Gained from Space Program. *Stateline*. Pew Charitable Trusts November 20, 2013.

Claxton, K. P.; Sculpher, M. J. Using Value of Information Analysis to Prioritise Health Research. *Pharmacoeconomics* **2006**, *24* (11), 1055–1068.

Clemen, R. T.; Winkler, R. L. Combining Probability Distributions From Experts in Risk Analysis. *Risk Anal.* **1999**, *19* (2), 187–203.

Clemen. *Making Hard Decisions: Introduction to Decision Analysis*; Brooks/Cole: Pacific Grove, CA, 1991.

Coeckelbergh, M. Risk and Public Imagination: Mediated Risk Perception as Imaginative Moral Judgement. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 202–219.

Cohen, J. E.; Oppenheim, J. Is a Limit to the Median Length of Human Life Imminent? *Genus* **2012**, *68* (1), 11–40.

Cohen, J. Updated: U.S. biosafety panel to come out of hibernation with new members | Science/AAAS | News http://news.sciencemag.org/biology/2014/07/updated-u-s-biosafety-panel-come-out-hibernation-new-members (accessed Jan 13, 2015).

Cohen, W. M.; Levinthal, D. A. Innovation and Learning: The Two Faces of R&D. *Econ. J.* **1989**, *99* (397), 569–596.

Colglazier, E. W.; Dungan, D. L.; Reaven, S. J. Value Issues and Stakeholders' Views in Radioactive Waste Management. In *Waste Management '87 Symposium*; Post, R., Wacks, M., Eds.; University of Arizona Press: Tuscon, 1987; pp 185–193.

Commoner, B. *The Closing Circle: Nature, Man, and Technology*; Knopf: New York, 1971.

Cooke, R. M. The Aggregation of Expert Judgment: Do Good Things Come to Those Who Weight? *Risk Anal.* **2015**, *35* (1), 12–15.

Costanza, R. Four Visions of the Century Ahead: Will It Be Star Trek , Ecotopia , Big Government, or Mad Max. *Futurist* **1999**, *33* (2), 23–28.

Costanza, R. Visions of Alternative (unpredictable) Futures and Their Use in Policy Analysis. *Conserv. Ecol.* **2000**, *4* (1), 5.

Costanza, R.; Kubiszewski, I.; Giovannini, E.; Lovins, H.; McGlade, J.; Pickett, K. E.; Ragnarsdóttir, K.; Roberts, D.; De Vogli, R.; Wilkinson, R. Time to Leave GDP behind. *Nature* **2014**, *505* (7583), 283–285.

Coull, S. E.; Kenneally, E. E. A Qualitative Risk Assessment Framework for Sharing Computer Network Data. *SSRN Electron. J.* **2012**.

Cox, L. A. What's Wrong with Risk Matrices? *Risk Anal.* **2008**, *28* (2), 497–512.

Cox, L. A. Community Resilience and Decision Theory Challenges for Catastrophic Events. *Risk Anal.* **2012a**, *32* (11), 1919–1934.

Cox, L. A. Confronting Deep Uncertainties in Risk Analysis. *Risk Anal.* **2012b**, *32* (10), 1607–1629.

Cox, L. A.; Babayev, D.; Huber, W. Some Limitations of Qualitative Risk Rating Systems. *Risk Anal.* **2005**, *25* (3), 651–662.

Cox, N. J.; Trock, S. C.; Burke, S. A. Pandemic Preparedness and the Influenza Risk Assessment Tool (IRAT). *Curr. Top. Microbiol. Immunol.* **2014**, *385*, 119–136.

Cragin, M. H.; Nichols, L.; Simon, M.; Watts, S. M. Measuring Science : Emerging Tools for Analysis of Federal R & D Investments. In *ASIST 2012: Proceedings of the 75th ASIS&T Annual Meeting*; Association for Information Science and Technology, 2012; p 4.

Cranor, C. F. Normative Nature of Risk Assessment: Features and Possibilities, The. *Risk Heal. Saf. Environ.* **1997**, *8*, 123–136.

Cranor, C. F. A Plea for a Rich Conception of Risks. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 27–39.

Culp, D. A Brain Drain due to Increased Regulation of Influenza Virus Research Is Highly Speculative. *MBio* **2014**, *5* (5), e01814–14.

Cumming, R. B. Is Risk Assessment A Science? *Risk Anal.* **1981**, *1* (1), 1–3.

Cyranoski, D. Ethics of Embryo Editing Divides Scientists. *Nature* **2015**, *519* (7543), 272–272.

Cyranoski, D.; Reardon, S. Embryo Editing Sparks Epic Debate. *Nature* **2015**, *520* (7549), 593–594.

Czarnitzki, D.; Lopes-Bento, C. Value for Money? New Microeconometric Evidence on Public R&D Grants in Flanders. *Res. Policy* **2013**, *42* (1), 76–89.

Damon, I. K.; Damaso, C. R.; McFadden, G. Are We There yet? The Smallpox Research Agenda Using Variola Virus. *PLoS Pathog.* **2014**, *10* (5), e1004108.

Davidson, M. D. Acceptable Risk to Future Generations. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 77–91.

Davis, C. T.; Chen, L.; Pappas, C.; Stevens, J.; Tumpey, T. M.; Gubareva, L. V; Katz, J. M.; Villanueva, J. M.; Donis, R. O.; Cox, N. J. Use of Highly Pathogenic Avian Influenza A (H5N1) Gain-Of-Function Studies for Molecular-Based Surveillance and Pandemic Preparedness. *MBio* **2014**, *5* (6), e02431–14.

Davis, P. M. *Journal Usage Half-Life*; American Association of Publishers: Ithaca, 2013.

DeLoache, W. C.; Russ, Z. N.; Narcross, L.; Gonzales, A. M.; Martin, V. J. J.; Dueber, J. E. An Enzyme-Coupled Biosensor Enables (S)-Reticuline Production in Yeast from Glucose. *Nat. Chem. Biol.* **2015**.

Dempster, A. P. Upper and Lower Probabilities Induced by a Mulitvalued Mapping. *Ann. Math. Stat.* **1967**, *38* (2), 325–339.

Dermody, T. S.; Casadevall, A.; Imperiale, M. J.; Sandri-Goldin, R. M.; Shenkh, T. The Decision to Publish an Avian H7N1 Influenza Virus Gain-of-Function Experiment. *MBio* **2014**, *5* (5), e01985–14.

Dickersin, K. The Existence of Publication Bias and Risk Factors for Its Occurrence. *JAMA J. Am. Med. Assoc.* **1990**, *263* (10), 1385.

DOJ. *Amerithrax Investigative Summary*; U.S. Department of Justice: Washington, DC, 2010.

Doorn, N. The Blind Spot in Risk Ethics: Managing Natural Hazards. *Risk Anal.* **2015**, *35* (3), 354–360.

Doorn, N.; Hansson, S. O. Should Probabilistic Design Replace Safety Factors? *Philos. Technol.* **2011**, *24* (2), 151–168.

Douglas, H. Inductive Risk and Values in Science. *Philos. Sci.* **2000**, *67* (4), 559–579.

Douglas, M.; Wildavsky, A. *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers*; University of California Press: Berkeley, CA, 1982.

Draper, D. Assessment and Propagation of Model Uncertainty. *J. R. Stat. Soc. Ser. B* **1995**, *57* (1), 45–97.

Dreyfus, H. *What Computers Can't Do: The Limits of Artificial Intelligence*; MIT Press: New York, 1972.

Driessen, E.; Hollon, S. D.; Bockting, C. L. H.; Cuijpers, P.; Turner, E. H. Does Publication Bias Inflate the Apparent Efficacy of Psychological Treatment for Major Depressive Disorder? A Systematic Review and Meta-Analysis of US National Institutes of Health-Funded Trials. *PLoS One* **2015**, *10* (9), e0137864.

Drinkwater, K.; Kuiken, T.; Lightfoot, S.; McNamara, J.; Oye, K. *Creating a Research Agenda for the Ecological Implications of Synthetic Biology*; Washington, DC, 2014.

Dubois, D. Possibility Theory and Statistical Reasoning. *Comput. Stat. Data Anal.* **2006**, *51* (1), 47–69.

Dubois, D.; Prade, H. *Possibility Theory: An Approach to Computerized Processing of Uncertainty*; Plenum Press: New York, 1988.

Duckett, D.; Wynne, B.; Christley, R. M.; Heathwaite, A. L.; Mort, M.; Austin, Z.; Wastling, J. M.; Latham, S. M.; Alcock, R.; Haygarth, P. Can Policy Be Risk-Based? The Cultural Theory of Risk and the Case of Livestock Disease Containment. *Sociol. Ruralis* **2015**, *55* (4), 379–399.

Eagly, A. H.; Chaiken, S. *The Psychology of Attitudes.*; Harcourt, Brace, & Jovanovich: Fort Worth, TX, 1993.

Edwards, D. W. Are We Too Risk-Averse for Inherent Safety? *Process Saf. Environ. Prot.* **2005**, *83* (2), 90–100.

Edwards, W.; von Winterfeldt, D. Public Values in Risk Debates. *Risk Anal.* **1987**, *7* (2), 141–158.

Ehrenberg, R. Engineered Yeast Paves Way for Home-Brew Heroin. *Nature* **2015**, *521* (7552), 267–268.

Elahi, S. Here Be Dragons… Exploring the "unknown Unknowns." *Futures* **2011**, *43* (2), 196–201.

Ellul, J. *The Technological Society*, English Tr.; Wilkinson, J., Ed.; Vintage Books: New York, 1964.

Engel-Glatter, S. Dual-Use Research and the H5N1 Bird Flu: Is Restricting Publication the Solution to Biosecurity Issues? *Sci. Public Policy* **2013**, *41* (3), 370–383.

Enserink, M. Dutch H5N1 Ruling Raises New Questions. *Science* **2013**, *342* (6155), 178.

Enserink, M. Infectious Diseases. In Guinea, a Long, Difficult Road to Zero Ebola Cases. *Science* **2015**, *348* (6234), 485–486.

Ernst, E.; Resch, K. L. Risk-Benefit Ratio or Risk-Benefit Nonsense? *J. Clin. Epidemiol.* **1996**, *49* (10), 1203–1204.

Ernst, H. Patent Applications and Subsequent Changes of Performance: Evidence from Time-Series Cross-Section Analyses on the Firm Level. *Res. Policy* **2001**, *30* (1), 143–157.

Espinoza, N. Incommensurability: The Failure to Compare Risks. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 128–143.

Evans, A. W.; Verlander, N. Q. What Is Wrong with Criterion FN-Lines for Judging the Tolerability of Risk? *Risk Anal.* **1997**, *17* (2), 157–168.

Evans, N. G. Great Expectations - Ethics, Avian Flu and the Value of Progress. *J. Med. Ethics* **2013**, *39* (4), 209–213.

Evans, N. G. Valuing Knowledge: A Reply to the Epistemological Perspective on the Value of Gain-of-Function Experiments. *MBio* **2014**, *5* (5), e01993–14.

Eypasch, E.; Lefering, R.; Kum, C. K.; Troidl, H. Probability of Adverse Events That Have Not yet Occurred: A Statistical Reminder. *BMJ* **1995**, *311* (7005), 619–620.

Faden, R. R.; Karron, R. A. Public Health and Biosecurity. The Obligation to Prevent the next Dual-Use Controversy. *Science* **2012**, *335* (6070), 802–804.

Fairchild, A. L.; Bayer, R. Smoke and Fire over E-Cigarettes. *Science* **2015**, *347* (6220), 375–376.

Falk, A.; Szech, N. Morals and Markets. *Science* **2013**, *340* (6133), 707–711.

Fang, J. Ecology: A World without Mosquitoes. *Nature* **2010**, *466* (7305), 432–434.

Farrell, R. M.; Falcone, T. Uterine Transplant: New Medical and Ethical Considerations. *Lancet* **2015**, *385* (9968), 581–582.

FDA. *Structured Approach to Benefit-Risk Assessment in Drug Regulatory Decision-Making*; U.S. Food and Drug Administration: Silver Springs, MD, 2013.

Federal Demonstration Partnership. STAR METRICS - Phase II http://sites.nationalacademies.org/PGA/fdp/PGA_057159.

Feinendegen, L. E. Evidence for Beneficial Low Level Radiation Effects and Radiation Hormesis. *Br. J. Radiol.* **2005**, *78* (925), 3–7.

Feleppa, R. Epistemic Utility and Theory Acceptance: Comments on Hempel. *Synthese* **1981**, *46* (3), 413–420.

Ferson, S. *Model Uncertainty in Risk Analysis*; Université de Technologie de Compiègne: Compiègne Cedex, France, 2014.

Ferson, S.; Ginzburg, L. R. Different Methods Are Needed to Propagate Ignorance and Variability. *Reliab. Eng. Syst. Saf.* **1996**, *54* (2-3), 133–144.

