From CBA to precautionary appraisal: practical responses to intractable problems

Article (Accepted Version)

From CBA to Precautionary Appraisal: practical responses to intractable problems

By Andy Stirling and Josie Coburn

Synthetic biology is the latest wave of technological exuberance in which '*pro-innovation policies'*[1] are held to yield '*science-based solutions.'*[2] Synbio is often said to be so promising that the regulatory issues for its many emerging applications are confined to *how*, not *whether*, to proceed. To contemplate deeper or more general questioning of these technologies is frequently treated as irrational, unscientific, or anti-technology.

But is this a good default basis for governance of emerging technologies in complex, dynamic, momentous contexts?[3] On either side, the stakes in answering this question are high.[4] There is little doubt that significant potential benefits are possible, at least for some.[5] But possible direct and indirect harms are also grave.[6] As a result, the pro-technology euphoria provokes equally intense counter-reactions, and all this polarized rhetoric presents further challenges that need to be addressed in governance of these technologies.[7]

The purpose of this essay is to critically review the design of methods for ethically robust forms of technology appraisal in the regulation of research and innovation in synthetic biology. It will focus in particular on the extent to which cost-benefit analysis (CBA) offers a basis for informing decisions about which technological pathways to pursue and which to discourage. A further goal is to consider what (if anything) 'the precautionary principle' might offer in enabling better decisions. And this in turn raises questions about why mention of precaution can excite accusations of unscientific bias or irrational, "*anti-innovation*" extremism[8]. What does the polarized debate tell us about the politics around synthetic biology? In seeking more rigorous, timely, and practical ways to govern these remarkable new technologies, what might we be missing?

The sophistication, diversity, and scope of synthetic biology may seem to make these technologies a rather idiosyncratic setting for exploring these general issues. It may seem to be a special case, with the bewildering pace of change amplifying the difficulties. But at root, some of the trickiest issues are just specific instances of familiar and long-standing conundrums in the governance of science and technology. The basic challenge is how to weigh up, for a wide range of potential options, the various pros and cons, as viewed from divergent perspectives, and find a way to justify the best course of action on behalf of society as a whole. This is the central problem addressed by a number of techniques in CBA.[9] On the face of it, synbio seems to present just one more application of these well-established and self-confident prescriptive methods.

But there do emerge several obstinate, even prohibitive, difficulties for CBA. How to deal with the intractability of real-world uncertainties, of kinds that defy the probabilistic reductions of CBA? How to address forms of ethical reasoning that go beyond the simple scalar trade-offs that lie at the core of CBA? When consequences are difficult to characterize in any one robust way, how rational is it to insist on doing so anyway? What about the ambiguities and irreconcilabilities in the divergent values and interests around synthetic biology, difficulties that make it misleading to produce any single, notionally definitive picture? And how to balance the often-invisible effects of power – operating as much within research and innovation systems (and appraisal itself) as outside?

None of these challenges are resolved – or often even properly addressed – by CBA. Indeed, although they are well acknowledged by the scholarly literature on and around this topic, they are often side-lined in practice.[10] Yet all are central to the case for applying the concept of precaution to a field like synthetic biology. This essay will briefly explore multicriteria mapping as one practical way to address them.[11] The essay focuses on MCM not because it presents any sort of panacea for appraisal, but because it is illustrative of the concrete implications of precaution. Setting out even just one among potentially many practical alternative methods at least refutes the last-ditch argument that CBA is the only operational choice.

Thorny Problems for Technology Appraisal

Before turning to the difficulties in CBA, it is important to acknowledge that some of the deepest and most important challenges in governance of synthetic biology apply no matter what assessment method is employed – be it risk analysis, cost-benefit valuation, decision theory, technology assessment, participatory engagement, various kinds of ethical deliberation, or multicriteria mapping. The challenges are institutionally hard-wired dilemmas concerning the nature and scope of governance attention in general. For instance: how to select and define the options for social, institutional, or technological innovation that the chosen method should consider in the first place? Should regulatory processes address only the proposed synthetic biology application itself, or should attention extend to other policies, practices, or products for addressing the same societal ends? Such fundamental questions in the framing of technology governance[12] are typically settled implicitly and at the earliest stages of regulation, deep inside the institutions conducting the appraisal and (for the most part) outside the canons of any particular appraisal method.

