

The Coming Breakthrough in Risk Research

Carlo Jaeger¹

Version of May 2, 2016

Abstract

The paper proposes a novel way to handle the relation between decision theory and uncertainty in the context of policy design. It starts from the fact that rich countries have developed a historically unprecedented capability to manage conventional risks – fire, floods, earthquakes etc., but also car accidents, many workplace risks, and more. It is based on two institutions – insurance markets and public risk governance – supported by a powerful theory: the expected utility approach to risk. Expected utility refines the utilitarian paradigm of rational action by combining the concept of utility functions with the concept of probability distributions. Probabilities may be used to represent the randomness of phenomena like the paradigmatic dice, as well as the incompleteness of information about a state of affairs like who won yesterday's horse race. One might think that future progress in risk research will consist mainly in refining this approach and spreading it to emerging and less developed countries. However, greater progress is necessary and possible. It is necessary because the global economy and technostructure we live in have generated new systemic risks –including financial crises, pandemics, climate change, nuclear war. These risks exceed the coping capacity of conventional risk management and call for new forms of integrated risk governance. Greater progress is possible because recent research has developed ways to address the basic difficulties of expected utility without losing its valuable insights. They involve three major advances. First, to introduce a risk function that generalizes expected utility so as to overcome well-known difficulties like the Allais paradox. Second, to embed expected utility in a framework of iterated network games so as to take into account the social learning processes that are essential for real world risk governance. And third, to accommodate the logic of complementary descriptions called for by the new systemic risks of the 21st century. The coming breakthrough in risk research may best be achieved by bringing these advances to bear on practical efforts aiming at integrated risk governance.

Keywords:

Risk, Uncertainty, Rationality, Integrated Risk Governance

JEL Classifications:

¹ Beijing Normal University, Arizona State University and Global Climate Forum. Comments from reviewers and readers are gratefully acknowledged, especially help by comments of reviewer 2 on non-standard logic and climate policy. Further comments are welcome to: carlo.jaeger@globalclimateforum.org.

B41, C73, D80

1. Introduction

In many areas of policy design, the relation between decision theory and uncertainty has been consolidated so as to provide indispensable tools for solving practical problems. In other areas – especially, but not only, in view of global problems – the situation is more problematic. In order to improve policy design in those areas, I propose a novel way to frame the relation between decision theory and uncertainty. I do not suggest to discard the conventional wisdom, quite the opposite. For further progress in this field, it is essential to deepen our understanding of the idea of rational action that informs current practices of risk management. Therefore, I will first consider this idea in some depth and then develop arguments that shall help to establish new research on and new practices of risk governance.

In today's high-income countries, risk governance is based on the capability to manage conventional disaster risks – earthquakes, floods, fires, etc. – with remarkable success (Kellenberg and Mobarak 2007). Even if hurricane Katrina caused a death toll of nearly 2'000 and the Tohoku earthquake caused more than 15'000 casualties, these are exceptional impacts in rich countries, while in developing countries much larger death tolls are still frequent.

The damages from Katrina were in the order of 100 billion \$, and those of Tohoku twice as large. But these figures actually represent rather small fractions of the American and Japanese capital stock in comparison with damages from natural disasters experienced in poor countries. Moreover, the impact of both Katrina and Tohoku would have been much smaller if the respective countries had used the best available and – for these countries – affordable technologies for early warning and protection.

The capability of industrialized countries to manage disaster risks is historically unprecedented. This success is due to the development of science and technology

together with the accumulation of capital and the establishment of professionally organized institutions. Altogether these are successes of Western rationality, a way to mobilize the self-interest of individuals, organizations and nations in order to solve problems one by one.

Over the past centuries, this kind of rationality has shaped the world. It has enabled industrialized nations to manage natural disasters with the help of insurance and re-insurance companies. And it has enabled all sorts of companies to manage business risks with the help of national governments and the financial sector. By the same means, it has become possible to manage a wide range of individual risks, from illnesses to car accidents, in ways that are far from perfect, but still extremely impressive by historical standards.

The growth path taken by today's high-income countries, however, has generated new systemic risks of planetary relevance – nuclear war, pandemics, mass extinction of species, climate change, global financial breakdowns, and more. Awareness of this situation goes back at least to the invention of the atomic bomb, that led Einstein to say (in a quote that has been reformulated in many ways): “Our situation is not comparable to anything in the past. It is impossible, therefore, to apply methods and measures which at an earlier age might have been sufficient” (Green, 2003, p.52).

The new systemic risks are unintended consequences of actions driven by a highly effective, but also problematic way of thinking (van der Leeuw 2012). As Charles Perrow (1984) observed long ago, it is the way of absolute rationality: a way of thinking that can be pursued by isolated actors optimizing the use of scarce resources for their respective ends. But the capability to manage the new risks cannot be based on the thinking and institutions that generated them. A key task of risk research, then, is to develop new ways of thinking that are adequate to this challenge and to investigate what kinds of institutions might sustain adequate patterns of risk governance (Renn 2008).