Ferson, S.; Hajagos, J. G. Arithmetic with Uncertain Numbers: Rigorous and (often) Best Possible Answers. *Reliab. Eng. Syst. Saf.* **2004**, *85* (1-3), 135–152.

Ferson, S.; Tucker, W. T. *Sensitivity in Risk Analysis with Uncertain Number*; Sandia National Laboratories: Albuquerque, NM, 2006.

Ferson, S.; Nelsen, R. B.; Hajagos, J.; Berleant, D. J.; Zhang, J. *Dependence in Probabilistic Modeling, Dempster-Shafer Theory, and Probability Bounds Analysis*; Sandia National Laboratories: Albuquerque, NM, 2004.

Festinger, L. *A Theory of Cognitive Dissonance*; Stanford University Press: Redwood City, CA, 1957.

Feyerabend, P. Consolations for the Specialist. In *Criticism and the Growth of Knowledge*; Lakatos, I., Musgrave, A., Eds.; Cambridge University Press: Cambridge, UK, 1970; pp 197–230.

Feyerabend, P. *Against Method*; New Left Books: London, 1975.

Feyerabend, P. *Science in a Free Society*; New Left Books: London, 1978.

Feyerabend, P. *The Tyranny of Science*; Oberheim, E., Ed.; Polity Press: Cambridge, UK, 2011.

Field Jr., T. G. The Science Court Symposia http://ipmall.info/risk/scict.htm (accessed May 6, 2015).

Fineberg, H. V. Swine Flu of 1976: Lessons from the Past. An Interview with Dr Harvey V Fineberg. *Bull. World Health Organ.* **2009**, *87* (6), 414–415.

Fineberg, H. V. Wider Attention for GOF Science. *Science* **2015**, *347* (6225), 929.

Fink, G. R. *Biotechnology Research in an Age of Terrorism*; National Academies Press: Washington, DC, 2004.

Finkel, A. M. EPA Underestimates, Oversimplifies, Miscommunicates, and Mismanages Cancer Risks by Ignoring Human Susceptibility. *Risk Anal.* **2014**, *34* (10), 1785–1794.

Fischhoff, B. Risk Perception and Communication Unplugged: Twenty Years of Process. *Risk Anal.* **1995**, *15* (2), 137–145.

Fischhoff, B. Scientific Management of Science? *Policy Sci.* **2000**, *33* (1), 73–87.

Fischhoff, B.; Slovic, P.; Lichtenstein, S.; Read, S.; Combs, B. How Safe Is Safe Enough? A Psychometric Study of Attitudes towards Technological Risks and Benefits. *Policy Sci.* **1978**, *9* (2), 127–152.

Flage, R.; Aven, T. Expressing and Communicating Uncertainty in Relation to Quantitative Risk Analysis. *Reliab. Risk Anal. Theory Appl.* **2009**, *2* (13), 9–18.

Flage, R.; Aven, T.; Zio, E.; Baraldi, P. Concerns, Challenges, and Directions of Development for the Issue of Representing Uncertainty in Risk Assessment. *Risk Anal.* **2014**, *34* (7), 1196–1207.

Flyvbjerg, B. Five Misunderstandings About Case-Study Research. *Qual. Inq.* **2006**, *12* (2), 219–245.

Fontanari, L.; Gonzalez, M.; Vallortigara, G.; Girotto, V. Probabilistic Cognition in Two Indigenous Mayan Groups. *Proc. Natl. Acad. Sci.* **2014**, *111* (48), 17075–17080.

Fossati, E.; Narcross, L.; Ekins, A.; Falgueyret, J.-P.; Martin, V. J. J. Synthesis of Morphinan Alkaloids in Saccharomyces Cerevisiae. *PLoS One* **2015**, *10* (4), e0124459.

Fouchier, R. A. M. Reply to " Comments on Fouchier's Calculation of Risk and Elapsed Time for Escape of a Laboratory-Acquired Infection from His Laboratory ." *MBio* **2015a**, *6* (2), e00407–e00415.

Fouchier, R. A. M. Studies on Influenza Virus Transmission between Ferrets: The Public Health Risks Revisited. *MBio* **2015b**, *6* (1), e02560–14.

Fouchier, R. A. M.; García-Sastre, A.; Kawaoka, Y. Pause on Avian Flu Transmission Studies. *Nature* **2012**, *481* (7382), 443.

Fouchier, R. A. M.; Kawaoka, Y.; Cardona, C.; Compans, R. W.; García-Sastre, A.; Govorkova, E. A.; Guan, Y.; Herfst, S.; Orenstein, W. A.; Peiris, J. S. M.; et al. Gain-of-Function Experiments on H7N9. *Science* **2013**, *341* (6146), 612–613.

Fountain, H. Six Vials of Smallpox Discovered in Laboratory Near Washington. *New York Times*. New York July 9, 2014, p A14.

Fox, C. R.; Tversky, A. Ambiguity Aversion and Comparative Ignorance. *Q. J. Econ.* **1995**, *110* (3), 585–603.

Fox, D. M. *Engines of Culture: Philanthropy and Art Museums*, 2nd ed.; Transaction Publishers: New Brunswick, 1995.

Franco, A.; Malhotra, N.; Simonovits, G. Publication Bias in the Social Sciences: Unlocking the File Drawer. *Science* **2014**, *345* (6203), 1502–1505.

Freeman, R.; Van Reenen, J. What If Congress Doubled R&D Spending on the Physical Sciences? In *Innovation Policy and the Economy, Volume 9*; Lerner, J., Stern, S., Eds.; University of Chicago Press, 2009; Vol. 9, pp 1–38.

Freud, S. *Civilization and Its Discontents*; Hogarth: Oxford, 1930.

Frisch, W.; Meschede, M.; Blakey, R. C. *Plate Tectonics: Continental Drift and Mountain Building*; Springer Science & Business Media, 2010; Vol. 2.

Funk, C.; Rainie, L.; Page, D. *Americans, Politics and Science Issues*; Pew Research Center, 2015.

Furmanski, M. The 1977 H1N1 Influenza Virus Reemergence Demonstrated Gain-of-Function Hazards. *MBio* **2015**, *6* (5), e01434–15.

Galizi, R.; Doyle, L. A.; Menichelli, M.; Bernardini, F.; Deredec, A.; Burt, A.; Stoddard, B. L.; Windbichler, N.; Crisanti, A. A Synthetic Sex Ratio Distortion System for the Control of the Human Malaria Mosquito. *Nat. Commun.* **2014**, *5*, 3977.

Gallucci, R. "What - Me Worry?" "Why so Serious?": A Personal View on the Fukushima Nuclear Reactor Accidents. *Risk Anal.* **2012**, *32* (9), 1444–1450.

Gardoni, P.; Murphy, C. A Scale of Risk. *Risk Anal.* **2014**, *34* (7), 1208–1227.

Gaudioso, J.; Caskey, S. A.; Burnett, L.; Heegaard, E.; Owens, J.; Stroot, P. *Strengthening Risk Governance in Bioscience Laboratories*; Sandia National Laboratories: Albuquerque, NM, 2009.

Gesang, B. Are Moral Philosophers Moral Experts? *Bioethics* **2010**, *24* (4), 153–159.

Gewin, V. Biotechnology: Independent Streak. *Nature* **2013**, *499* (7459), 509–511.

Gibson, D. G.; Venter, J. C. Construction of a Yeast Chromosome. *Nature* **2014**, *509* (7499), 168–169.

Gillette, C. P.; Krier, J. E. Risk, Courts, and Agencies. *Univ. PA. Law Rev.* **1990**, *138* (4), 1027–1109.

Glass, K.; Xia, Y.; Grenfell, B. T. Interpreting Time-Series Analyses for Continuous-Time Biological Models—measles as a Case Study. *J. Theor. Biol.* **2003**, *223* (1), 19–25.

Glenn, J. The Radium Age. *Nature* **2012**, *489* (7415), 204–205.

Gluckman, P. The Art of Science Advice to Government. *Nature* **2014**, *507* (7491), 163–165.

Goldston, D. Time to Stop Relying on Things Past. *Nature* **2012**, *492* (7429), 313.

Gomes, M. F. C.; Pastore y Piontti, A.; Rossi, L.; Chao, D.; Longini, I.; Halloran, M. E.; Vespignani, A. Assessing the International Spreading Risk Associated with the 2014 West African Ebola Outbreak. *PLoS Curr.* **2014**, *6*.

Goodwin, P.; Wright, G. The Limits of Forecasting Methods in Anticipating Rare Events. *Technol. Forecast. Soc. Change* **2010**, *77* (3), 355–368.

Grady, D.; McNeil Jr, D. G. Ebola Sample Is Mishandled at C.D.C. Lab. *New York Times*. New York December 25, 2014, p A1.

Griliches, Z. Issues in Assessing the Contribution of Research and Development to Productivity Growth. *Bell J. Econ.* **1979**, *10*, 92–116.

Griliches, Z. Productivity, R&D, and the Data Constraint. *Am. Econ. Rev.* **1994**, *84* (1), 1–23.

Grinstead, C. M.; Snell, J. L. *Introduction to Probability*; American Mathematical Society, 1997.

Gross, L. A Broken Trust: Lessons from the Vaccine--Autism Wars. *PLoS Biol.* **2009**, *7* (5), e1000114.

Gruner, S. New Ethics of Technology: A Discourse in Germany in the 1990s. In *Annual Conference of the Philosophical Society of Southern Africa*; Pretoria, 2008; p 15.

Gurwitz, D. Gene Drives Raise Dual-Use Concerns. *Science* **2014**, *345* (6200), 1010.

Gutierrez, R.; Giner-Sorolla, R. Anger, Disgust, and Presumption of Harm as Reactions to Taboo-Breaking Behaviors. *Emotion* **2007**, *7* (4), 853–868.

Haas, C. N. The Role of Risk Analysis in Understanding Ebola. *Risk Anal.* **2015**, *35* (2), 183–185.

Haimes, Y. Y. Hierarchical Holographic Modeling. *IEEE Trans. Syst. Man. Cybern.* **1981**, *11* (9), 606–617.

Haimes, Y. Y.; Kaplan, S.; Lambert, J. H. Risk Filtering, Ranking, and Management Framework Using Hierarchical Holographic Modeling. *Risk Anal.* **2002**, *22* (2), 383–397.

Hall, B. H.; Mairesse, J.; Mohnen, P. *Measuring the Returns to R&D*; 15622; National Bureau of Economic Research: Cambridge, MA, 2009.

Hall, C. A. S. An Assessment of Several of the Historically Most Influential Theoretical Models Used in Ecology and of the Data Provided in Their Support. *Ecol. Modell.* **1988**, *43* (1-2), 5–31.

Hamilton, G. The Mitochondria Mystery. *Nature* **2015**, *525* (7570), 444–446.

Hammitt, J. K.; Shlyakhtel, A. I. The Expected Value of Information and the Probability of Surprise. *Risk Anal.* **1999**, *19* (1), 135–152.

Hammitt, J. K.; Zhang, Y. Combining Experts' Judgments: Comparison of Algorithmic Methods Using Synthetic Data. *Risk Anal.* **2013**, *33* (1), 109–120.

Hansson, S. O. The Moral Significance of Indetectable Effects. *Risk Heal. Saf. Environ.* **1999**, *10*, 101–108.

Hansson, S. O. Are Natural Risks Less Dangerous than Technological Risks? *Philos. Nat.* **2003**, *40*, 43–54.

Hansson, S. O. Philosophical Perspectives on Risk. *Techné Res. Philos. Technol.* **2004**, *8* (1), 10–35.

Hansson, S. O. Seven Myths of Risk. *Risk Manag.* **2005**, *7* (2), 7–17.

Hansson, S. O. Do We Need Second-Order Probabilities? *Dialectica* **2008**, *62* (4), 525–533.

Hansson, S. O. An Agenda for the Ethics of Risk. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009a; pp 11–23.

Hansson, S. O. From the Casino to the Jungle. *Synthese* **2009b**, *168* (3), 423–432.

Hansson, S. O. Promoting Inherent Safety. *Process Saf. Environ. Prot.* **2010a**, *88* (3), 168–172.

Hansson, S. O. Risk: Objective or Subjective, Facts or Values. *J. Risk Res.* **2010b**, *13* (2), 231–238.

Hansson, S. O. A Panorama of the Philosophy of Risk. In *Handbook of Risk Theory*; Roeser, S., Hillerbrand, R., Sandin, P., Peterson, M., Eds.; Springer: Dordrecht, 2012; pp 27–54.

Hansson, S. O.; Rudén, C. Evaluating the Risk Decision Process. *Toxicology* **2006**, *218* (2-3), 100–111.

Hansson, S. O.; Aven, T. Is Risk Analysis Scientific? *Risk Anal.* **2014**, *34* (7), 1173–1183.