The scope of governance is a question that is strikingly neglected in much technology governance. Regulatory attention is typically limited to appraisal of a particular technology or application, rather than widening to include a comparative appraisal of alternatives.[13] The fact is, however, that new technologies tend to enjoy support because of their benefits for powerful sponsoring academic, governmental, or corporate interests. These power dynamics are not necessarily negative, but it is not clear why they should be treated as satisfactory determinants of societal goods and harms.[14] To do this (even if only implicitly) is to perpetuate a serious asymmetry in how appraisal is conducted. Yet this is typically exactly the effect of regulatory assessment when it is framed as addressing questions merely about whether the new technology that happens to be favored by the most powerful interests is '*Safe?*' '*Safe enough?*' '*Acceptable?*' …or even just '*tolerable*'?[15]

The problem with this restricted approach to technology governance, of course, is that innovation in any given sector is not rationally understood as a one-track 'race to the future',[16] in which the next advance in any given field will be singular and self-evident. If it were, then regulation could more reasonably reduce to circumscribed questions about how fast or safely to proceed. But the evidence from every discipline that has examined innovation is that scientific and technological advance is not linear and determined, but rather involves emergent and contingent branching evolutionary patterns.[17] At every stage, there are multiple, continually bifurcating, alternative pathways.[18] The main issues in regulating emerging technologies, then, are not just about *how safe* or *how fast* the chosen technological trajectories should be.[19] In order to be robust, governance should involve far more difficult questions about the choices themselves – concerning '*which way?*', '*who says?*' and '*why?*'

Why is the purpose of regulatory appraisal so often – and strongly – framed in unreasonably circumscribed ways? Part of the answer lies in how feedback processes channel technological change in particular directions rather than others.[20] Innovations that are favored by powerful interests in research or innovation systems can be even more strongly favored by processes of lock-in, entrenchment, and entrapment.[21] Even just neutral historical path dependencies and network externalities can help inferior technologies rise to market ascendancy.[22] This is how processes of technology choice that are prospectively radically open and flexible, can nonetheless look closed and determined in retrospect. The constraining of scope in regulatory appraisal is merely one among many manifestations of these pressures.[23]

But the practical implications for innovation governance go even further. Even if some alternatives are widely agreed to offer superior benefits, circumscribed styles of technology regulation may end up crowding them out. It is therefore insufficient to address qualities like 'responsibility' or 'ethics' in innovation governance merely as a matter of selecting an appraisal method. Whatever method is chosen, arguably the most fundamental issues lie in the questions that it is charged with answering. To restrict regulatory attention to questions about '*how fast?*' risks reducing innovation governance to an elaborate legitimation of whatever technology happens to be most powerfully backed.[24]

To remedy this, technology governance in general – and regulatory appraisal in particular – should recognize that innovation in all the fields to which synthetic biology might be applied is as much about nurturing and exploring a space of different possibilities as about making any given trajectory *safe* or *tolerable*.[25] This means thinking about more than the costs and benefits of a given technology and taking a full and balanced view of the comparative pros and cons of a wide array of alternative possible innovation pathways, including social innovations and completely different technologies in each given setting.[26] And it implies no particular skepticism about synbio to note that it is simply one among a number of potentially positive innovations in the fields in which it might be applied. Whatever the benefits of diversity, not all can be equally pursued – let alone acquire momentum – together.

For instance, synbio is one many other potential options for addressing challenges of sustainable food provision. Even within the biosciences, transgenics and cisgenics open up quite radically different political-economic implications. Marker-assisted selection and advanced conventional plant breeding are also often effective. And beyond the biosciences, a host of practices like open source production, support for farmer seed selection, and various kinds of ecological agriculture may, in particular settings, often offer more effective responses to particular challenges. Each of these innovation pathways may display radically contrasting political-economic, environmental, and safety characteristics [27].

A central dilemma in the governance of synthetic biology is thus to ensure balanced attention. There is no guarantee that the innovation trajectories offering greatest benefits for private actors in commercial markets will be the same pathways that offer the best mix of pros and cons for society at large. Factors that drive successful pathways tend to involve considerations like prestige in elite scientific disciplines; rent on intellectual property; control of commercial value chains; realizing economies of scope and scale; and consolidating market share. Such incentives offer little to the poorest people, who are usually marginalized in decision-making, yet it is trajectories favored in precisely these ways that conventional framing of regulatory appraisal tends to entrench. And the language of objective rationality and 'sound scientific' decision-making so often used around CBA can compound the entrenchment. When particular incumbent trajectories are challenged by ‘precautionary’ criticism that favors some alternative innovation pathway, it is by reference to ostensibly rational techniques like CBA that advocates of precaution are labelled ‘*anti-science*’ or ‘*anti-innovation*’.[28]