A first characteristic of the new kind of risks is that they affect humankind as a whole. This holds for the risks of global environmental change, for global financial instabilities, for the danger of pandemics, and more generally for what the OECD (2011) calls future global shocks.

A second characteristic of the new kind of risks is that they are the result of incredibly successful scientific and technological advances. This has created a profound tension between knowledge claims by scientifically trained experts and the doubts and beliefs entertained by other stakeholders and the public at large. Renn (2008, see also Renn and Klinke, 2004, and Renn and Walker 2008) has reacted to this tension by developing the concept of integrated risk governance, and by successfully testing procedures to reduce and sometimes overcome the problems it creates (see also Shi et al., 2013).

A third characteristic is the difficulty to consider those risks in isolation. This is due to a considerable extent to the fact that they arise as unintended consequences of a way of thinking that works precisely by considering problems in isolation (van der Leeuw, 2008). Non-linearities, stochasticities, complex feedbacks, lack of adequate concepts and various combinations of these can result in severe problems for this way of thinking – the rational actor paradigm (Jaeger et al., 2001).

In the 21st century, humankind will struggle with those new systemic risks. Hopefully, they will be dealt with successfully. This will require a major breakthrough in risk governance, and of course also in risk research. Next, I outline key features of the rational actor paradigm and its application to uncertainty, the expected utility approach (2). Then, I rehearse the anomalies of the paradigm that have been found by risk researchers (3). I then discuss three generalizations that hold promise to overcome those anomalies without losing the wealth of insights gained with the paradigm. The first generalization concerns the relation between risk and rationality within the expected utility approach (4). The second one

concerns the relation between optimization and social norms in the rational actor paradigm (5). The third one concerns the logic of complementary descriptions in risk research (6). As an example, I sketch how these generalizations can be used when addressing the risks of climate change (7). I conclude with an assessment of how the three generalizations hang together and how they may inspire future research (8).

2. Uncertainty and the Paradigm of Rational Action

Modern risk management emerged by the confluence of a theoretical and a practical development. On the one hand, mathematicians on the European continent laid the groundwork of probability theory in the 17th century. On the other hand, insurance contracts became increasingly common in London in the same century. The insurance business expanded hand in hand with the expansion of the transatlantic slave trade, and then experienced a major boost with the great London fire of 1666.

Since then, the concepts and methods used to manage risks have been greatly expanded, sustained by the core ideas developed in those times. These ideas have shaped the paradigm of rational action in the face of risk and uncertainty. They can be summarized in the following way. Faced with a risky prospect, a decision-maker should try to be as rational as possible. For this purpose, she needs first of all to identify three things (Savage, 1954):

- the set of outcomes, X , she is interested in (e.g. selling a successful new product, selling a new product without success, selling an old product with limited success).
- the set of Actions, A , she can undertake (e.g. making a new product or not)
- the set of conditions, S , that determine which actions may have which outcomes (e.g. the market may or may not be ready for a new product); this set then comes with a function E that takes an action and an outcome, (a,x) , and indicates under which condition $s=E(a,x)$ action a will have outcome x .

Then she needs to assess two things:

- how likely it is for any consequence to occur for the different actions; this should then be represented by a probability, $P(E(a,x))$.
- how strongly she prefers one consequence over another; this should then be represented by a utility index, $U(a)$.

On this basis, the expected utility of an act is defined as:

$$EU(a) = \sum_{h=1}^n U(x_h) \cdot P(E(a, x_h)) \quad (1)$$

In this view, rational actors will try to choose actions that maximize expected utility. Moreover, one should expect most actors to behave this way, because it provides the best strategy and actors who don't use it will be gradually marginalized.

In fact, the expected utility approach to action under uncertainty is a special case – and a very important one – of the rational actor paradigm that has shaped modern societies since centuries (Jaeger et al. 2001). In a way it started in the 17th century with the work of Thomas Hobbes, who introduced many ideas that later led to game theory (Eggers, 2011). One of the most important ideas of the rational actor paradigm is the strict separation of facts and values, that later led to the idea of exogenously given utility functions (Read, 2004).

According to this understanding of rationality, facts are what rational actors will agree about once they have access to the same information. To what amount probabilities can count as facts is a matter of dispute (see Eagle, 2010, for an overview). If I don't know whether the famous horse California Chrome won yesterday's race I may still be willing to place a bet implying a 90 percent probability for his victory. This probability expresses an attitude due to my incomplete information rather than a fact about yesterday's horserace. On the other hand it is usually treated as a plain fact that throwing a standard poker dice will yield an ace with probability one sixth. We need not go into the debate about the meaning of probability any further here. It suffices to clarify that stochasticity in the

mathematical sense – i.e. a family of random variables indexed by a semi-group like the natural or real numbers – can be used to study probabilities regardless of whether they are understood as “objective” – i.e. representing properties of phenomena independent from the observer – or “subjective” – i.e. representing states of mind of actors faced with different possibilities.