Harsanyi, J. Can the Maximin Principle Serve as a Basis for Morality? *Am. Polit. Sci. Rev.* **1975**, *69* (2), 594–606.

Hayden, E. C. Synthetic-Biology Firms Shift Focus. *Nature* **2014**, *505* (7485), 598.

Hayes, K. R.; Barry, S.; Beebe, N.; Dambacher, J. M.; Barro, P. De; Ferson, S.; Ford, J.; Foster, S.; Goncalves da Silva, A.; Hosack, G. R.; et al. *Risk Assessment for Controlling Mosquito Vectors with Engineered Nucleases: Sterile Male Construct Final Report*; CSIRO Biosecurity Flagship: Hobart, Australia, 2015.

Heidegger, M. *The Question Concerning Technology and Other Essays*, English Tr.; Lovitt, W., Ed.; Garland: New York, 1977.

Hempel, C. G. Inductive Inconsistencies. *Synthese* **1960**, *12* (4), 439–469.

Henderson, D. A.; Arita, I. The Smallpox Threat: A Time to Reconsider Global Policy. *Biosecurity and bioterrorism* **2014**, *12* (3), 1–5.

Henkel, R. D.; Miller, T.; Weyant, R. S. Monitoring Select Agent Theft, Loss and Release Reports in the United States — 2004-2010. *Appl. Biosaf.* **2012**, *17* (4), 171–180.

Henkler, F.; Luch, A. Safety Regulations: More Extensive Tests for E-Cigarettes. *Nature* **2015**, *525* (7568), 187.

Henry, H.; Taylor, A. Re-Thinking Apollo: Envisioning Environmentalism in Space. *Sociol. Rev.* **2009**, *57* (s1), 190–203.

Herfst, S.; Schrauwen, E. J. a; Linster, M.; Chutinimitkul, S.; de Wit, E.; Munster, V. J.; Sorrell, E. M.; Bestebroer, T. M.; Burke, D. F.; Smith, D. J.; et al. Airborne Transmission of Influenza A/H5N1 Virus between Ferrets. *Science* **2012**, *336* (6088), 1534–1541.

Hermansson, H.; Hansson, S. O. A Three-Party Model Tool for Ethical Risk Analysis. *Risk Manag.* **2007**, *9* (3), 129–144.

Hertzfeld, H. R. Measuring the Economic Impact of Federal Research and Development Investments in Civilian Space Activities. In *Papers Commissioned for a Workshop on the Federal Role in Research and Development*; National Academies Press, 1985; p 265.

HHS. Possession, Use, and Transfer of Select Agents and Toxins; Addition of Certain Influenza Virus Strains to the List of Select Agents and Toxins https://www.federalregister.gov/articles/2015/07/16/2015-17435/possession-use-and-transfer-of-select-agents-and-toxins-addition-of-certain-influenza-virus-strains (accessed Oct 17, 2015).

Hinxton Group. *Consensus Statement on Genome Editing Technologies and Human Germline Genetic Modification*; The Johns Hopkins Berman Institute of Bioethics: Baltimore, MD, 2015.

Hites, R. A.; Foran, J. A.; Carpenter, D. O.; Hamilton, M. C.; Knuth, B. A.; Schwager, S. J. Global Assessment of Organic Contaminants in Farmed Salmon. *Science* **2004**, *303* (5655), 226–229.

Hoeting, J. A.; Madigan, D.; Raftery, A. E.; Volinsky, C. T. Bayesian Model Averaging: A Tutorial (with Comments by M. Clyde, David Draper and E. I. George, and a Rejoinder by the Authors). *Stat. Sci.* **1999**, *14* (4), 382–417.

Holden, W. L. Benefit-Risk Analysis: A Brief Review and Proposed Quantitative Approaches. *Drug Saf.* **2003**, *26* (12), 853–862.

Hollanders, H.; Es-Sadki, N. *Innovation Union Scoreboard 2014*; European Commission, UNU-MERIT, 2014.

Hotez, P. J. Combating the next Lethal Epidemic. *Science* **2015**, *348* (6232), 296–297.

Howard, R. A.; Matheson, J. E. Influence Diagram Retrospective. *Decis. Anal.* **2005**, *2* (3), 144–147.

Hulme, P. E. Biosecurity's Weakest Link. *Science* **2013**, *340* (6129), 141–142.

Hvistendahl, M. Enigmatic Bird Flu Strain Races across the U.S. Midwest. *Science* **2015**, *348* (6236), 741–742.

Imai, M.; Watanabe, T.; Hatta, M.; Das, S. C.; Ozawa, M.; Shinya, K.; Zhong, G.; Hanson, A.; Katsura, H.; Watanabe, S.; et al. Experimental Adaptation of an Influenza H5 HA Confers Respiratory Droplet Transmission to a Reassortant H5 HA/H1N1 Virus in Ferrets. *Nature* **2012**, *486* (7403), 420–428.

Imperiale, M. J.; Casadevall, A. Vagueness and Costs of the Pause on Gain-of-Function (GOF) Experiments on Pathogens with Pandemic Potential, Including Influenza Virus. *MBio* **2014**, *5* (6), e02292–14.

Ioannidis, J. P. A. Why Most Published Research Findings Are False. *PLoS Med.* **2005**, *2* (8), e124.

Jaffe, A. B.; Trajtenberg, M.; Henderson, R. Geographic Localization of Knowledge Spillovers as Evidence by Patent Citations. *Q. J. Econ.* **1993**, *108* (3), 577–598.

Jaffe, H.; Patterson, A. P.; Lurie, N. Extra Oversight for H7N9 Experiments. *Science* **2013**, *341* (6147), 713–714.

Jasanoff, S. Bridging the Two Cultures of Risk Analysis. *Risk Anal.* **1993**, *13* (2), 123–129.

Jasanoff, S. Technologies of Humility: Citizen Participation in Governing Science. *Minerva* **2003**, *41* (3), 223–244.

Jasanoff, S. *States of Knowledge: The Co-Production of Science and Social Order*; Routledge: London, 2004.

Jasanoff, S. Technologies of Humility. *Nature* **2007**, *450* (7166), 33.

Jebari, K. Existential Risks: Exploring a Robust Risk Reduction Strategy. *Sci. Eng. Ethics* **2014**.

Jeffrey, R. C. Valuation and Acceptance of Scientific Hypotheses. *Philos. Sci.* **1956**, *23* (3), 237–246.

Jha, A. Dan Shechtman: "Linus Pauling said I was talking nonsense" http://www.theguardian.com/science/2013/jan/06/dan-shechtman-nobel-prize-chemistry-interview.

Johnson, M. *Moral Imagination: Implications of Cognitive Science for Ethics*; University of Chicago Press: Chicago, 1993; Vol. 1.

Johnson, R. W. Analyze Hazards, Not Just Risks. *Chem. Eng. Prog.* **2000**, *96* (7), 31–40.

Johnston, C. Womb transplants: first 10 British women given go-ahead http://www.theguardian.com/lifeandstyle/2015/sep/29/10-women-receive-go-ahead-for-first-ever-womb-transplants-in-uk.

Jonas, H. *The Imperative of Responsibility: In Search of an Ethics for the Technological Age*; University of Chicago Press: Chicago, 1984.

Jones, B. F. Age and Great Invention. *Rev. Econ. Stat.* **2010**, *92* (1), 1–14.

Jones, C. I.; Williams, J. C. Measuring the Social Return to R&D. *Q. J. Econ.* **1998**, *113* (4), 1119–1135.

Jones, S. E. *Against Technology: From the Luddites to Neo-Luddism*; Routledge: New York, 2006.

Joy, B. Why the Future Doesn't Need Us. *Wired.* April 2000, p 8.04.

Kafouros, M. I.; Buckley, P. J. Under What Conditions Do Firms Benefit from the Research Efforts of Other Organizations? *Res. Policy* **2008**, *37* (2), 225–239.

Kahan, D. Fixing the Communications Failure. *Nature* **2010**, *463* (7279), 296–297.

Kahan, D.; Slovic, P.; Braman, D.; Gastil, J. Fear of Democracy: A Cultural Evaluation of Sunstein on Risk. *Harv. Law Rev.* **2006**, *119* (4), 1071–1109.

Kahan, D.; Braman, D.; Cohen, G. L.; Gastil, J.; Slovic, P. Who Fears the HPV Vaccine, Who Doesn't, and Why? An Experimental Study of the Mechanisms of Cultural Cognition. *Law Hum. Behav.* **2010**, *34* (6), 501–516.

Kahan, D.; Jenkins- Smith, H.; Braman, D. Cultural Cognition of Scientific Consensus. *J. Risk Res.* **2011**, *14* (2), 147–174.

Kahan, D.; Jenkins-Smith, H.; Tarantola, T.; Silva, C. L.; Braman, D. Geoengineering and Climate Change Polarization: Testing a Two-Channel Model of Science Communication. *Ann. Am. Acad. Pol. Soc. Sci.* **2015**, *658* (1), 192–222.

Kahlor, L.; Dunwoody, S.; Griffin, R. J.; Neuwirth, K.; Giese, J. Studying Heuristic-Systematic Processing of Risk Communication. *Risk Anal.* **2003**, *23* (2), 355–368.

Kahn, L. H. Can Biosecurity Be Embedded into the Culture of the Life Sciences? *Biosecurity and bioterrorism* **2012**, *10* (2), 241–246.

Kaiser, D.; Moreno, J. D. Self-Censorship Is Not Enough. *Nature* **2012**, *492* (7429), 345–347.

Kaiser, J. A Call for NIH Youth Movement. *Science* **2014a**, *346* (6206), 150–151.

Kaiser, J. NIH Institute Considers Broad Shift to "people" Awards. *Science* **2014b**, *345* (6195), 366–367.

Kaiser, J. The Catalyst. *Science* **2014c**, *345* (6201), 1112–1115.

Kaiser, J. U.S. Halts Two Dozen Risky Virus Studies. *Science* **2014d**, *346* (628), 404.

Kaiser, J.; Normile, D. Embryo Engineering Study Splits Scientific Community. *Science* **2015**, *348* (6234), 486–487.

Kanji, G. K. *100 Statistical Tests*, Third Edit.; SAGE Publications: London, 2006.

Kant, I. *Groundwork for the Metaphysics of Morals*, 2002 Engli.; Wood, A. W., Ed.; Yale University Press: New Haven, CT, 1785.

Kant, L.; Mourya, D. T. Managing Dual Use Technology: It Takes Two to Tango. *Sci. Eng. Ethics* **2010**, *16* (1), 77–83.

Kantrowitz, A. Proposal for an Institution for Scientific Judgment. *Science* **1967**, *156* (3776), 763–764.

Kasperson, R. E.; Renn, O.; Slovic, P.; Brown, H. S.; Emel, J.; Goble, R.; Kasperson, J. X.; Ratick, S. The Social Amplification of Risk: A Conceptual Framework. *Risk Anal.* **1988**, *8* (2), 177–187.

Ke, Q.; Ferrara, E.; Radicchi, F.; Flammini, A. Defining and Identifying Sleeping Beauties in Science. *Proc. Natl. Acad. Sci.* **2015**, 1–6.

Keisler, J. M.; Collier, Z. A.; Chu, E.; Sinatra, N.; Linkov, I. Value of Information Analysis: The State of Application. *Environ. Syst. Decis.* **2013**, *34* (1), 3–23.

Keith, D. W. When Is It Appropriate to Combine Expert Judgments? *Clim. Change* **1996**, *33* (2), 139–143.

Kelle, A. Beyond Patchwork Precaution in the Dual-Use Governance of Synthetic Biology. *Sci. Eng. Ethics* **2013**, *19* (3), 1121–1139.

Kelly, K. *What Technology Wants*; Viking Press: New York, 2010.

Kelman, S. Cost-Benefit Analysis: An Ethical Critique. *AEI J. Gov. Soc. Regul.* **1981**, *5* (1), 33–40.

Keren, A. Kitcher on Well-Ordered Science: Should Science Be Measured against the Outcomes of Ideal Democratic Deliberation? *Theoria* **2013**, *28* (77), 233–244.

Kerr, R. A. Planetary Scientists Casting Doubt On Feasibility of Plan to Corral Asteroid. *Science* **2013**, *340* (6133), 668–669.

Keynes, J. M. *The Principle of Indifference*; Macmillian and Co.: London, UK, 1921.

Khan, F. I.; Amyotte, P. R. How to Make Inherent Safety Practice a Reality. *Can. J. Chem. Eng.* **2003**, *81* (1), 2–16.

Kidd, D. C.; Castano, E. Reading Literary Fiction Improves Theory of Mind. *Science* **2013**, *342* (6156), 377–380.

Kim, J.; Yeo, S. K.; Brossard, D.; Scheufele, D. A.; Xenos, M. A. Disentangling the Influence of Value Predispositions and Risk/benefit Perceptions on Support for Nanotechnology among the American Public. *Risk Anal.* **2014**, *34* (5), 965–980.