So governance of technology must resist pressures to narrow its own focus. It must ask questions about whether synthetic biology is necessary or justified and which of a range of (possibly radically) alternative technology or policy options might be preferable. For this to be achieved, what is needed are appraisal methods that deliberately broaden out attention to a wider variety of options, issues, uncertainties, and perspectives than are convenient to incumbent interests.[29] These can open up social choice in a field like this in order to facilitate wider, more democratic, attention than is otherwise conditioned by the ‘body language’ of reductive analysis and technocratic expertise.[30] If CBA can do this, then it will form part of the solution. If it cannot, then it will be part of the problem.

Lessons of Precaution for CBA

The first point to note is just that, even though incumbency and power imbalances are inevitable and not necessarily all bad, they must be effectively discussed. At root, it is the danger of airbrushing power from regulatory appraisal methods like risk assessment and CBA, that the heuristics and practices of precaution address most directly. Precaution, in all its forms, is a cue to attend to these issues of power and bias in regulation of emerging technologies. Precaution takes seriously that uncertainties cannot satisfactorily be reduced to probabilistic risk or expected values in CBA, as power typically prefers. Rather than pretending at ostensibly comprehensive and objective analysis (as CBA is often presented), precaution enjoins that regulatory appraisal be recognized as inherently normative and that it includes key roles for social values. Science is of course absolutely necessary, but it is never sufficient. Values and subjective framings are always part of the analysis.

A practical illustration of how this point bears on methods like CBA can be found in the energy sector, which is arguably the most influential and mature field for application of these techniques. Applications of CBA-related methods have perhaps attained their most sophisticated forms in energy, and the results have often been adopted in regulatory law. The sixty-three peer-reviewed CBA studies in Figure 1, for instance, each express their results precisely, with rather little attention to uncertainty. Yet, when we look at them alongside each other, then remarkable uncertainties and ambiguities are apparent. The ranges across the analyses are so large that it is not just the absolute magnitudes of the ‘costs’ and ‘benefits’ that are at issue; it turns out that CBA can justify radically different ordinal rankings across whatever might be included as the contending alternatives for appraisal.[31] In other words, while evidence-based policy appraisal methods like CBA can provide valuable insights, they cannot generate a single definitive ‘sound scientific’ policy choice.

Figure 1. Concealed variability in a well-established field regulatory CBA

It is easy to explain the general causes of this variability. No matter how finely the methodological protocols are specified, their framing assumptions differ in apparently minor ways, but with potentially major effects on their recommendations. Figure 2 shows a variety of dimensions in which contrasting positions may be taken in the implementation of any method (like CBA) concerning the framing of options, issues, contexts, and uncertainties. Slight variations in any of these will typically yield significantly contrasting pictures of the relative performance of different alternatives.

Of course, standardized methodological conventions could be adopted, such that different studies might seek to apply the same framing assumptions. But there would always remain the question, whether any given ordering is simply an artefact of particular contingent decisions about standardization. The problem remains, that standardization can be based equally reasonably around different sets of framing assumptions, which hold contrasting implications for the ordering of policy options. Of course, these difficulties are not unique to CBA. Albeit differing in their details, the very general nature of the many kinds of framing assumptions mean that similar challenges apply equally across all quantitative, qualitative, and hybrid methods in regulatory appraisal. But the more assertive the presentation of prescriptive results, the more serious the resulting problems.

The appropriate role for regulatory appraisal methods like CBA, then, is not to pretend at deriving a single definitive 'science-based' picture of contrasting technology and policy options. Instead, the value of assessment lies in the clarity and rigor with which it can show which specific assumptions and perspectives lead to which conclusions. What precaution calls for in the governance of emerging technology, then, is use of methods that resist the technocratic approach to regulatory appraisal and avoid attempting to claim a singular definitive output. Focusing on the implications of various kinds of uncertainty, precaution urges greater transparency and conditionality – and associated deliberation and accountability – in the justification of why one technological pathway should be preferred to another.[32]

Figure 2. Framing parameters under which diverging assumptions can radically affect results in regulatory appraisal