Things are somewhat simpler with utilites: to the rational actor paradigm, there are no moral facts, values are just what an actor happens to prefer. This leads to the problem whether individually rational actions by separate actors are not self-defeating. They would be so if they should lead to aggregate outcomes that some or all actors could improve upon according to their own preferences. The classical example for this problem is the prisoner’s dilemma. Among the many uses of the rational actor paradigm is the analysis of this kind of situations (Stewart and Plotkin, 2012).

The influence of the paradigm is based mainly on the claim that modern societies have discovered an institutional setting where individual rationality does indeed lead to the best possible outcome, as if guided by Adam Smith’s “invisible hand”. It is the setting of competitive markets, based on private property and binding contracts of exchange. According to this claim, competitive markets lead to the same outcome a fictitious benevolent planner would implement. In reality, however, the claim goes, no planner would be able to take into account all the information that decentralized markets can and must process in order to achieve that outcome.

Over the past centuries a huge amount of scientific research and practical application has gone into the refinement of the rational actor paradigm in general and of the expected utility approach to risk in particular. The latter has become the cognitive basis on which modern societies have learned to manage naturally given risks – from earthquakes to infant mortality – with amazing success, and to successfully engage in all sorts of projects – from shipping trade to space flight – that would have been impossible without new techniques of risk management.

It is useful to distinguish four kinds of elaboration of the expected utility approach (for an overview see Machina and Viscusi, 2014).

- 1) The collection of empirical data on all sorts of risks, from the first mortality tables compiled in London in the 17th century to the natural disaster database of Munich Re (2013).
- 2) The development of sophisticated mathematical tools to expand the scope of application of the approach. This includes methods to define probabilities over infinite sets, ways to use them as measures of uncertain beliefs, theorems about conditions for existence and uniqueness of solutions to the maximization problem and many more.
- 3) The linking of expected utility to public policy. This includes techniques for probabilistic cost-benefit analysis as well as ways to deal with the thorny problem of defining utility functions for political bodies, from problems in city planning to those of national and even international policy-making.
- 4) The linking of expected utility to the analysis of market economies. It is often overlooked that expected utility has become deeply interwoven with the study not only of insurance and financial markets, but with the development of marketing strategies like product bundling, with the study of investment decisions etc. What is more, the pervasive role of game theory in shaping present understanding of economic, social and biological phenomena would hardly be possible without the expected utility approach.

3. Anomalies

The expected utility approach enabled a cumulative growth of knowledge that provided the cognitive basis for modern risk management. Since several decades, however, an impressive amount of evidence has emerged that increasingly

challenges the whole approach (again, see Machina and Viscusi, 2014). Two research lines deserve special attention in this regard:

- 1) Empirical research highlighting serious limitations of the expected utility approach. E.g. it has been shown that decision-makers – including successful professionals with explicit training in using expected utility – often don't fit the expected utility format in their actual decisions. This has led to a series of paradoxes and anomalies, often named after their discoverers, like the Allais paradox, the Ellsberg paradox, the framing effects identified by Kahneman and Tversky, and more. These scientifically challenging results acquire significant practical relevance in view of the new systemic risks that have been triggered by the piecemeal approach of conventional risk management.
- 5) Theoretical research proposing alternatives to or generalizations of the expected utility approach in order to take those findings into account. This includes weighted utility theory, prospect theory and more. While these proposals still focus on an isolated decision-maker, advances in evolutionary game theory are opening up new avenues for the transition from the imaginary absolute rationality of an isolated actor to the social rationality advocated by Perrow (1984) in his seminal book on technological risks.

4. Risk Avoidance

Against this background, three kinds of generalizations seem necessary in order to conserve the insights gained from the expected utility approach without getting trapped by its limitations. The first generalization takes into account that there are good reasons for people to anchor their expectations in a particular kind of outcome, e.g. the worst case. Based on this anchoring, the probabilities of utility differences get specific weights. These weights can express the fact that a decision-maker may see a particular action as too risky because the spread between the different outcomes is too big, even if the expected utility looks quite attractive.

Buchak (2013) has shown that such anchoring and weighting effects can be elegantly captured by defining a “Risk adjusted expected utility” as follows. One orders the outcomes from worst to best and gives weights to probabilities by means of a risk function. However, the relevant probabilities now are not those of outcomes as such, but of the differences between outcomes. The result is function (2):

$$REU(a) = \sum_{h=1}^n [U(x_h) - U(x_{h-1})] \cdot R(P(E(a, \{x_h..x_n\}))) \quad (2)$$

Here, x_1 is the worst case and x_0 is a “null event” with utility and probability zero. R is the risk function giving weights to the different probabilities, with $R(0) = 0$ and $R(1) = 1$. In (1), the function E assigned one condition, a “state of the world” to pairs of an action and an outcome. Now, the function E takes pairs of an action and a set of outcomes. Accordingly, it yields whole sets of possible states of the world.