Kiran, A. H.; Verbeek, P.-P. Trusting Our Selves to Technology. *Knowledge, Technol. Policy* **2010**, *23* (3-4), 409–427.

Kitcher, P. *Science, Truth, and Democracy*; Oxford University Press: Oxford, 2001.

Klemes, V. Operational Testing of Hydrological Simulation Models. *Hydrol. Sci. J.* **1986**, *31* (1), 13–24.

Kletz, T. A. What You Don't Have, Can't Leak. *Chem. Ind.* **1978**, *6*, 287–292.

Kletz, T. A. Inherently Safer Plants. *Plant/Operations Prog.* **1985**, *4* (3), 164–167.

Klotz, L. C. Comments on Fouchier's Calculation of Risk and Elapsed Time for Escape of a Laboratory-Acquired Infection from His Laboratory. *MBio* **2015**, *6* (2), e00268–15.

Klotz, L. C.; Sylvester, E. J. The unacceptable risks of a man-made pandemic http://thebulletin.org/unacceptable-risks-man-made-pandemic.

Klotz, L. C.; Sylvester, E. J. The Consequences of a Lab Escape of a Potential Pandemic Pathogen. *Front. Public Heal.* **2014**, *2*, 116.

Koehn, D. A Role for Virtue Ethics in the Analysis of Business Practice. *Bus. Ethics Q.* **1995**, *5* (3), 533–539.

Kraemer, J.; Gostin, L. The Limits of Government Regulation of Science. *Science* **2012**, *335* (6072), 1047–1049.

Krayer von Krauss, M. P.; Walker, W. E.; van der Sluijs, J. P.; Janssen, P.; van Asselt, M. B. A.; Rotmans, J. Response to "To What Extent, and How, Might Uncertainty Be Defined" by Norton, Brown, and Mysiak. *Integr. Assess.* **2006**, *6* (1), 89–94.

Krueger, A.; Stone, A. Progress in Measuring Subjective Well-Being. *Science* **2014**, *346* (6205), 42–43.

Kuhn, T. S. *The Structure of Scientific Revolutions.*; University of Chicago Press: Chicago, 1962.

Kuiken, T. Shaping the Future of Synthetic Biology. *Science* **2015**, *348* (6232), 296.

Kun, H. Moving Away from Metrics. *Nature* **2015**, *520* (7549), S18–S20.

Kupferschmidt, K. Rules of the Name. *Science* **2015**, *348* (6236), 745.

Kuzma, J.; Besley, J. C. Ethics of Risk Analysis and Regulatory Review: From Bio- to Nanotechnology. *Nanoethics* **2008**, *2* (2), 149–162.

Labbé, C.; Labbé, D. Duplicate and Fake Publications in the Scientific Literature: How Many SCIgen Papers in Computer Science? *Scientometrics* **2013**, *94* (1), 379–396.

Lach, S.; Schankerman, M. Dynamics of R&D and Investment in the Scientific Sector. *J. Polit. Econ.* **1989**, *97* (4), 880–904.

Lakatos, I. Falsification and the Methodology of Scientific Research Programmes. In *Criticism and the Growth of Knowledge*; Lakatos, I., Musgrave, A., Eds.; Cambridge University Press: Cambridge, UK, 1970; pp 91–196.

Lake, J. A. Lynn Margulis (1938–2011). *Nature* **2011**, *480* (7378), 458–458.

Lane, J. Assessing the Impact of Science Funding. *Science* **2009**, *324* (5932), 1273–1275.

Lane, J.; Bertuzzi, S. Measuring the Results of Science Investments. *Science* **2011**, *331* (6018), 678–680.

Langlois, R. A.; Albrecht, R. A.; Kimble, B.; Sutton, T.; Shapiro, J. S.; Finch, C.; Angel, M.; Chua, M. A.; Gonzalez-Reiche, A. S.; Xu, K.; et al. MicroRNA-Based Strategy to Mitigate the Risk of Gain-of-Function Influenza Studies. *Nat. Biotechnol.* **2013**, *31* (9), 844–847.

Lanphier, E.; Urnov, F.; Haecker, S. E.; Werner, M.; Smolenski, J. Don't Edit the Human Germ Line. *Nature* **2015**, *519* (7544), 410–411.

Largent, M. A.; Lane, J. I. STAR METRICS and the Science of Science Policy. *Rev. Policy Res.* **2012**, *29* (3), 431–438.

Ledford, H. CRISPR, the Disruptor. *Nature* **2015**, *522* (7554), 20–24.

Lee, J.; Schmidt, A. G. Research and Development Satellite Account Update Estimates for 1959 – 2007. *Surv. Curr. Bus.* **2010**, *90* (12), 16–27.

Lee, S.-W.; Markham, P. F.; Coppo, M. J. C.; Legione, A. R.; Markham, J. F.; Noormohammadi, A. H.; Browning, G. F.; Ficorilli, N.; Hartley, C. A.; Devlin, J. M. Attenuated Vaccines Can Recombine to Form Virulent Field Viruses. *Science* **2012**, *337* (6091), 188.

Leek, J. T.; Peng, R. D. Statistics: P Values Are Just the Tip of the Iceberg. *Nature* **2015**, *520* (7549), 612–612.

Leitenberg, M.; Zilinskas, R. A.; Kuhn, J. H. *The Soviet Biological Weapons Program: A History*; Harvard Univsersity Press: Cambridge, MA, 2012.

Leopold, A. *A Sand County Almanac: And Sketches Here and There*; Oxford University Press, 1949.

Leshner, A. I. Maintaining Science Eminence. *Science* **2013**, *341* (6148), 820.

Lewnard, J. A.; Ndeffo Mbah, M. L.; Alfaro-Murillo, J. A.; Altice, F. L.; Bawo, L.; Nyenswah, T. G.; Galvani, A. P. Dynamics and Control of Ebola Virus Transmission in Montserrado, Liberia: A Mathematical Modelling Analysis. *Lancet. Infect. Dis.* **2014**, *14* (12), 1189–1195.

Li, D.; Agha, L. Big Names or Big Ideas: Do Peer-Review Panels Select the Best Science Proposals? *Science* **2015**, *348* (6233), 434–438.

Li, W. C. Y. *Depreciation of Business R&D Capital*; Bureau of Economic Analysis/U.S. Department of Commerce, 2012.

Liang, P.; Xu, Y.; Zhang, X.; Ding, C.; Huang, R.; Zhang, Z.; Lv, J.; Xie, X.; Chen, Y.; Li, Y.; et al. CRISPR/Cas9-Mediated Gene Editing in Human Tripronuclear Zygotes. *Protein Cell* **2015**, *6* (5), 363–372.

Linkov, I.; Bridges, T.; Creutzig, F.; Decker, J.; Fox-Lent, C.; Kröger, W.; Lambert, J. H.; Levermann, A.; Montreuil, B.; Nathwani, J.; et al. Changing the Resilience Paradigm. *Nat. Clim. Chang.* **2014**, *4* (6), 407–409.

Linster, M.; van Boheemen, S.; de Graaf, M.; Schrauwen, E. J. A.; Lexmond, P.; Mänz, B.; Bestebroer, T. M.; Baumann, J.; van Riel, D.; Rimmelzwaan, G. F.; et al. Identification, Characterization, and Natural Selection of Mutations Driving Airborne Transmission of A/H5N1 Virus. *Cell* **2014**, *157* (2), 329–339.

Lipsitch, M. Can Limited Scientific Value of Potential Pandemic Pathogen Experiments Justify the Risks? *MBio* **2014**, *5* (5), e02008–e02014.

Lipsitch, M.; Galvani, A. P. Ethical Alternatives to Experiments with Novel Potential Pandemic Pathogens. *PLoS Med.* **2014**, *11* (5), e1001646.

Lipsitch, M.; Inglesby, T. V. Moratorium on Research Intended To Create Novel Potential Pandemic Pathogens. *MBio* **2014**, *5* (6), e02366–14.

Lipsitch, M.; Inglesby, T. V. Reply to " Studies on Influenza Virus Transmission between Ferrets : The Public Health Risks Revisited ." *MBio* **2015**, *6* (1), e00041–15.

Lipsitch, M.; Esvelt, K.; Inglesby, T. Calls for Caution in Genome Engineering Should Be a Model for Similar Dialogue on Pandemic Pathogen Research. *Ann. Intern. Med.* **2015**.

Loewenstein, G. F.; Weber, E. U.; Hsee, C. K.; Welch, N. Risk as Feelings. *Psychol. Bull.* **2001**, *127* (2), 267–286.

Loewenstein, G.; Moore, D. A.; Weber, R. A. Paying $1 to Lose $2: Misperceptions of the Value of Information in Predicting the Performance of Others. *Acad. Manag. Proc.* **2003**, *2003* (1), C1–C6.

Lowrie, H.; Tait, J. *Policy Brief: Guidelines for the Appropriate Risk Governance of Synthetic Biology*; International Risk Governance Council: Geneva, 2010.

Lund, E.; Engeset, D.; Alsaker, E.; Skeie, G.; Hjårtaker, A.; Lundebye, A.-K.; Niebor, E. Cancer Risk and Salmon Intake. *Science* **2004**, *305* (5683), 477–478; author reply 477–478.

Lunshof, J. Regulate Gene Editing in Wild Animals. *Nature* **2015**, *521* (7551), 127.

Lymn, N.; Houston, T. Taking Aim at Science. *Front. Ecol. Environ.* **2013**, *11* (5), 227–227.

Lyttleton, R. A. The Nature of Knowledge. In *The Encyclopaedia of Ignorance: Everything you ever wanted to know about the unknown*; Duncan, R., Weston-Smith, M. M., Eds.; Pocket Books: New York, 1978; p 443.

MacGillivray, B. H. Heuristics Structure and Pervade Formal Risk Assessment. *Risk Anal.* **2014**, *34* (4), 771–787.

Macilwain, C. Halt the Avalanche of Performance Metrics. *Nature* **2013a**, *500* (7462), 255.

Macilwain, C. The Unlikely Wisdom of Chairman Mao. *Nature* **2013b**, *495* (7440), 143–143.

Macilwain, C. Beware of Backroom Deals in the Name of "Science". *Nature* **2014**, *508* (7496), 289.

MacKenzie, C. A. Summarizing Risk Using Risk Measures and Risk Indices. *Risk Anal.* **2014**, *34* (12), 2143–2162.

MacLean, D. *Values at Risk*; Rowman & Littlefield: Totawa, NJ, 1986.

MacLean, D. Ethics, Reasons and Risk Analysis. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 115–143.

Maher, B. Bias Accusation Rattles US Biosecurity Board. *Nature* **2012a**.

Maher, B. The Biosecurity Oversight. *Nature* **2012b**, *485* (7399), 431–434.

Mahmoud, A. Gain-of-Function Research: Unproven Technique. *Science* **2013**, *342* (6156), 310–311.

Malakoff, D. U. S. Agencies to Start Screening Biomedical Proposals for Dual Use. *Science* **2012**, *336* (6077), 21.

Malakoff, D.; Enserink, M. New U . S . Rules Increase Oversight Of H5N1 Studies, Other Risky Science. *Science* **2013**, *339* (6123), 1025.

Malyshev, D. a; Dhami, K.; Lavergne, T.; Chen, T.; Dai, N.; Foster, J. M.; Corrêa, I. R.; Romesberg, F. E. A Semi-Synthetic Organism with an Expanded Genetic Alphabet. *Nature* **2014**, *509* (7500), 385–388.

Mandell, D. J.; Lajoie, M. J.; Mee, M. T.; Takeuchi, R.; Kuznetsov, G.; Norville, J. E.; Gregg, C. J.; Stoddard, B. L.; Church, G. M. Biocontainment of Genetically Modified Organisms by Synthetic Protein Design. *Nature* **2015**, *518* (7537), 55–60.

Mander, J. *Four Arguments for the Elimination of Television*; HarperCollins: New York, 1978.

Marburger, J. H. I. *Keynote Address*; Annual American Association for the Advancement of Science Forum on Science and Technology Policy: Washington, DC, 2005.

Marcus, D. M.; Grollman, A. P. Botanical Medicines - The Need for New Regulations. *N. Engl. J. Med.* **2002**, *347* (25), 2073–2076.

Margottini, L. Police Claim Leading Flu Researcher Illegally Sold Virus. *Science* **2014**, *345* (6201), 1105–1106.

Marraffini, L. A. CRISPR-Cas Immunity in Prokaryotes. *Nature* **2015**, *526* (7571), 55–61.

Marschak, J. Toward an Economic Theory of Organization and Information. In *Decision Processes*; Thrall, R. M., Coombs, C. H., Davis, R. L., Eds.; Wiley: New York, 1954; pp 187–220.

Marschak, J. Efficient and Viable Organization Forms. In *Modern Organization Theory*; Haire, M., Ed.; Wiley: New York, 1959; pp 307–320.