Figure 3 presents a field of different possible methods for undertaking more precautionary appraisal in innovation governance. The two axes of the diagram show the varying degrees of confidence that may exist concerning two fundamental dimensions of knowledge. These twin parameters of 'likelihood' and 'outcome' constitute the logical structure of expected utility calculus, which in turn forms the core of risk assessment and CBA. So it is in these two respects that precautionary imperatives for transparency, deliberation, and accountability are arguably most acute. The usual forms of CBA involve quantifying both parameters (probabilities and magnitudes) as single scalar numbers. Thus, CBA typically operates in the top lefthand quadrant of Figure 3, in which all regulatory dilemmas are effectively conceived of as '***risk.***'

What Figure 3 shows is how this practice marginalizes appreciation for other important ways in which the knowledge underpinning governance of emerging technologies may be justifiably queried. There may, for instance, be important ***ambiguities*** concerning the defining, bounding, or prioritizing of contending issues, perspectives, and institutional, policy or technology options under appraisal.[33] These resulting dilemmas may concern the evaluative criteria to prioritize in appraisal, as well as the pros and cons of different innovation options under these criteria.[34] Even before attempting to determine how likely different outcomes may be, there may be radical disagreements over how to define and order the relevant options and issues. This includes many of the questions raised in the last section: what are the options to be evaluated in the first place; what is actually to count as 'benefit' or 'harm' – and to whom? And framed in which ways? In short, ambiguity addresses the dilemma (well established in the rational choice theory that underpins CBA) that in a modern pluralistic society, there can be no guarantee that there can exist a single rational ordering of preferences concerning these kinds of questions, let alone that any such ordering can be derived in any operational method.

On the vertical axis, there arise separate but related issues of ***uncertainty*** – in the strict sense of this term, referring to conditions under which it is impossible confidently to assign single probability distributions across potential outcomes. This may be due, for instance to limitations in historical evidence or the completeness of models; or to complexity, dynamism, nonlinearity or variability in open-ended systems; or to the diversity of issues often rolled into 'the human factor'. Under these conditions, according to the probability theorist de Finetti, "*probabilities do not exist.*"[35] – so to proceed as if they do is to succumb to "*pretence of knowledge.*"[36] And where confidence may be lacking equally in the determination of possibilities and probabilities alike (on both horizontal and vertical axes), there applies a condition of ***ignorance,*** reflecting basic limits to knowledge and the potential for surprise. This state under which '*we don't know what we don't know*' – is even more fundamentally intractable to probabilistic 'expected utility' methods.[37]

These deeper dilemmas of incertitude are not necessarily self-evident; they may be perceived to different degrees. Nor are they mutually exclusive; they may all arise together. Yet, even though they are sometimes acknowledged in the methodological fine print and in more academic risk policy literatures, they are typically downplayed in the regulatory practice of CBA. So, the central point emphasized in Figure 3 is that – while the different aspects of incertitude typically overlap – each may be better addressed by different methods, practices, and institutions.

Figure 3. Incertitude matrix

The practical upshot of all this is that methods that reduce outcomes to quantified "costs" and "benefits" may sound reassuringly scientific but are simplistic and misleading. When these methods attempt to compare different innovation options, they may seem to generate concrete orderings. But what Figure 1 shows is that, under uncertainty, ambiguity, and ignorance, not only the absolute numbers but even clear ordinal sequences are typically not possible. Given these problems, adherence to CBA as if it were a "sound science" response to the complexities of a challenge like synthetic biology is as irrational as the proverbial carpenter with only a hammer, for whom every problem looks like a nail.

As indicated in the lower and right hand quadrants of Figure 3, there are many concrete methods for regulatory appraisal that can help to address the wider dilemmas in the governance of synthetic biology. These methods offer means to broaden out the 'inputs' to appraisal – to enable consideration of alternative innovations, wider societal issues, divergent perspectives, and deeper uncertainties and ambiguities than are typically addressed in CBA. These more precautionary methods also aim to open up the 'outputs' of appraisal, such as to acknowledge that a single optimal resolution to these complexities is very rare. In order to effectively address this point, appraisal practices should enable exploration of a plurality of divergent possible conclusions concerning the pros and cons of different options in contrasting contexts, each explicitly qualified with its associated conditions.

To the extent that CBA makes use of (or is accompanied by), techniques that enable these qualities of 'broadening out' and 'opening up', then it too can be part of an effective – more rational and precautionary – response to the challenges presented by an emerging technology like synthetic biology. Indeed, the problem with CBA is not so much that the method itself is incapable of being used this way, but that associated disciplinary and organisational cultures tend to prevent this. Prevailing academic and political commitments can often be very strong to the supposed scientific (but actually political) ‘necessity’ of reduced numerical compressions of complexity.