For the special case where $R(y) = y$ for all y , the risk adjusted expected utility of (2) reduces to the expected utility of (1). In the general case it gives decisions like those known from the Allais paradox as perfectly reasonable outcomes. They are due to the fact that on top of the kind of risk aversion captured in traditional decision theory by concave utility functions, humans also try to avoid risks with especially threatening worst cases – a behavior sometimes labelled as risk avoidance.

At first sight, it seems that REU cannot accommodate the ambiguity aversion documented in the Ellsberg paradox. There, one is faced e.g. with an urn containing one black and one white ball and a second urn containing two balls with an unknown combination of black and white (as usual in these problems, one considers an ideal setting without any biases). If told that in both cases picking a white ball yields, say, 100\$, and a black one nothing, most people prefer the urn with known proportions. The two cases yield not only the same expected utility, but also the same risk adjusted expected utility.

If one considers risk adjusted expected utilities in a social setting with repeated bets, however, this behavior turns out to be perfectly reasonable. This point was examined among others by Schneeweiss (1999). One can reformulate the problem in terms of a big urn containing a set of small urns. First, a small urn is picked from the big one, then a black or white ball from the small one. In the case where the distribution of black and white balls is known, the set of small urns is a singleton and one knows its distribution. When the distribution is unknown, the big urn contains a different small urn for each possible distribution. The result is multiplicative noise, which matters in many situations leading to log-normal or power-law distributions.

Now consider the task of picking a ball twice with replacement in an Ellsberg setting. Consider the case where one knows that the relevant urn contains a black and white ball. I pick the singleton urn, take a ball, put it back, the singleton urn is shaken by the master of ceremonies (or experimenter or whatever), I take a ball again: the probability of the worst case, i.e. picking black twice, is one quarter. When the composition of the urn in terms of color is unknown, there are three small urns (two black balls, one black and one white, two white balls). I pick one of the urns, take a ball, put it back, the selected urn is shaken, I take a ball again. The probability for the first ball to be black is still one half, but the first ball can only be black if the composition is not both white. So the probability of getting black again from the selected urn is larger than one half, and the probability of the worst case is larger than when the composition of the urn was known from the outset. The same holds for the best case, but if risk avoidance anchors expectations in the worst case, it is reasonable to prefer the urn with known composition. That repetition of action situations is a crucial ingredient of the human condition, in turn, leads to the second generalization of the expected utility approach.

5. Iterated Games

The second generalization recognizes that human beings live in social networks characterized by repeated action situations – from greeting each other to building airports and much more (Collins, 2004). Through these repetitions, people learn from experience and – usually even more – by observing each other. Moreover, they try to repeat action situations that they experienced as gratifying and to avoid those they experienced as frustrating. While the generalization of risk avoidance remains in the framework of individual rationality, the second generalization leads to the world of social rationality. Optimization becomes but one leg of evolving human action, the other being learning from experience, especially from the experiences of others.

As a result of this generalization, all the symbols in (2) implicitly or explicitly carry two additional indices: one for the actor and one for the time under consideration. This allows to combine the description of individual decisions with the indispensable description of how the individuals themselves as well as the networks they live in may evolve from one action situation to the next.

The result is a format of network games with a transition function that is more general than the replicator functions used in evolutionary game theory inspired from biology. The difference is due to the fact that at each stage, the actors do perform a boundedly rational optimization. The bounds of rationality are set by their past experiences and observations, and they shift as a result of current experience and observation (Gintis, 2010).

An important example of a transition function results if one considers a population of actors that at each stage revise their probability estimates (perhaps by Bayesian updating) as a function of the outcomes resulting from their previous actions. In (3), such a function takes the actions and probability assessments of all agents, (\mathbf{a}, \mathbf{P}) , at time $t-1$, together with a random variable, ξ , and yields the outcomes for all agents, \mathbf{x} ,

at time t together with their revised probability assessments, \mathbf{P} . The random variable, ξ , here represents all sorts of random factors, including changes in the network connecting the different agents, changes that in turn may influence the information available to them and the updating of their probability assessments.

$$(\mathbf{x}, \mathbf{P})_t = \Phi(\mathbf{a}, \mathbf{P}, \xi)_{t-1} \quad (3)$$

By setting a particular value of ξ , e.g. zero, as the default representing the system without random shocks, one can then decompose the dynamics of the system into a deterministic and a stochastic component. If one defines the fixed points of (3) with $\xi = 0$ as the solutions of the iterated game, then these solutions do not depend on the form of ξ . Still, this form can greatly affect the dynamic behavior of the system. However, the iterated game represented by (3) need not be meaningfully decomposable into a single deterministic structure (that may be characterized by a particular value of ξ) and random shocks perturbing it. If not, the solutions of the iterated game will depend on the form of ξ as well.