Martikainen, M.; Pedersen, O. Germline Edits: Heat Does Not Help Debate. *Nature* **2015**, *520* (7549), 623.

Mazzucato, M. *The Entrepreneurial State*; Demos: London, 2011.

Mcculla, S. H.; Holdren, A. E.; Smith, S. Improved Estimates of the National Income and Product Accounts. *Surv. Curr. Bus.* **2013**, *93* (9), 14–45.

McGilvray, A. The Limits of Excellence. *Nature* **2014**, *511* (5710), S64–S66.

Mcleary, P. Pentagon Stops Anthrax Shipments As Scandal Grows https://foreignpolicy.com/2015/06/02/pentagon-lethal-anthrax-shipments/ (accessed Mar 6, 2015).

McNeil Jr, D. G. C.D.C. Shuts Labs After Accidents With Pathogens. *New York Times*. July 12, 2014a, p A1.

McNeil Jr, D. G. Leader of Troubled Lab Steps Down, C.D.C. Says. *New York Times*. New York July 24, 2014b, p A17.

Medford, J. I.; Prasad, A. Plant Synthetic Biology Takes Root. *Science* **2014**, *346* (6206), 162–163.

Meredith, W. A Workforce out of Balance. *Nature* **2012**, *486*, 304.

Merler, S.; Ajelli, M.; Fumanelli, L.; Vespignani, A. Containing the Accidental Laboratory Escape of Potential Pandemic Influenza Viruses. *BMC Med.* **2013**, *11*, 252.

Merton, R. K. Science and Technology in a Democratic Order. *J. Leg. Polit. Sociol.* **1942**, *1*, 115–126.

Mervis, J. A Radical Change in Peer Review. *Science* **2014a**, *345* (6194), 248–249.

Mervis, J. Does Science Suffer an "innovation Deficit"? *Science* **2014b**, *346* (6208), 408–409.

Mervis, J. Peering Into Peer Review. *Science* **2014c**, *343* (6171), 596–598.

Mervis, J. Play It Again, Uncle Sam. *Science* **2014d**, *345* (6203), 1442.

Ter Meulen, V. Time to Settle the Synthetic Controversy. *Nat. News* **2014**, *509* (7499), 135.

Meyer, A. On the Nature of Scientific Progress: Anarchistic Theory Says "Anything Goes"—But I Don't Think So. *PLoS Biol.* **2011**, *9* (10), e1001165.

Miller, H. I. Germline Gene Therapy: We're Ready. *Science* **2015**, *348* (6241), 1325–1325.

Miller, S.; Oda, L. M.; Chakhava, G.; Rao, C. K.; Temsamani, K. R.; Lentzos, F.; Nasim, A.; Makalino, I. R.; Gorski, A.; Nampala, P. Sustaining Progress in the Life Sciences: Strategies for Managing Dual Use Research of Concern--Progress at the National Level. *Biosecur. Bioterror.* **2009**, *7* (1), 93–100.

Minteer, B. Is It Right to Reverse Extinction? *Nature* **2014**, *509* (7500), 261.

Molas-Gallart, J. Which Way to Go? Defence Technology and the Diversity of "Dual-Use" Technology Transfer. *Res. Policy* **1997**, *26* (3), 367–385.

Möller, N. The Concepts of Risk and Safety. In *Handbook of Risk Theory*; Roeser, S., Hillerbrand, R., Sandin, P., Peterson, M., Eds.; Springer: Dordrecht, 2012; p 1187.

Möller, N.; Hansson, S. O. Principles of Engineering Safety: Risk and Uncertainty Reduction. *Reliab. Eng. Syst. Saf.* **2008**, *93* (6), 798–805.

Mollison, D.; Din, S. U. Deterministic and Stochastic Models for the Seasonal Variability of Measles Transmission. *Math. Biosci.* **1993**, *117* (1-2), 155–177.

Montibeller, G.; von Winterfeldt, D. Cognitive and Motivational Biases in Decision and Risk Analysis. *Risk Anal.* **2015**, *35* (7), 1230–1251.

Moodie, M. Managing the Security Risks of Emerging Technologies. *Arms Control Today* **2012**, *42* (September).

Mooney, C. *Do Scientists Understand the Public?*; American Academy of Arts & Sciences: Cambridge, MA, 2010.

Morello, L. More Cuts Loom for US Science. *Nature* **2013**, *501* (7466), 147–148.

Morens, D. M.; Subbarao, K.; Taubenberger, J. K. Engineering H5N1 Avian Influenza Viruses to Study Human Adaptation. *Nature* **2012**, *486* (7403), 335–340.

Morens, D. M.; Taubenberger, J. K.; Fauci, A. S. Pandemic Influenza Viruses — Hoping for the Road Not Taken. *N. Engl. J. Med.* **2013**, *368* (25), 2345–2348.

Morgan, M. G. Our Knowledge of the World Is Often Not Simple: Policymakers Should Not Duck That Fact, But Should Deal with It. *Risk Anal.* **2015**, *35* (1), 19–20.

Morgan, M. G.; Henrion, M.; Small, M. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*; Cambridge University Press: New York, 1990.

Moser, C.; Stauffacher, M.; Krütli, P.; Scholz, R. W. The Crucial Role of Nomothetic and Idiographic Conceptions of Time: Interdisciplinary Collaboration in Nuclear Waste Management. *Risk Anal.* **2012**, *32* (1), 138–154.

Mosleh, A.; Bier, V. M. Uncertainty about Probability: A Reconciliation with the Subjectivist Viewpoint. *IEEE Trans. Syst. Man, Cybern. - Part A Syst. Humans* **1996**, *26* (3), 303–310.

Muir, J. *My First Summer in the Sierra*; Houghton Mifflin: Boston, 1911.

Muller, N. Boosting GDP Growth by Accounting for the Environment. *Science* **2014**, *345* (6199), 873–874.

Naess, A. The Shallow and the Deep, Long-Range Ecology Movement. A Summary. *Inquiry* **1973**, *16* (1-4), 95–100.

Nakayachi, K. The Unintended Effects of Risk-Refuting Information on Anxiety. *Risk Anal.* **2013**, *33* (1), 80–91.

Nelson, B. Hacking into the Cyberworld. *Nature* **2014**, *506* (7489), 517–519.

Nichols, A. L.; Zeckhauser, R. J. The Perils of Prudence: How Conservative Risk Assessments Distort Regulation. *Regul. Toxicol. Pharmacol.* **1988**, *8* (1), 61–75.

Nightingale, P.; Martin, P. The Myth of the Biotech Revolution. *Trends Biotechnol.* **2004**, *22* (11), 564–569.

NIH. Conducting Risk and Benefit Analysis on Gain-of Function Research Involving Pathogens with Pandemic Potential https://www.fbo.gov/?s=opportunity&mode=form&id=c134018fd1d008c582b7755be1fc1c06&tab=core&_cview=0 (accessed Apr 8, 2015).

NIH. Conducting Risk and Benefit Analysis on Gain-of Function Research Involving Pathogens with Pandemic Potential https://www.fbo.gov/index?s=opportunity&mode=form&id=88318d41facba57c7a947b90223cea3d&tab=core&tabmode=list&= (accessed Apr 8, 2015).

Van Noorden, R. Seven Thousand Stories Capture Impact of Science. *Nature* **2015**, *518* (7538), 150–151.

Nordhaus, W. D. A Review of the Stern Review on the Economics of Climate Change. *J. Econ. Lit.* **2007**, *45* (3), 686–702.

Norton, J. P.; Brown, J. D.; Mysiak, J. To What Extent, and How, Might Uncertainty Be Defined? Comments Engendered by "Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support." *Integr. Assess.* **2006**, *6* (1), 83–88.

Norton, M. I.; Mochon, D.; Ariely, D. The IKEA Effect: When Labor Leads to Love. *J. Consum. Psychol.* **2012**, *22* (3), 453–460.

Nozick, R. *Anarchy, State, and Utopia*; Basic Books: Philidelphia, 1974.

NSABB. *Proposed Framework for the Oversight of Dual Use Life Sciences Research : Strategies for Minimizing the Potential Misuse of Research Information*; National Science Advisory Board for Biosecurity: Bethesda, MD, 2007.

NSABB. *Attributes of Highly Pathogenic Avian Influenza H5N1 Research That May Warrant Alternative Venues or Modes of Communication*; National Science Advisory Board for Biosecurity: Bethesda, MD, 2012a.

NSABB. *March 29- 30, 2012 Meeting of the National Science Advisory Board for Biosecurity to Review Revised Manuscripts on Transmissibility of A/H5N1 Influenza Virus*; National Science Advisory Board for Biosecurity: Bethesda, MD, 2012b.

NSABB. NSABB Meeting November 25 http://osp.od.nih.gov/office-biotechnology-activities/event/2014-11-25-160000-2014-11-25-180000/nsabb-meeting.

NSABB. NSABB Meeting October 22 http://osp.od.nih.gov/office-biotechnology-activities/event/2014-10-22-121500-2014-10-22-200000/nsabb-meeting.

NSABB. NSABB May 5 Draft Meeting Agenda http://osp.od.nih.gov/office-biotechnology-activities/event/2015-05-05-120000-2015-05-05-210000/national-science-advisory-board-biosecurity-nsabb-meeting.

NSTC. *The Science of Science Policy: A Federal Research Roadmap*; National Science And Technology Council, 2008.

Nussbaum, M. C. *Love's Knowledge*; Oxford University Press: New York, 1990.

Nussbaum, M. C. Virtue Ethics: A Misleading Category? *J. Ethics* **1999**, *3* (3), 163–201.

Nussbaum, M. C. Moral Expertise?: Constitutional Narratives and Philosophical Argument. *Metaphilosophy* **2002**, *33* (5), 502–520.

Nuzzo, R. Scientific Method: Statistical Errors. *Nature* **2014**, *506* (7487), 150–152.

Nuzzo, R. How Scientists Fool Themselves – and How They Can Stop. *Nature* **2015**, *526* (7572), 182–185.

O'Hagan, A.; Oakley, J. E. Probability Is Perfect, but We Can't Elicit It Perfectly. *Reliab. Eng. Syst. Saf.* **2004**, *85* (1-3), 239–248.

Oeppen, J.; Vaupel, J. W. Broken Limits to Life Expectancy. *Science* **2002**, *296* (5570), 1029–1031.

OMB. Definitions of Research and Development: An Annotated Compilation of Official Sources http://www.nsf.gov/statistics/randdef/fedgov.cfm (accessed May 12, 2015).

Open Science Collaboration, T. Estimating the Reproducibility of Psychological Science. *Science* **2015**, *349* (6251), aac4716–aac4716.

Orange, R. Medicine. Pioneering Womb Transplant Trial Highlights Risks and Ethical Dilemmas. *Science* **2014**, *343* (6178), 1418–1419.

Ord, T.; Hillerbrand, R.; Sandberg, A. Probing the Improbable: Methodological Challenges for Risks with Low Probabilities and High Stakes. *J. Risk Res.* **2010**, *13* (2), 191–205.

Oreskes, N.; Shrader-Frechette, K.; Belitz, K. Verification, Validation, and Confirmation of Numerical Models in the Earth Sciences. *Science* **1994**, *263* (5147), 641–646.

OSTP. White House Office of Science and Technology Policy. US Government Gain-of-Function Deliberative Process and Research Funding Pause on Selected Gain-of-Function Research Involving Influenza, MERS, and SARS Viruses. October 17, 2014. http://www.phe.gov/s3/dualuse/Documents/ (accessed Oct 21, 2014).

OTA. *Research Funding as an Investment: Can We Measure the Returns?*; Office of Technology Assessment: Washington, DC, 1986.

Otway, H. Experts, Risk Communication, and Democracy. *Risk Anal.* **1987**, *7* (2), 125–129.

Owens, B. Judgement Day. *Nature* **2013**, *502* (7471), 288–290.

Oye, K. A.; Esvelt, K.; Appleton, E.; Catteruccia, F.; Church, G.; Kuiken, T.; Lightfoot, S. B.-Y.; Mcnamara, J.; Smidler, A.; Collins, J. P. Regulating Gene Drives. *Science* **2014**, *345* (6197), 626–628.

Oye, K. A.; Lawson, J. C. H.; Bubela, T. Regulate "Home-Brew" Opiates. *Nature* **2015**, *521* (7552), 281–283.

Palese, P.; Wang, T. T. H5N1 Influenza Viruses: Facts, Not Fear. *Proc. Natl. Acad. Sci.* **2012**, *109* (7), 2211–2213.

Pandey, A.; Atkins, K. E.; Medlock, J.; Wenzel, N.; Townsend, J. P.; Childs, J. E.; Nyenswah, T. G.; Ndeffo-Mbah, M. L.; Galvani, A. P. Strategies for Containing Ebola in West Africa. *Science* **2014**, *346* (6212), 991–995.