The Case of MCM

What’s necessary to address these challenges in concrete ways can be illustrated by considering the method called multicriteria mapping. The important point lies not just in the technical details of MCM, but in its organizing norms, overall architecture, and associated body of practice. The guiding principles are listed in Figure 4, which summarizes key lessons from past debates about precaution in innovation governance.[38] In order to legitimately be considered precautionary, any technology appraisal process must possess these qualities. And an associated possible organizational architecture is illustrated in Figure 5, which shows a procedural framework for implementing a variety of regulatory appraisal methods like MCM. This framework could incorporate conventional risk-based methods (as shown in Figure 3) while still doing justice to the challenges of precaution.[39] Without respecting the wider principles outlined in Figure 4, appraisal is not likely to properly address the kinds of issues raised here. Without following a systematic process like that outlined in Figure 5, the process is not likely to be manageable.

Figure 4. Some general lessons of precaution by reference to historic regulatory episodes

Figure 5. An illustrative framework for precautionary regulation of emerging technologies

With these provisos understood, MCM has been used in a wide variety of contexts, including the appraisal of options for energy strategy,[40] food production,[41] environmental policy consultation,[42] and public health responses to the shortage of kidney donors.[43] It has been favorably reviewed as an academic research tool[44] and as a framework for policy appraisal.[45] Forming part of a wider “deliberative mapping” process, it has been recommended as an aid in high-level innovation governance processes,[46] including in roles like those described as “deliberative process” and “precautionary appraisal” in Figure 5.

What MCM offers is a hybrid analytic method-deliberative process that seeks (at least in principle) to integrate all the key features of precautionary appraisal outlined in Figure 4. The most fundamental principle in MCM is that participants (not facilitators, analysts, designers, or the sponsors of analysis) should be in the driver’s seat. This can be achieved in MCM partly by means of software that allows participants to develop their own appraisal and to interact with each other as they work. The process starts with an effort to characterize options. Attention then moves to defining the evaluative criteria. Each option is assessed under each criterion. Uncertainties are expressed by systematically distinguishing possible ‘pessimistic’ and ‘optimistic’ conditions. At every stage, great care is taken to elicit the reasons for the quantified judgements. Then criteria are weighted – also noting evaluative discussion - to reflect their relative importance. The final stage is to consider the resulting patterns in overall performance ranks.

The ability of other appraisal methods (like CBA) to 'broaden out' and 'open up' representation of diversity and complexity is often limited by structural features of those methods. With the principle that the participant is in the driving seat, MCM seeks to reduce such constraints. Perhaps the most important example of this is the way MCM extends the focus away from a single option (like a synthetic biology solution to a problem), in order to give balanced attention to a range of alternatives. It is a basic principle of MCM that a diverse array of options are selected at the outset and that participants can add new options at any time in the process.

Another common constraint is use of a predefined set of evaluative criteria. In the case of CBA, the ease with which different issues can be considered is biased by the metrics that are favored in this method (such as monetary values). Rather than weighing priorities directly and transparently, it is real or modelled money values that mediate implicit trade-offs. In MCM, by contrast, participants are challenged all the time qualitatively to justify their inputs; but they can select, define, measure, and prioritize their criteria as they wish. Nor is there any attempt to impose a single shared value tree on divergent criteria schemes. MCM also avoids imposing any dependence on expert assessments, instead allowing participants to undertake their own appraisal – giving reasons why they diverge from established evidence or analysis.

MCM allows different dimensions of outcomes to be traded off against each other, but it also allows participants to stipulate (with reasons) that some aspects are not subject to trade-offs. And the expression of uncertainty in MCM is also more open than is typically the case in CBA. MCM elicits a performance range between whatever participants consider to be reasonably 'pessimistic' or 'optimistic' scenarios. Again, as much attention is given to documenting qualitative reasons behind these scenarios as to quantifying scores. And at the end of a session, MCM allows each participant to review how their results will be reported. Unless a participant expresses satisfaction with how their findings are summarized, the results cannot be used.

Broadening Out and Opening Up

Conventional innovation governance, as exemplified in many uses of regulatory CBA, is deeply flawed. The perspectives of the most vulnerable and marginalized people are typically entirely excluded. Technologies and policies that might offer a preferable mix of pros and cons are routinely sidelined. Deep uncertainties, ambiguities, and ignorance are systematically understated and treated simply as "risk." And the enormous latitude for variability in the findings of "evidence-based" analysis is typically artificially reduced. In all these and other ways, regulatory appraisal tends to be shaped in advance so as to reinforce, rather than challenge, the powerful processes of lock-in and closure in research and innovation.