Iterated games of the form given by (3) can be analysed mathematically and implemented computationally in evolutionary agent-based models (Hallier, 2014). It can be used to study systems with a single stable equilibrium, as usually assumed in economics, as well as metastable systems with several locally stable equilibria, as seems particularly relevant for disaster risks. Catastrophic events like an earthquake, an industrial accident or a financial crisis can push a social system from one equilibrium into another one, and recovery from a catastrophe may lead to a still different equilibrium.

A metastable equilibrium consists in an action profile, \mathbf{a} , together with a neighborhood of similar action profiles such that the system stays in this neighborhood for a significant amount of time and would converge towards \mathbf{a} in the absence of the random shocks ξ .

If a society or social network is unable to maintain at least one such equilibrium for sufficient amounts of time, it is bound to disintegrate, as patterns of communication and interaction will break down. Social conventions, rules, norms and the like are means by which societies establish metastable equilibria. Along these lines, the format of iterated games allows to study the interaction between the dynamics of individual strategies and social structures (Young, 1998).

By generalizing the expected utility approach to the framework of iterated games the sets of possible conditions, $S_{i,t}$, that matter for each agent include the possible actions by the other agents. Moreover, it is straightforward to consider objects as part of actor networks (Latour, 2005). If their behavior is deterministic, for each possible condition, $s_{i,t}$, their utility function give a positive utility only to the outcome resulting from that condition together with the resulting behavior, and their probability assessments give probability one to the actual actions of the last period. If their behavior is stochastic, this is captured by the random variable included in the standard description of agents according to the second generalization.

Finally, from the point of view of the focal actor, a whole set of actors can be grouped as the environment (the word environment here becomes a technical term, to be distinguished from everyday notions of environment). These are the actors that are relevant for the problem at hand and somehow affect the focal agent, while the focal agent has no significant effect on them.

The shift to risk adjusted expected utility and the one to iterated games with transition (not just replicator) functions has four major implications for risk management. First, risk adjustment means that worst case considerations that could be considered negligible from an expected utility perspective may now call for much greater attention. Second, risk management must be seen much more systematically as an iterative process, often making reversible decisions more attractive and making explicit learning during implementation the rule, not the exception. Third,

risk managers must see themselves as partners in complex webs of risk governance, paying attention to the effects of their actions on the future strategies of the other agents. And fourth, it is crucial to distinguish between marginal and inframarginal measures. The former are based on the assumption that the system one is part of remains in the same basin of attraction over the time one is interested in. By contrast, the latter aim at a transition from one such basin to another one.

All four implications can be spelled out using the specific format of transition functions considered in (3). This format, however, presupposes that for each agent the decision problem of maximizing (2) has a unique solution – otherwise the iteration of (3) could in general only lead to a set-valued function. For the same reason (3) presupposes that action spaces and utility functions don't change in the course of time. Network structures at time t in turn must be fully determined by the triple of outcomes, probability assessments and random factors at the same time – if they would also depend on the structure at time $t-1$ an additional time lag would be needed in (3).

By relaxing these assumptions, one gets the following format for a transition function:

$$(\mathbf{x}, \mathbf{X}, \mathbf{A}, \mathbf{U}, \mathbf{P}, \mathbf{R}, \mathbf{\Gamma})_t = \Psi (\mathbf{x}, \mathbf{X}, \mathbf{a}, \mathbf{A}, \mathbf{U}, \mathbf{P}, \mathbf{R}, \mathbf{\Gamma}, \xi)_{t-1} \quad (4)$$

The variables are as follows (with time indices left out for simplicity):

\mathbf{x} , set of realized outcomes

\mathbf{X} , set of possible outcomes

\mathbf{A} , set of possible actions

\mathbf{U} , set of utility functions

\mathbf{P} , set of probability distributions

\mathbf{R} , set of risk functions

$\mathbf{\Gamma}$, network structure

\mathbf{a} , set of realized actions

ξ , random variable

(4) leads to the kind of structures often described as complex systems. It can be used to organize one's thinking about a wide array of important phenomena and problems. Moreover, it leads to an additional, critical insight for risk management and governance. This is the acknowledgement that in order to become practically useful, (4) has to be specified with the help of major additional assumptions about the particular problem at hand. In other words, it is dangerous to try to manage risks without substantial knowledge about the specific domain in which these risks arise. The study of complex systems can become quite misleading without acknowledging the indispensability of domain specific knowledge.