Paté-Cornell, E. On "Black Swans" and "Perfect Storms": Risk Analysis and Management When Statistics Are Not Enough. *Risk Anal.* **2012**, *32* (11), 1823–1833.

Paté-Cornell, E.; Cox, L. A. Improving Risk Management: From Lame Excuses to Principled Practice. *Risk Anal.* **2014**, *34* (7), 1228–1239.

Patterson, A. P.; Tabak, L. A.; Fauci, A. S.; Collins, F. S.; Howard, S. A Framework for Decisions About Research with HPAI H5N1 Viruses. *Science* **2013**, *339* (6123), 1036–1037.

Pennisi, E. Building the Ultimate Yeast Genome. *Science* **2014**, *343* (6178), 1426–1429.

Pielke Jr, R. A. *The Honest Broker: Making Sense of Science in Policy and Politics*; Cambridge University Press, 2007.

Pilkey, O. H.; Pilkey-Jarvis, L. *Useless Arithmetic: Why Environmental Scientists Can't Predict the Future*; Columbia University Press: New York, 2007.

Pohjola, M. V.; Leino, O.; Kollanus, V.; Tuomisto, J. T.; Gunnlaugsdóttir, H.; Holm, F.; Kalogeras, N.; Luteijn, J. M.; Magnússon, S. H.; Odekerken, G.; et al. State of the Art in Benefit-Risk Analysis: Environmental Health. *Food Chem. Toxicol.* **2012**, *50* (1), 40–55.

Polka, J. K.; Krukenberg, K. A. Making Science a Desirable Career. *Science* **2014**, *346* (6215), 1422.

Pollack, R. Eugenics Lurk in the Shadow of CRISPR. *Science* **2015**, *348* (6237), 871.

Pontin, J. Why We Can't Solve Big Problems. *MIT Technol. Rev.* **2012**, *115* (6).

Popp, D. Induced Innovation and Energy Prices. *Am. Econ. Rev.* **2002**, *92* (1), 160–180.

Porter, J. E. Time to Speak up for Research. *Science* **2014**, *344* (6189), 1207.

Postman, N. *Technopoly: The Surrender of Culture to Technology*; Knopf: New York, 1992.

Poulin, B.; Lefebvre, G.; Paz, L. Red Flag for Green Spray: Adverse Trophic Effects of Bti on Breeding Birds. *J. Appl. Ecol.* **2010**, *47* (4), 884–889.

Pratt, J. W.; Raïffa, H.; Schlaifer, R. *Introduction to Statistical Decision Theory*; MIT Press: Cambridge, MA, 1995.

Press, W. H. What's So Special About Science (And How Much Should We Spend on It?). *Science* **2013**, *342* (6160), 817–822.

Prewitt, K. Is Any Science Safe? *Science* **2013**, *340* (6132), 525.

Price. [Editorial] The Price of Progress. *Nature* **2012**, *490* (7418), 5.

Quigley, J.; Revie, M. Estimating the Probability of Rare Events: Addressing Zero Failure Data. *Risk Anal.* **2011**, *31* (7), 1120–1132.

Quinn, S. C.; Kumar, S.; Freimuth, V. S.; Musa, D.; Casteneda-Angarita, N.; Kidwell, K. Racial Disparities in Exposure, Susceptibility, and Access to Health Care in the US H1N1 Influenza Pandemic. *Am. J. Public Health* **2011**, *101* (2), 285–293.

Rabesandratana, T. Europe's Rifts over Transgenic Crops Deepen at Key Deadline. *Science* **2015**, *350* (6256), 18–19.

Rabinovich, A.; Morton, T. A. Unquestioned Answers or Unanswered Questions: Beliefs about Science Guide Responses to Uncertainty in Climate Change Risk Communication. *Risk Anal.* **2012**, *32* (6), 992–1002.

Ran, F. A.; Cong, L.; Yan, W. X.; Scott, D. A.; Gootenberg, J. S.; Kriz, A. J.; Zetsche, B.; Shalem, O.; Wu, X.; Makarova, K. S.; et al. In Vivo Genome Editing Using Staphylococcus Aureus Cas9. *Nature* **2015**, *520* (7546), 186–191.

Rappert, B. Why Has Not There Been More Research of Concern? *Front. public Heal.* **2014**, *2* (74).

Rask, M. The Tragedy of Citizen Deliberation – Two Cases of Participatory Technology Assessment. *Technol. Anal. Strateg. Manag.* **2013**, *25* (1), 39–55.

De Rassenfosse, G.; van Pottelsberghe de la Potterie, B. A Policy Insight into the R&D–patent Relationship. *Res. Policy* **2009**, *38* (5), 779–792.

Rawls, J. *A Theory of Justice*; Harvard University Press, 1971.

Reardon, S. Lobbying Sways NIH Grants. *Nature* **2014a**, *515* (7525), 19.

Reardon, S. US Biodefence Facilities Ramp Up. *Nature* **2014b**, *509* (7498), 16–17.

Reardon, S. US Suspends Risky Disease Research. *Nature* **2014c**, *514* (7523), 411–412.

Reardon, S. "Forgotten" NIH Smallpox Virus Languishes on Death Row. *Nature* **2015**, *514* (7524), 544.

Reaven, S. J. The Concept of Net Energy I: The Problems of Net Energy Analysis. *Explor. Knowl.* **1984**, *1* (1), 191–231.

Reaven, S. J. The Concept of Net Energy II: Physical and Pragmatic Aspects of Net Energy. *Explor. Knowl.* **1985a**, *2* (1), 21–39.

Reaven, S. J. The Concept of Net Energy III: Does a Technology Have a Net Energy? *Explor. Knowl.* **1985b**, *2* (2), 25–45.

Reaven, S. J. Science Literacy Needs of Public Involvement Programs. In *Technological Literacy. Proceedings of the National Science, Technology and Society (STS) Conference*; Waks, L. J., Ed.; STS Press: Washington, DC, 1987a; pp 347–356.

Reaven, S. J. The Reliability of Individual and Pooled Expert Judgments. In *Uncertainty in Risk Assessment, Risk Management, and Decision Making*; Covello, V. T., Lave, L. B., Moghissi, A., Uppuluri, V. R. R., Eds.; Plenum Press: London, 1987b; pp 189–206.

Reaven, S. J. The Methodology of Probabilistic Risk Assessment: Completeness, Subjective Probabilities, and the "Lewis Report." *Explor. Knowl.* **1988**, *5* (1), 11–32.

Reaven, S. J. *How Sure Is Sure Enough: Scientific Controversies in Disposal of High-Level Radioactive Waste - Report to National Academy of Sciences Board on Radioactive Waste Management, Keystone Center, Colorado*; 1989a.

Reaven, S. J. The New Frontiers: Science and Technology at the World's Fair. In *Remembering the Future*; Bletter, Dickstein, Miller, Reaven, Eds.; Rizzoli: New York, 1989b; pp 74–103.

Rechard, R. P. Historical Relationship between Performance Assessment for Radioactive Waste Disposal and Other Types of Risk Assessment. *Risk Anal.* **1999**, *19* (5), 763–807.

Rees, M. Denial of Catastrophic Risks. *Science* **2013**, *339* (6124), 1123.

Refsgaard, J. C.; van der Sluijs, J. P.; Brown, J.; van der Keur, P. A Framework for Dealing with Uncertainty due to Model Structure Error. *Adv. Water Resour.* **2006**, *29* (11), 1586–1597.

Reich, E. S. Science Publishing: The Golden Club. *Nature* **2013**, *502* (7471), 291–293.

Rembold, C. M. The Health Benefits of Eating Salmon. *Science* **2004**, *305* (5683), 475; author reply 475.

Renn, O. White Paper on Risk Governance: Toward an Integrative Framework. In *Global Risk Governance: Concept and Practice Using the IRGC Framework*; Renn, O., Walker, K. D., Eds.; Springer Netherlands: Dordrecht, 2008; pp 3–73.

Renn, O.; Klinke, A. Systemic Risks: A New Challenge for Risk Management. *EMBO Rep.* **2004**, *5* (S1), S41–S46.

Renn, O.; Graham, P. *Risk Governance - Towards an Integrative Approach*; 1; International Risk Governance Council: Geneva, 2005.

Renn, O.; Benighaus, C. Perception of Technological Risk: Insights from Research and Lessons for Risk Communication and Management. *J. Risk Res.* **2013**, *16* (3-4), 293–313.

Resnik, D. B. Can Scientists Regulate the Publication of Dual Use Research? *Stud. Ethics. Law. Technol.* **2010**, *4* (1).

Resnik, D. B.; Barner, D. D.; Dinse, G. E. Dual-Use Review Policies of Biomedical Research Journals. *Biosecur. Bioterror.* **2011**, *9* (1), 49–54.

Revesz, R. L.; Howard, P. H.; Arrow, K.; Goulder, L. H.; Kopp, R. E.; Livermore, M. A.; Oppenheimer, M.; Sterner, T. Global Warming: Improve Economic Models of Climate Change. *Nature* **2014**, *508* (7495), 173–175.

Revill, J.; Jefferson, C. Tacit Knowledge and the Biological Weapons Regime. *Sci. Public Policy* **2013**, *41* (5), 597–610.

Rey, F.; Schwartz, O.; Wain-Hobson, S. Gain-of-Function Research: Unknown Risks. *Science* **2013**, *342* (6156), 311.

Ridley, M. *The Rational Optimist: How Prosperity Evolves*; HarperCollins: New York, 2010.

Rifkin, J. *The End of Work: The Decline of the Global Labor Force and the Dawn of the Post-Market Era.*; G.P. Putnam's Sons: New York, 1995.

Robberstad, B. QALY vs DALY vs LYs Gained: What Are the Difference, and What Difference Do They Make for Health Care Priority Setting? *Nor. Epidemiol.* **2005**, *15* (2), 183–191.

Roberts, R.; Relman, D. A. Letter to NSABB Chair, February 24, 2015 http://news.sciencemag.org/sites/default/files/Roberts Relman Letter, as sent to NSABB Chair, February 24, 2015, 14h30 U.S. Eastern.pdf.

Rockey, S. Announcement on Level I Activities https://www.starmetrics.nih.gov/Star/News (accessed May 15, 2015).

Roeser, S. Nuclear Energy, Risk, and Emotions. *Philos. Technol.* **2011**, *24* (2), 197–201.

Roman, H. A.; Hammitt, J. K.; Walsh, T. L.; Stieb, D. M. Expert Elicitation of the Value per Statistical Life in an Air Pollution Context. *Risk Anal.* **2012**, *32* (12), 2133–2151.

Rooney, P. On Values in Science: Is the Epistemic/Non-Epistemic Distinction Useful? In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*; University of Chicago Press: Chicago, 1992; pp 13–22.

Rosa, E. A. Metatheoretical Foundations for Post-Normal Risk. *J. Risk Res.* **1998**, *1* (1), 15–44.

Rosenthal, R. The “File Drawer Problem” and Tolerance for Null Results. *Psychol. Bull.* **1979**, *86* (3), 638–641.

Rossi, J. Participation Run Amok: The Costs of Mass Participation for Deliberative Agency Decisionmaking. *Northwest. Univ. Law Rev.* **1997**, *92* (1), 173–249.

Rotman, D. How Technology Is Destroying Jobs. *MIT Technol. Rev.* **2013**, *116* (4), 28–35.

Rozell, D. J. Protecting Privacy for Dual-Use Researchers. *Science* **2013**, *341* (6152), 1342.

Rozell, D. J. Fair Dynamic Pricing for Advanced Metering Infrastructure. *Strateg. Plan. Energy Environ.* **2014**, *34* (2), 26–38.

Rozell, D. J. A Cautionary Note on Qualitative Risk Ranking of Homeland Security Threats. *Homel. Secur. Aff.* **2015**, *11*, Article 3.

Rozell, D. J.; Reaven, S. J. Water Pollution Risk Associated with Natural Gas Extraction from the Marcellus Shale. *Risk Anal.* **2012**, *32* (8), 1382–1393.

Rozo, M.; Gronvall, G. K. The Reemergent 1977 H1N1 Strain and the Gain-of-Function Debate. *MBio* **2015**, *6* (4).

Rubinstein, A. Dilemmas of an Economic Theorist. *Econometrica* **2006**, *74* (4), 865–883.

Ruckelshaus, W. D. Risk in a Free Society. *Risk Anal.* **1984**, *4* (3), 157–162.

Rudner, R. The Scientist Qua Scientist Makes Value Judgments. *Philos. Sci.* **1953**, *20* (1), 1–6.

Rudski, J. M.; Osei, W.; Jacobson, A. R.; Lynch, C. R. Would You Rather Be Injured by Lightning or a Downed Power Line? Preference for Natural Hazards. *Judgm. Decis. Mak.* **2011**, *6* (4), 314–322.

Rumsfeld, D. H. Dept. of Defense News Briefing Feb. 12, 2002 http://www.defense.gov/transcripts/transcript.aspx?transcriptid=2636.