But there exist practical alternatives for enabling more rigorous and democratically accountable challenges to these processes. This has been shown here with the example of multicriteria mapping. It is not necessary entirely to reject improved forms of risk assessment or CBA – adjusted to address the challenges documented here. But in order to do justice to the imperatives of precaution under deep uncertainties, what any serious response must entail is considered use of a wide variety of alternative principles, practices, and methods that can broaden out and open up appraisal.

The task is not only to speak truth *to* power in the governance of emerging technologies but also to speak *about* power – and balance the effects it can have in regulatory appraisal. In the end, all of the methods mentioned here will only be as effective in addressing these challenges as the wider democratic structures and discourses within which they sit. And it is in helping to catalyze more rigorous and vibrant forms of innovation democracy that the deepest value of precautionary practice in innovation governance arguably lies.

References

[1] Commission of the European Communities, “Presidency Conclusions from the Lisbon European Council, March 23-24, 2000,” 2000; G. Brown, “Speech of 26th January 2004,” 2004; Commission of the European Communities, “Working Document on the Contributions of the Future Regional Policy to the Innovative Capacity of the European Union Committee on Regional Development” (Brussels, 2006).
[2] D. M. Byrd and C. R. Cothern, *Introduction to Risk Analysis: A Systematic Approach to Science Based Decision Making* (Oxford: Government Institutes Press, 2000).
[3] S. Funtowicz and J. R. Ravetz, *Uncertainty and Quality in Science for Policy* (Dordrecht, The Netherlands: Kluwer Academic Publishers, 1990).
[4] U. Felt et al., *Taking European Knowledge Society Seriously: Report of the Expert Group on Science and Governance to the Science, Economy and Society Directorate, Directorate-General for Research, European Commission*¸ ed. U. Felt and B. Wynne (Brussels: European Commission, 2008).
[5] IRGC, *Guidelines for the Appropriate Risk Governance of Synthetic Biology (Policy Brief),* Geneva: IRGC - International Risk Governance Council, 2010.
[6] Nuffield Council on Bioethics, *Emerging Biotechnologies: Technology, Choice and the Public Good - Summary,* London: Nuffield Council on Bioethics, 2012.
[7] F. Wickson and B. Wynne, “The Anglerfish Deception: The Light of Proposed Reform in the Regulation of GM Crops Hides Underlying Problems in EU Science and Governance,” *EMBO reports* 13, no. 2 (2012): 100–105; S. Jasanoff, *Designs on Nature: Science and Democracy in Europe and the United States* (Princeton, N.J. : Princeton University Press, 2005).
[8] Commission of the European Communities, “Analysis of Innovation Drivers and Barriers in Support of Better Policies: Economic and Market Intelligence on Innovation Social attitudes to innovation and entrepreneurship - social attitudes to innovation and entrepreneurship” (Brussels, 2012).
[9] R. Layard and S. Glaister, eds., *Cost-Benefit Analysis* (Cambridge: Cambridge University Press, 1994); E. J. Mishan and E. Quah, *Cost-Benefit Analysis* (London: Routledge, 2007).
[10] A. Stirling, “Precaution in the Governance of Technology,” in *Oxford Handbook on the Law and Regulation of Technology*, ed. K. Yeung (Oxford: Oxford University Press, 2017); B. Wynne, “Uncertainty and Environmental Learning: Reconceiving Science and Policy in the Preventive Paradigm,” *Global Environmental Change*, June 1992, pp. 111–127.
[11] A. Stirling, “Multi-Criteria Mapping: Mitigating the Problems of Environmental Valuation?,” in Valuing Nature? Ethics, Economics and the Environment, ed. J Foster (London: Routledge, 1997), 186- 210; A. Stirling and S. Mayer, “A Novel Approach to the Appraisal of Technological Risk: A Multicriteria Mapping Study of a Genetically Modified Crop,” Environmental Planning C: Government Policy 19, no. 4 (2001): 529-55; A. Stirling, P. Simmons, and C. Spash, Approaches to the Mapping of Values: A Review of Q-Methodology, Multi-Criteria Mapping and Attitudinal Scales , report of the European Commission CIVIC project, Brussels (Brussels: CIVIC project, 2003).
[12] A. Stirling, “Keep It Complex,” *Nature* 468 (2010): 1029–1031.