The need for additional assumptions results from a key difference between (3) and (4). In (3), the relevant dynamic mechanisms are reasonably familiar. Micro-economic optimization, bayesian updating of subjective probabilities, payoffs from interacting strategies by different players have all been investigated in a plethora of empirical and theoretica studies. In (4), a whole array of dynamical patterns is introduced about which very little can be said in general terms. Examples are the dynamics of utility functions or of networks in multi-agent models.

The question then arises of whether it is reasonable to look for practically relevant additional assumptions at this level of generality. If such attempts should succeed, that would be an astonishing achievement in the social sciences. And even if they should fail, attempts in this direction may turn out to be fruitful in other respects (for an interesting example of such an attempt, see Gintis, 2010).

6. Complementary Descriptions

It is likely, however, that many more practically relevant insights can be gained by a different strategy: specializing and modifying the generic transition function (4) in view of particular practical problems. In a way, this is less in line with the attempt to

use mathematics according to the template of physics, and closer to the spirit of “phronetic research” advocated and practiced by Flyvbjerg et al. (2003) in their work on the specific risks arising in the management of megaprojects.

Why then not discard generic mathematical structures like (4) altogether and limit the role of mathematics to help fitting equations to statistical data and to provide intuitively appealing metaphors like the prisoner’s dilemma? Because the fabric of insights and beliefs that is currently available to orient ourselves in the global economic system is firmly rooted in mathematical structures: in the supply, demand, production, utility and other functions that inform economic discourse as well as decision-making. If risk research were to develop so as to ignore these structures, it would become at best irrelevant, more likely misleading. Therefore, if risk research is to be useful in the face of the risks of the future it is vital to preserve mathematical formats while embracing the two generalizations introduced in this section: from expected utility to risk adjusted expected utility and from plain optimization problems to a framework of iterated games in complex networks.

In the perspective developed here, risk research needs mathematics to overcome the limitations of current economic analysis in the face of new systemic risks, and can do so best in combination with well-documented stories about particular practical experiences, be they the management of megaprojects or experiences of droughts, failed strategies to overcome organized crime or successful instances of environmental policy. So far, in-depth case studies have rarely involved sophisticated mathematical research, and path-breaking mathematical modeling has rarely related to the “thick descriptions” provided by sophisticated narratives (Kaploun, 2013). Risk research will provide plenty of opportunities to overcome this impasse in what will be a true breakthrough in risk research.

When using mathematical structures to represent, analyze and tackle problems of risk and uncertainty, however, one of the most elementary of those structures, the one of classical logic, brings pitfalls of its own. Given an arbitrary proposition “A”,

this structure suggests that while we may not know whether “A” or “not A” is true, one of the two is always the case. And given a second proposition “B” and writing “~” for “not”, there are exactly four possibilities: “A and B”, “A and ~B”, “~A and B”, “~A and ~B”. We may not know which one is the case, but we do know that one of them is. To use a venerable philosophical terminology, uncertainty is always an epistemological, never an ontological issue.

Already Aristotle, father of classical logic, struggled with the fact that this view makes it very difficult to think about the openness of the future that we take for granted in real life (Belnap et al., 2001). The openness of the future is of course ubiquitous in risk research. Modern science, however, has produced an impressive image of a deterministic world, an image that makes it hard to take the openness of the future rigorously into account when thinking about risk.

Fortunately, the findings of quantum mechanics have led to serious and on-going discussions about the relations between uncertainty, indeterminacy and logic. An important outcome of those discussions is the understanding of complementary descriptions made possible by the work of Kochen and Specker (1967). They proved a seminal theorem about quantum mechanics, showing that not all variables that one may introduce to describe a standard quantum system can have well-defined values simultaneously (Liang et al., 2011). Here we are interested in the generalization of classical logic that they introduced, because it is relevant for tackling uncertainty and ambiguity in decision-making (the key idea is introduced beautifully in Specker, 1960; see also Vetterlein, 2011, and Dowker et al., 2105).

Consider a language L consisting of propositions $\alpha, \alpha', \alpha'', \alpha''', \dots$. They are based on elementary propositions a, a', a'', a''', \dots . Considered in isolation, each elementary proposition is either true or false. Elementary propositions can be transformed into secondary propositions by the prefix “not” (\sim for short). If a proposition α is true, $\sim\alpha$ is false, if α is false, $\sim\alpha$ is true. So far, we operate within the language of classical sentential logic.

But now we add a connective for compatibility (\diamond for short) so as to formulate further propositions like: α is compatible with α' , or $\alpha \diamond \alpha'$ for short. And just as some elementary propositions may be true or false, propositions about compatibility may be true or false, too. Compatibility is reflexive, so we always have $\alpha \diamond \alpha$. It is also symmetric, so if $\alpha \diamond \alpha'$ is true, $\alpha' \diamond \alpha$ is true as well.