Russell, C. A.; Kasson, P. M.; Donis, R. O.; Riley, S.; Dunbar, J.; Rambaut, A.; Asher, J.; Burke, S.; Davis, C. T.; Garten, R. J.; et al. Improving Pandemic Influenza Risk Assessment. *Elife* **2014**, *3*, e03883.

Russell, S. Ethics of Artificial Intelligence: Take a Stand on AI Weapons. *Nature* **2015**, *521* (7553), 415–418.

Sagan, L. On the Origin of Mitosing Cells. *J. Theor. Biol.* **1967**, *14* (3), 225–IN6.

Sahlin, N. Unreliable Probabilities, Paradoxes, and Epistemic Risks. In *Handbook of Risk Theory*; Roeser, S., Hillerbrand, R., Sandin, P., Peterson, M., Eds.; Springer Netherlands: Dordrecht, 2012; pp 477–498.

Sale, K. *Rebels Against the Future*; Addison-Wesley Publishing: Reading, MA, 1995.

Salter, A. J.; Martin, B. R. The Economic Benefits of Publicly Funded Basic Research: A Critical Review. *Res. Policy* **2001**, *30* (3), 509–532.

Sandman, P. M. Science versus Spin: How Ron Fouchier and Other Scientists Miscommunicated about the Bioengineered Bird Flu Controversy http://www.psandman.com/articles/Fouchier.htm (accessed Feb 2, 2015).

Sarewitz, D. How Science Makes Environmental Controversies Worse. *Environ. Sci. Policy* **2004**, *7* (5), 385–403.

Sarewitz, D. Beware the Creeping Cracks of Bias. *Nature* **2012**, *485* (7397), 149.

Sarewitz, D. Science's Rightful Place Is in Service of Society. *Nature* **2013**, *502* (7473), 595–595.

Sarewitz, D. Reproducibility Will Not Cure What Ails Science. *Nature* **2015a**, *525* (7568), 159–159.

Sarewitz, D. Science Can't Solve It. *Nature* **2015b**, *522* (7557), 413–414.

Sarewitz, D.; Pielke Jr, R. A. The Neglected Heart of Science Policy: Reconciling Supply of and Demand for Science. *Environ. Sci. Policy* **2007**, *10* (1), 5–16.

Savulescu, J.; Gyngell, C.; Douglas, T. Germline Edits: Trust Ethics Review Process. *Nature* **2015**, *520* (7549), 623.

Scheiner, S. M.; Bouchie, L. M. The Predictive Power of NSF Reviewers and Panels. *Front. Ecol. Environ.* **2013**, *11* (8), 406–407.

Scheufele, D. A. Messages and Heuristics: How Audiences Form Attitudes about Emerging Technologies. In *Engaging science: Thoughts, deeds, analysis and action*; Turney, J., Ed.; The Wellcome Trust: London, 2006; pp 20–25.

Schoch-Spana, M. Public Engagement and the Governance of Gain-of-Function Research. *Heal. Secur.* **2015**, *13* (2), 1–5.

Schultz-Cherry, S.; Webby, R. J.; Webster, R. G.; Kelso, A.; Barr, I. G.; McCauley, J. W.; Daniels, R. S.; Wang, D.; Shu, Y.; Nobusawa, E.; et al. Influenza Gain-of-Function Experiments: Their Role in Vaccine Virus Recommendation and Pandemic Preparedness. *MBio* **2014**, *5* (6), e02430–14.

Scotchmer, S. *Innovation And Incentives*; MIT Press: Cambridge, MA, 2004.

Seidl, R.; Moser, C.; Stauffacher, M.; Krütli, P. Perceived Risk and Benefit of Nuclear Waste Repositories: Four Opinion Clusters. *Risk Anal.* **2013**, *33* (6), 1038–1048.

Selgelid, M. J. Governance of Dual-Use Research: An Ethical Dilemma. *Bull. World Health Organ.* **2009**, *87* (9), 720–723.

Sen, A. *Poverty and Famines: An Essay on Entitlement and Deprivation*; Oxford: Clarendon Press, 1981.

Servick, K. Unsettled Questions Trail IVF's Success. *Science* **2014**, *345* (6198), 744–746.

Shafer, G. *A Mathematical Theory of Evidence*; Princeton University Press: Princeton, NJ, 1976.

Shafer, G. Perspectives on the Theory and Practice of Belief Functions. *Int. J. Approx. Reason.* **1990**, *4* (5-6), 323–362.

Sharples, F.; Husbands, J.; Mazza, A.-M.; Thevenon, A.; Hook-Barnard, I. *Potential Risks and Benefits of Gain-of-Function Research: Summary of a Workshop*; National Academies Press: Washington, DC, 2015.

Shechtman, D.; Blech, I.; Gratias, D.; Cahn, J. W. Metallic Phase with Long-Range Orientational Order and No Translational Symmetry. *Phys. Rev. Lett.* **1984**, *53* (20), 1951–1953.

Sheil, M. On the Verge of a New ERA. *Nature* **2014**, *511* (7510), S67.

Shen, H. Grants: Funder Storm. *Nature* **2014**, *508* (7497), 555–556.

Shrader-Frechette, K. S. The Conceptual Risks of Risk Assessment. *IEEE Technol. Soc. Mag.* **1986**, *5* (2), 4–11.

Shrader-Frechette, K. S. *Risk and Rationality: Philosophical Foundations for Populist Reforms*; University of California Press: Berkeley, 1991.

Siebert, S.; Machesky, L. M.; Insall, R. H. Overflow in Science and Its Implications for Trust. *Elife* **2015**, *4*, e10825.

Siler, K.; Lee, K.; Bero, L. Measuring the Effectiveness of Scientific Gatekeeping. *Proc. Natl. Acad. Sci.* **2014**, *112* (2), 201418218.

Singer, P. W.; Friedman, A. *Cybersecurity and Cyberwar*; Oxford University Press: Oxford, 2014.

Singer, P. Moral Experts. *Analysis* **1972**, *32* (4), 115–117.

Singh, D. Publication Bias - a Reason for the Decreased Research Output in Developing Countries. *Afr. J. Psychiatry* **2006**, *9* (3), 153–155.

Sjöberg, L. Factors in Risk Perception. *Risk Anal.* **2000**, *20* (1), 1–11.

Sjöberg, L. Attitudes toward Technology and Risk: Going beyond What Is Immediately given. *Policy Sci.* **2002**, *35* (4), 379–400.

Slakey, F.; Cohen, L. R. Stop Laser Uranium Enrichment. *Nature* **2010**, *464* (7285), 32–33.

Slovic, P. Perception of Risk. *Science* **1987**, *236* (4799), 280–285.

Slovic, P. *The Perception of Risk*; Earthscan: London, 2000.

Slovic, P.; Finucane, M. L.; Peters, E.; MacGregor, D. G. Risk as Analysis and Risk as Feelings: Some Thoughts about Affect, Reason, Risk, and Rationality. *Risk Anal.* **2004**, *24* (2), 311–322.

Slovic, P.; Finucane, M. L.; Peters, E.; MacGregor, D. G. The Affect Heuristic. *Eur. J. Oper. Res.* **2007**, *177* (3), 1333–1352.

Small, M. J.; Güvenç, Ü.; DeKay, M. L. When Can Scientific Studies Promote Consensus among Conflicting Stakeholders? *Risk Anal.* **2014**, *34* (11), 1978–1994.

Smith, J.; Gangadharan, D.; Weyant, R. Review of Restricted Experiment Requests, Division of Select Agents and Toxins, Centers for Disease Control and Prevention, 2006-2013. *Heal. Secur.* **2015**, *13* (5), 307–316.

Smith, M. R.; Marx, L. *Does Technology Drive History?: The Dilemma of Technological Determinism*; MIT Press: Cambridge, MA, 1994.

Solow, A. R. Power Laws without Complexity. *Ecol. Lett.* **2005**, *8* (4), 361–363.

Solow, R. M. Technical Change and the Aggregate Production Function. *Rev. Econ. Stat.* **1957**, *39* (3), 312–320.

Song, G. Understanding Public Perceptions of Benefits and Risks of Childhood Vaccinations in the United States. *Risk Anal.* **2014**, *34* (3), 541–555.

Specter, M. A Life of Its Own: Where Will Synthetic Biology Lead Us? *The New Yorker*. New York 2009, pp 56–65.

Srinivasan, R.; Natarajan, S. Developments in Inherent Safety: A Review of the Progress during 2001–2011 and Opportunities Ahead. *Process Saf. Environ. Prot.* **2012**, *90* (5), 389–403.

Stein, R. Disgraced Scientist Clones Dogs, And Critics Question His Intent http://www.npr.org/sections/health-shots/2015/09/30/418642018/disgraced-scientist-clones-dogs-and-critics-question-his-intent (accessed Oct 9, 2015).

Stern, N. *Stern Review on the Economics of Climate Change*; Stern, N. N. ., Ed.; Cambridge University Press: Cambridge, UK, 2006.

Stern, N.; Taylor, C. Climate Change: Risk, Ethics, and the Stern Review. *Science* **2007**, *317* (5835), 203–204.

Stevens, A. J.; Jensen, J. J.; Wyller, K.; Kilgore, P. C.; Chatterjee, S.; Rohrbaugh, M. L. The Role of Public-Sector Research in the Discovery of Drugs and Vaccines. *N. Engl. J. Med.* **2011**, *364*, 535–541.

Stöhr, K. Perspective: Ill Prepared for a Pandemic. *Nature* **2014**, *507* (7490), S20–S21.

Stokey, E.; Zeckhauser, R. J. *A Primer for Policy Analysis*; Norton: New York, 1978.

Stone, D. *Policy Paradox: The Art of Political Decision Making*, Revised Ed.; Norton: New York, 2002.

Stone, R. Arms Control. Scientists Campaign against Killer Robots. *Science* **2013**, *342* (6165), 1428–1429.

Suk, J. E.; Zmorzynska, A.; Hunger, I.; Biederbick, W.; Sasse, J.; Maidhof, H.; Semenza, J. C. Dual-Use Research and Technological Diffusion: Reconsidering the Bioterrorism Threat Spectrum. *PLoS Pathog.* **2011**, *7* (1), e1001253.

Sumathipala, A.; Siribaddana, S.; Patel, V. Under-Representation of Developing Countries in the Research Literature: Ethical Issues Arising from a Survey of Five Leading Medical Journals. *BMC Med. Ethics* **2004**, *5* (1), E5.

Sunstein, C. R. *Laws of Fear: Beyond the Precautionary Principle*; Cambridge University Press: Cambridge, 2005.

Sutherland, W. J.; Burgman, M. Policy Advice: Use Experts Wisely. *Nature* **2015**, *526* (7573), 317–318.

Sutherland, W. J.; Spiegelhalter, D.; Burgman, M. Twenty Tips for Interpreting Scientific Claims. *Nature* **2013**, *503* (7476), 335–337.

Sutton, T. C.; Finch, C.; Shao, H.; Angel, M.; Chen, H.; Capua, I.; Cattoli, G.; Monne, I.; Perez, D. R. Airborne Transmission of Highly Pathogenic H7N1 Influenza Virus in Ferrets. *J. Virol.* **2014**, *88* (12), 6623–6635.

Taleb, N. N.; Goldstein, D. G.; Spitznagel, M. W. The Six Mistakes Executives Make in Risk Management. *Harv. Bus. Rev.* **2009**, *87* (10), 78–81.

Tannert, C.; Elvers, H.-D.; Jandrig, B. The Ethics of Uncertainty. *EMBO Rep.* **2007**, *8* (10), 892–896.

Tavernise, S.; McNeil Jr, D. G. Lapse Sets Off Anthrax Scare for Scientists at C.D.C. Labs. *New York Times*. New York June 20, 2014, p A11.

Taylor, P. R. The Mismeasure of Risk. In *Handbook of Risk Theory*; Roeser, S., Hillerbrand, R., Sandin, P., Peterson, M., Eds.; Springer Netherlands: Dordrecht, 2012; pp 441–475.

Teitelbaum, M. S. *Falling Behind? Boom, Bust and the Global Race for Scientific Talent*; Princeton University Press: Princeton, NJ, 2014.

Tetlock, P. E. Thinking the Unthinkable: Sacred Values and Taboo Cognitions. *Trends Cogn. Sci.* **2003**, *7* (7), 320–324.

Thompson, P. B. Risking or Being Willing: Hamlet and the DC-10. *J. Value Inq.* **1985**, *19* (4), 301–310.

Thyer, R.; Ellefson, J. New Letters for Life's Alphabet. *Nature* **2014**, *509* (7500), 291–292.

Tierney, J. Doomsayers Beware, A Bright Future Beckons. *New York Times*. New York May 18, 2010, p D1.