[13] M. O'Brien, *Making Better Environmental Decisions: An Alternative to Risk Assessment* (Cambridge Mass: MIT Press, 2000).
[14] N. Pollock and R. Williams, "The Business of Expectations: How Promissory Organizations Shape Technology and Innovation," *Soc. Stud. Sci.*, vol. 40, no. 4 (2010): 525–48.
[15] F. Bouder, D. Slavin, and R. E. Lofstedt, The Tolerability of Risk: A New Framework for Risk Management (London: Earthscan, 2007); U.K. Health and Safety Executive, Reducing Risks, Protecting People: HSE's Decision-Making Process (London: Her Majesty's Stationery Office, 2001); Office for Nuclear Regulation, Safety Assessment Principles for Nuclear Facilities (London: Office for Nuclear Regulation, 2006); B. Wynne, "Risk and Environment as Legitimatory Discourses of Technology: Reflexivity Inside Out," Current Sociology, 50 (2002): 459-77.
[16] A. Broers, *The Triumph of Technology* (Cambridge: Cambridge Univ. Press, 2005).
[17] R. Nelson and S. Winter, *An Evolutionary Theory of Economics Change* (Cambridge, Ma.: Harvard University Press, 1982); W. B. Arthur, *The Nature of Technology: What It Is and How It Evolves* (London: Penguin, 2009); G. O. for Science, "Annual Report of the Government Chief Scientific Adviser 2014 - Innovation: managing risk, not avoiding it - report overview," 2014.
[18] G. Dosi, *Innovation, Organization and Economic Dynamics - Selected Essays*. Cheltenham: Edward Elgar, 2000.
[19] A. Stirling, "Pluralising Progress: From Integrative Transitions to Transformative Diversity," *Environ. Innov. Soc. Transitions*. 1, no. 1 (2011): 82–88.
[20] A. Stirling, "Towards Innovation Democracy: Participation, Responsibility and Precaution in Innovation Governance," in Annual Report of the Government Chief Scientific Adviser 2014, Innovation: Managing Risk, Not Avoiding It. Evidence and Case Studies (London: Government Office for Science, 2014), 49-62, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/381906/14-1190b-innovation-managing-risk-evidence.pdf
[21] W. B. Arthur, "'Competing Technologies, Increasing Returns, and Lock-in by Historical Events,'" *Econ. J.* 99, no. 394 (1989): 116–131.
[22]. Ibid; P. A. David, "Clio and the Economics of QWERTY," *Econ. Hist.* 75, no. 2 (1985): 332–37; P. A. David, "Path Dependency and the Quest for Historical Economics: one more chorus of the ballad of QWERTY," Oxford, 1997.
[23] P. Sabatier, "Social Movements and Regulatory Agencies: Toward a More Adequate and Less Pessimistic Theory of 'Clientele Capture'," *Policy Sciences* 6 (1975): 301–342.
[24] A. Stirling, "Towards Innovation Democracy: Participation, Responsibility and Precaution in Innovation Governance," in *Annual Report of the Government Chief Scientific Adviser 2014, Innovation: Managing Risk, Not Avoiding It. Evidence and Case Studies, London: Government Office of Science, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/376505/14*, London: UK Government, 2014, pp. 49–62.
[25] K. Frenken, "The Early Development of the Steam Engine: An Evolutionary Interpretation Using Complexity Theory," *Ind. Corp. Chang.* 13, no. 2 (2004): 419–50; P. P. Saviotti and K. Frenken, "Export Variety and the Economic Performance of Countries," *J. Evol. Econ.* 18, no. 2 (2008): 201–218; K. Frenken, P. P. Saviotti, and M. Trommetter, "Variety and Niche Creation in Aircraft Helicopters, Motorcycles and Microcomputers," *Res. Policy* 28, no. 5 (1999): 469–88.
[26] A. Stirling, "Towards Innovation Democracy: Participation, Responsibility and Precaution in Innovation Governance," in *Annual Report of the Government Chief Scientific Adviser 2014, Innovation: Managing Risk, Not Avoiding It. Evidence and Case Studies, London: Government Office of Science, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/376505/14*, London: UK Government, 2014, pp. 49–62.
[27] Ibid.