For compatible propositions we define the connective “and” (\wedge for short). Just as in classical logic, we add the rule that $\alpha \wedge \alpha'$ is true if both α and α' are true, false otherwise. And just as in classical logic, by using brackets one can then formulate still more complex propositions like $\sim(\sim\alpha \wedge \sim\alpha')$. The last proposition can be used to define a further connective “or” (\vee for short), by defining $\alpha \vee \sim\alpha'$ as a shorthand for $\sim(\sim\alpha \wedge \sim\alpha')$. The other classical connectives for implication and equivalence can be constructed analogously. The difference with classical logic is that connectives are defined only for compatible propositions.

The resulting logic may be called compatibility logic. It contains classical logic as a special case, arising when all propositions are compatible with each other.

If a language L contains incompatible propositions, it has at least two disjoint subsets of maximally compatible propositions. These subsets form complementary descriptions of whatever domain of discourse L may refer to. Complementary descriptions are not inconsistent, but if we take one of them to be true or false, the other ones become meaningless. Let $A := \alpha, \alpha', \alpha'', \dots$ and $B := \beta, \beta', \beta'', \dots$ be two incompatible descriptions. Then the propositions $\alpha, \alpha \vee \alpha', \dots$ may be true or false, and so may $\beta, \beta \vee \beta', \dots$. But $\alpha \wedge \beta, \alpha \vee \beta, \dots$ are not meaningful propositions. Given α , there is no β to be known and the other way round.

Historically, the first and in a way still paradigmatic example of complementary descriptions is given by the fact that electrons, photons etc. sometimes have to be

described as particles, sometimes as waves. The question of how the findings of quantum mechanics and the formulas expressing those findings can be understood is notoriously difficult, and we will not enter this highly controversial discussion here. From the point of view of decision theory, what matters is the fact that an experimenter can decide which one of several complementary propositions she wants to test. E.g. she can decide whether she wants to measure position or momentum (or some combination of probable positions and probable momentums). Moreover, she can switch back and forth between complementary descriptions. But each time she goes for one she loses, as it were, the other ones.

Now while it is often said that classical mechanics is close to common sense while quantum mechanics is highly counter-intuitive, in this respect the opposite seems true. Everyday life is full of situations where by choosing one course of action we will find out certain things while losing the ability to find out others. A classical example is the question of how one's life would have continued if one had married another person than one's actual spouse. Another one is the question of what would be the name of, say, today's emperor of Germany if that country had won World War I. The point is not that we cannot be certain about possible events, in many cases we can. The point is that often possible consequences of human choices are only defined once a choice has been made.

Risk research is faced with such situations quite often, although of course not always. We know a lot about the future, and often we know enough to assess the relevant consequences of different courses of actions we are faced with. The critical word here is "relevant". Suppose I have well-defined preferences over combinations of two variables A and B . But now I am faced with a choice where one action makes it possible to forecast A and some variable C , with B being undefined, while another action makes it possible to forecast B and C , with A undefined. Then, variables A and B are relevant, while C is not. Then there is no rational choice between the two actions. What human beings do in such situations is to develop some new preferences, perhaps in an *ad-hoc* way just for the particular situation, perhaps in

more far-reaching ways. But about how to do this the rational actor paradigm remains silent.

7. The Case of Climate Change

An important example for the relevance of complementary descriptions is the current challenge of anthropogenic climate change. From a rational actor perspective, the standard approach is to start with the fiction of a benevolent planner, in order to gain a benchmark against which to assess actual events and possible policies. This has led to the following list of specifications (for one of the most comprehensive attempts in this direction see Nordhaus, 2013).

- 1) The set of possible *outcomes* involves possible present values of future global GDP per capita, sometimes complemented by monetary equivalents of non-monetary damages like losses of human lives and species extinction, rarely by some measure of fairness as well.
- 2) The set of possible *actions* is specified by different trajectories of global emissions, usually given as sequences of amounts of greenhouse gas equivalents emitted in the atmosphere. Notice that this assumes the benevolent planner to be in a position to implement such trajectories.
- 3) In theory, the set of possible *conditions* would be specified as features of the climate system, the biosphere, global human population, world economy and world society that together would determine the effect of different emissions trajectories on outcomes like the present value of future GDP trajectories. In practice, one assumes three functions:
 - a. A first one, given by a discount factor, to get a net present values from trajectories of GDP per capita or a similar variable like consumption per capita.
 - b. A second one to get trajectories of GDP or similar variables from trajectories of temperature. Pindyck (2011), e.g., uses the function $C = C^* / (1+T)$, where C^* is aggregate consumption without global

warming, T the increase in global mean temperature and C aggregate consumption under global warming.

- c. A third function to get trajectories of global mean temperature from trajectories of emissions. A possible approach here is to use some model for how carbon concentrations depend on emissions and then rely on the hypothesis that the “annual rate of temperature increase is [...] linearly related to the rate of increase of cumulative emissions” (Matthews et al., 2012).