Tonn, B.; Stiefel, D. Human Extinction Risk and Uncertainty: Assessing Conditions for Action. *Futures* **2014**, *63*, 134–144.

Toulmin, S. The Complexity of Scientific Choice: A Stocktaking. *Minerva* **1964**, *2* (3), 343–359.

Trevan, T. Do Not Censor Science in the Name of Biosecurity. *Nature* **2012**, *486* (7403), 295.

Trock, S. C.; Burke, S. A.; Cox, N. J. Development of an Influenza Virologic Risk Assessment Tool. *Avian Dis.* **2012**, *56* (4s1), 1058–1061.

Tucker, J. B. *Innovation, Dual Use, and Security: Managing the Risks of Emerging Biological and Chemical Technologies*; MIT Press: Cambridge, MA, 2012.

Tucker, W. T.; Ferson, S. Strategies for Risk Communication: Evolution, Evidence, Experience. *Ann. N. Y. Acad. Sci.* **2008**, *1128* (631), ix – xii.

Tumpey, T. M.; Basler, C. F.; Aguilar, P. V; Zeng, H.; Solórzano, A.; Swayne, D. E.; Cox, N. J.; Katz, J. M.; Taubenberger, J. K.; Palese, P.; et al. Characterization of the Reconstructed 1918 Spanish Influenza Pandemic Virus. *Science* **2005**, *310* (5745), 77–80.

Tuomisto, J. T.; Tuomisto, J.; Tainio, M.; Niittynen, M.; Verkasalo, P.; Vartiainen, T.; Kiviranta, H.; Pekkanen, J. Risk-Benefit Analysis of Eating Farmed Salmon. *Science* **2004**, *305* (5683), 476–477; author reply 476–477.

Turner, E. H.; Matthews, A. M.; Linardatos, E.; Tell, R. A.; Rosenthal, R. Selective Publication of Antidepressant Trials and Its Influence on Apparent Efficacy. *N. Engl. J. Med.* **2008**, *358* (3), 252–260.

Tversky, A.; Kahneman, D. Availability: A Heuristic for Judging Frequency and Probability. *Cogn. Psychol.* **1973**, *5* (2), 207–232.

Tversky, A.; Kahneman, D. Judgment under Uncertainty: Heuristics and Biases. *Science* **1974**, *185* (4157), 1124–1131.

Uhlenhaut, C.; Burger, R.; Schaade, L. Protecting Society. *EMBO Rep.* **2013**, *14* (1), 25–30.

Unwin, S. D. A Fuzzy Set Theoretic Foundation for Vagueness in Uncertainty Analysis. *Risk Anal.* **1986**, *6* (1), 27–34.

Usher, S. *Letters of Note: Correspondence Deserving of a Wider Audience*; Cannongate: Edinburgh, 2013.

Vesely, W. E.; Goldberg, F. F.; Roberts, N. H.; Haasl, D. F. *Fault Tree Handbook*; Nuclear Regulatory Commission: Washington, DC, 1981.

Viscusi, W. K.; Hakes, J. K. Synthetic Risks, Risk Potency, and Carcinogen Regulation. *J. policy Anal. Manag.* **1998**, *17* (1), 52–73.

Visser, H.; Folkert, R. J. M.; Hoekstra, J.; Wolff, J. J. de. Identifying Key Sources of Uncertainty in Climate Change Projections. *Clim. Change* **2000**, *45* (3-4), 421–457.

Visser't Hooft, H. P. *Justice to Future Generations and the Environment*; Kluwer Academic Publishers: Dordrecht, 1999.

Vogel, G. Embryo Engineering Alarm. *Science* **2015**, *347* (6228), 1301.

Vogel, K. M. Intelligent Assessment: Putting Emerging Biotechnology Threats in Context. *Bull. At. Sci.* **2013**, *69* (1), 43–52.

Vrugt, J. A.; Diks, C. G. H.; Gupta, H. V.; Bouten, W.; Verstraten, J. M. Improved Treatment of Uncertainty in Hydrologic Modeling: Combining the Strengths of Global Optimization and Data Assimilation. *Water Resour. Res.* **2005**, *41* (1), n/a – n/a.

Wachs, M. When Planners Lie with Numbers. *J. Am. Plan. Assoc.* **1989**, *55*, 476–479.

Wade, N. Scientists Seek Ban on Method of Editing the Human Genome. *New York Times*. New York March 20, 2015, p A1.

Wadhwa, V.; Freeman, R.; Rissing, B. *Education and Tech Entrepreneurship*; Ewing Marion Kauffman Foundation: Kansas City, MO, 2008.

Wagner, C. S. Innovation Goes Global. *Science* **2014**, *343* (6172), 730.

Wain-Hobson, S. An Avian H7N1 Gain-of-Function Experiment of Great Concern. *MBio* **2014**, *5* (5), e01882–14.

Walker, W. E.; Harremoes, P.; Rotmans, J.; van der Sluijs, J. P.; van Asselt, M. B. A.; Janssen, P.; Krayer von Krauss, M. P. Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integr. Assess.* **2003**, *4* (1), 5–17.

Walley, P. *Statistical Reasoning with Imprecise Probabilities*; Chapman and Hall: London, 1991.

Wang, D.; Song, C.; Barabási, A.-L. Quantifying Long-Term Scientific Impact. *Science* **2013**, *342* (6154), 127–132.

Wang, D.; Song, C.; Shen, H.-W.; Barabási, A.-L. Science Communication. Response to Comment on "Quantifying Long-Term Scientific Impact". *Science* **2014a**, *345* (6193), 149.

Wang, J.; Mei, Y.; Hicks, D. Science Communication. Comment on "Quantifying Long-Term Scientific Impact". *Science* **2014b**, *345* (6193), 149.

Wang, Z. *In Sputnik's Shadow: The President's Science Advisory Committee and Cold War America*; Rutgers University Press: Piscataway, NJ, 2008; Vol. 40.

Watanabe, T.; Zhong, G.; Russell, C. A.; Nakajima, N.; Hatta, M.; Hanson, A.; McBride, R.; Burke, D. F.; Takahashi, K.; Fukuyama, S.; et al. Circulating Avian Influenza Viruses Closely Related to the 1918 Virus Have Pandemic Potential. *Cell Host Microbe* **2014**, *15* (6), 692–705.

Watson, T. Private Asteroid Hunt Lacks Cash to Spy Threats in Orbit. *Nature* **2015**, *522* (7557), 402–403.

Watts, S. Society Needs More than Wonder to Respect Science. *Nature* **2014**, *508* (7495), 151.

Weaver, D. E. Contaminant Levels in Farmed Salmon. *Science* **2004**, *305* (5683), 478; author reply 478.

Webster, R. G.; Bean, W. J.; Gorman, O. T.; Chambers, T. M.; Kawaoka, Y. Evolution and Ecology of Influenza A Viruses. *Microbiol. Rev.* **1992**, *56* (1), 152–179.

Weichselberger, K. The Theory of Interval-Probability as a Unifying Concept for Uncertainty. *Int. J. Approx. Reason.* **2000**, *24* (2-3), 149–170.

Weiman, C. The “Curse of Knowledge” or Why Intuition about Teaching Often Fails. *Am. Phys. Soc. News* **2007**, *16* (10), 8.

Wein, L. M.; Liu, Y. Analyzing a Bioterror Attack on the Food Supply: The Case of Botulinum Toxin in Milk. *Proc. Natl. Acad. Sci.* **2005**, *102* (28), 9984–9989.

Weinberg, A. M. Science and Trans-Science. *Minerva* **1972**, *10* (2), 209–222.

Weinberg, A. M. Reflections on Risk Assessment. *Risk Anal.* **1981**, *1* (1), 5–7.

Weinberg, B. A.; Owen-Smith, J.; Rosen, R. F.; Schwarz, L.; Allen, B. M.; Weiss, R. E.; Lane, J. Science Funding and Short-Term Economic Activity. *Science* **2014**, *344* (6179), 41–43.

Weinberger, S. Laser Plant Offers Cheap Way to Make Nuclear Fuel. *Nature* **2012**, *487* (7405), 16–17.

Wekerle, H. Vaccination and Narcolepsy: Immune Link Found? *Sci. Transl. Med.* **2015**, *7* (294), 294fs27.

Werner, R. The Focus on Bibliometrics Makes Papers Less Useful. *Nature* **2015**, *517* (7534), 245.

WHO. *Report on Technical Consultation on H5N1 Research Issues*; World Health Organization: Geneva, 2012.

WHO. *Health Risk Assessment from the Nuclear Accident after the 2011 Great East Japan Earthquake and Tsunami, Based on a Preliminary Dose Estimation*; World Health Organization: Geneva, 2013.

WHO. Ebola Virus Disease in West Africa — The First 9 Months of the Epidemic and Forward Projections. *N. Engl. J. Med.* **2014**, *371* (16), 1481–1495.

Williams, P. R.; Hammitt, J. K. Perceived Risks of Conventional and Organic Produce: Pesticides, Pathogens, and Natural Toxins. *Risk Anal.* **2001**, *21* (2), 319–330.

Wilson, R.; Crouch, E. A. C. *Risk-Benefit Analysis*, 2nd ed.; Harvard Univsersity Press: Cambridge, MA, 2001.

Windbichler, N.; Papathanos, P. A.; Catteruccia, F.; Ranson, H.; Burt, A.; Crisanti, A. Homing Endonuclease Mediated Gene Targeting in Anopheles Gambiae Cells and Embryos. *Nucleic Acids Res.* **2007**, *35* (17), 5922–5933.

Winkler, R. L. Uncertainty in Probabilistic Risk Assessment. *Reliab. Eng. Syst. Saf.* **1996**, *54* (2-3), 127–132.

Winkler, R. L. Equal versus Differential Weighting in Combining Forecasts. *Risk Anal.* **2015**, *35* (1), 16–18.

Winkler, R. L.; Smith, J. E.; Fryback, D. G. The Role of Informative Priors in Zero-Numerator Problems: Being Conservative versus Being Candid. *Am. Stat.* **2002**, *56* (1), 1–4.

Von Winterfeldt, D.; Edwards, W. Defining a Decision Analytic Structure. In *Advances in Decision Analysis*; Edwards, W., Miles, R. F., von Wintereldt, D., Eds.; Cambridge University Press: New York, 2007; pp 489–513.

Wolinetz, C. D. Implementing the New U.S. Dual-Use Policy. *Science* **2012**, *336* (6088), 1525–1527.

Woyke, T.; Rubin, E. M. Searching for New Branches on the Tree of Life. *Science* **2014**, *346* (6210), 698–699.

Wray, K. B. Older Scientists Get Their Due. *Science* **2014**, *346* (6212), 929.

Wright, S. Terrorists and Biological Weapons. Forging the Linkage in the Clinton Administration. *Polit. Lfe Sci.* **2006**, *25* (1-2), 57–115.

Wyatt, E. U.S. Sets 21st-Century Goal: Building a Better Patent Office. *New York Times*. New York February 21, 2011, p A1.

Xie, X.-F.; Wang, M.; Zhang, R.-G.; Li, J.; Yu, Q.-Y. The Role of Emotions in Risk Communication. *Risk Anal.* **2011**, *31* (3), 450–465.

Yamamoto, Y. T. Values, Objectivity and Credibility of Scientists in a Contentious Natural Resource Debate. *Public Underst. Sci.* **2012**, *21* (1), 101–125.

Yarborough, M. Openness in Science Is Key to Keeping Public Trust. *Nature* **2014**, *515* (7527), 313.

Yokota, F.; Thompson, K. M. Value of Information Analysis in Environmental Health Risk Management Decisions: Past, Present, and Future. *Risk Anal.* **2004**, *24* (3), 635–650.

Young, A.; Penzenstadler, N. Inside America's secretive biolabs http://www.usatoday.com/story/news/2015/05/28/biolabs-pathogens-location-incidents/26587505/ (accessed Mar 6, 2015).

Young, N. S.; Ioannidis, J. P. A.; Al-Ubaydli, O. Why Current Publication Practices May Distort Science. *PLoS Med.* **2008**, *5* (10), e201.

Zadeh, L. A. Fuzzy Sets. *Inf. Control* **1965**, *8* (3), 338–353.

Zandvoort, H. Requirements for the Social Acceptability of Risk-Generating Technological Acitivities. In *The Ethics of Technological Risk*; Asveld, L., Roeser, S., Eds.; Earthscan: London, 2009; pp 40–54.

Zetsche, B.; Gootenberg, J. S.; Abudayyeh, O. O.; Slaymaker, I. M.; Makarova, K. S.; Essletzbichler, P.; Volz, S. E.; Joung, J.; van der Oost, J.; Regev, A.; et al. Cpf1 Is a Single RNA-Guided Endonuclease of a Class 2 CRISPR-Cas System. *Cell* **2015**.

Zimmerman, J. B.; Anastas, P. T. Toward Designing Safer Chemicals. *Science* **2015**, *347* (6219), 215.