[28] Ibid.; A. Stirling, "Precaution in the Governance of Technology"; J. D. Graham, "The Perils of the Precautionary Principle: Lessons from the American and European Experience," Heritage Lectures, no. 818 (2004): 1-4.
[29] A. Stirling, "Precaution, Foresight and Sustainability: Reflection and Reflexivity in the Governance of Science and Technology," in *Reflexive Governance for Sustainable Development*, ed. J.-P. Voss and R. Kemp (Cheltenham: Edward Elgar, 2006), pp. 225–72.
[30] A. Stirling, "Keep It Complex. ," *Nature* 468 (2010): 1029–1031.
[31] A. Stirling, "The Appraisal of Sustainability: Some Problems and Possible Responses," *Local Environment: The International Journal of Justice and Sustainability* 4, no. 2, (1999): 111-135; A. Stirling, "Renewables, Sustainability and Precaution : Beyond Environmental Cost-Benefit and Risk Analysis," *Issues in Environmental Science and Technology* no. 19 (2003): 113–134.
[32] A. Stirling, "On the Economics and Analysis of Diversity," Science Policy Research Unit, Brighton, 1998; Stirling, "Precaution in the Governance of Technology."
[33] A. Stirling, "Risk at a Turning Point?" Journal of Environmental Medicine [AUS: correct title?] 1, no. 3 (1999): 119-26; A. Stirling, "Risk, Uncertainty and Precaution: Some Instrumental Implications from the Social Sciences," in Negotiating Change: New Perspectives from the Social Sciences, ed. F. Berkhout, M. Leach, and I. Scoones (Cheltenham, U.K.: Edward Elgar, 2003).
[34] F. H. Knight, *"Risk, Uncertainty and Profit"* (Boston, MA: Houghton Mifflin, 1921); J. M. Keynes, "A Treatise on Probability.," *Philos. Rev.*31, no. 2 (1922): 180.
[35] B. de Finetti, *Theory of Probability - A Critical Introductory Treatment* (Chichester: Wiley, 1974).
[36] F. A. Hayek, *New Studies in Philosophy, Politics, Economics and the History of Ideas*, London: Routledge, 1978.
[37] B. Wynne, "Uncertainty and Environmental Learning"; S. Dovers and J. W. Handmer, "Ignorance, the Precautionary Principle, and Sustainability," *Ambio* 24, no. 2 (1995): 92-97, 1995; H. Brooks, "'The Typology of Surprises in Technology, Institutions and Development.'" IIASA / CUP, Cambridge, 1986.
[38] A. Stirling, On Science and Precaution in the Management of Technological Risk—Volume I: A Synthesis Report of Case Studies (Seville, Spain: Institute for Prospective Technological Studies, 1999); P. Harremoës et al., Precautionary Principle in the Twentieth Century (London: Earthscan, 2002); D. Gee et al., eds., Late Lessons from Early Warnings: Science, Precaution, Innovation (European Environment Agency Report No. 1) (Copenhagen, Denmark: EEA, 2013).
[39] A. Klinke et al., "Precautionary Risk Regulation in European Governance," *J. Risk Res.* 9, no. 4 (2006): 373–92.
[40] A. Stirling, *Multi-Criteria Mapping: Mitigating the Problems of Environmental Valuation* (Routledge, London, 1997); W. McDowall and M. Eames, "Towards a Sustainable Hydrogen Economy: A Multi-Criteria Mapping of the UKSHEC Hydrogen Futures," London: Policy Studies Institute, 2006, pp. 1–63.
[41] A. Stirling and S. Mayer, "Precautionary approaches to the appraisal of risk: a case study of a genetically modified crop" International Journal of Occupational and Environmental Health 6, no. 4 (2000): 296-311 ; Stirling and Mayer, "A Novel Approach to the Appraisal of Technological Risk"; S. Mayer and A. Stirling, "Finding a Precautionary Approach to Technological Developments: Lessons for the Evaluation of GM Crops," Journal of Agricultural and Environmental Ethics 15, no. 1 (2002): 57-71.
[42] L. S. Clark et al., "Local Outreach. Report to the UK Environment Agency," R&D Technical Report SWCON 204, University College Environment and Society Research Unit, London, 2001.

[43] G. Davies et al., "Deliberative Mapping : Appraising Options for Addressing 'the Kidney Gap,'" October, 2003.
[44] S. Yearley, "Mapping and Interpreting Societal Responses to Genetically Modified Crops and Food," *Soc. Stud. Sci*. 31, no. 1 (2001): 151–160.
[45] W. McDowall and M. Eames, "Towards a Sustainable Hydrogen Economy: A Multi-Criteria Mapping of the UKSHEC Hydrogen Futures," London: Policy Studies Institute, 2006.
[46] T. Horlick-Jones et al., "A Deliberative Future?," Norwich: Understanding Risk Programme, University of East Anglia, 2004.