The set of relevant conditions then corresponds to the functional forms and parameter values that one considers possible for this concatenation of functions.

- 4) In the case of climate change, the assignment of *probabilities* to the different states of the world that are considered relevant can hardly be justified by statistical analyses, because we are dealing mostly with unique events in the future that are not subject to any well-known law. So the probabilities are subjective ones, supposedly expressing the more or less reasonable attitudes of the benevolent planner towards the risks at stake. Of course, in fact they express the attitudes chosen by the researchers when setting up a model for the fictitious benevolent planner.
- 5) A similar situation arises with the *utility* assigned to the different outcomes considered, e.g. the present value of future aggregate consumption. Not only are such utility assignments necessarily subjective – in the rational actor paradigm there is no doubt about that – but the only subjectivity that can be projected on the global benevolent planner is again the one chosen by the researchers trying to describe that fictitious figure.

The emphasis on the subjective aspects of both probability and utility may sound like a weakening in comparison with standard practices in statistics, atmospheric physics and other fields relevant for climate research. But it can provide decision-makers with less misleading arguments than those resulting from misplaced application of those standard practices.

A good example is given by the measurement problems that loom large in climate related research. Understanding the inevitably subjective character of the probabilities and utilities used in the analysis makes it easier to realize that many measurement problems are more important for the internal needs of the scientific enterprise than for actual decision-making in climate policy. For the remaining ones, researchers can define and document a-priori distributions and improve them stepwise as useful evidence becomes available.

The fundamental difficulty with the attempt to analyse the climate challenge in terms of the rational actor paradigm lies elsewhere. It is the one highlighted by complementarity logic: our actions enable us to know certain things, actual or possible, while making others unknowable, not because of our cognitive limitations, but because they simply are not there to be found.

For our present purposes it is useful to distinguish two kinds of issues in climate policy (for background on each of these see the monumental IPCC 2013-2014, which however does not offer the grouping proposed here). The first kind is given by issues we know beyond reasonable doubt.

1. Anthropogenic climate change is causing impacts many people consider undesirable, like:
 - 1.1. Glaciers visibly melting all over the world, although with exceptions
 - 1.2. Heat waves increasing in many regions, including densely populated areas
 - 1.3. Traditional landscapes changing because of species migration
 - 1.4. And more
2. Anthropogenic climate change can trigger global climate catastrophes – like sea-level rise of 15 meters and more – in the coming centuries and millennia
3. Global climate catastrophes will hardly materialize before the end of the present century, to say that they will be unavoidable if we continue current emissions trends for another ten or even fifty years underestimates the long-term

possibilities of technological solutions – whether those will be desirable is another question.

4. To stabilize the global climate systems global net emissions would need to shrink to near-zero levels
5. Ill designed climate policies can cause massive economic, social and geopolitical disruption
6. With well-designed climate policies, the world economy can reach such emissions levels at negligible costs and possibly even with overall gains in the course of the current century

Further research on this kind of issues is certainly needed. It will refine the statements made above, not overturn them (like research about the exact speed of light was and is important to get better estimates, without putting the order of magnitude of about 300'000 km/sec already estimated by Bradley in 1728 into question).

The second kind of issues is given by problems about which very little is known but where we can gather substantive knowledge by embarking on particular paths of socio-economic development:

- I. Causal paths for climate change influencing global GDP. If humankind stays on the present growth path for another century, we will be able to learn a lot about such mechanisms. The knowledge so gathered will make formulas presently used to represent those mechanisms look as little more than oddly naïve.
- II. Reasonable priors for dates of global climate catastrophes. Somebody betting on a toss of a normal coin with priors assigning probability 1 to heads and 0 to tails betrays a lack of understanding of what regular coins are. Presently, the very act of stating priors on possible dates for climate catastrophes betrays a similar lack of understanding. If humankind should turn a large-enough amount of the fossil fuel reserves available on Earth into net greenhouse gas

emissions, however, we would learn what reasonable priors might be, and from there could elaborate better and better posteriors.

- III. Causal paths for ambitious global emissions reductions affecting global GDP. In learning studies, the idea that learning is fostered by suitable challenges has a long-standing tradition (Ingle, Dagworth, 2013). Related phenomena in economic growth are poorly understood. Examples are the fact that the long-term economic performance of The Netherlands has not been hampered by the challenge of maintaining and expanding below sea-level territory, or that economic growth in Switzerland has not been impaired by introducing in 1876 the simple rule that forest area should not be reduced (Lehmann et al. 2007, p.20ff). If in the coming decades humankind should embark on the collective experience of drastically reducing greenhouse gas emissions, we would have the opportunity of learning from successes and failures what are the mechanisms governing the impact of global climate policy on world GDP.
- IV. Opportunities and pitfalls in developing long-term global risk governance. As emphasized in the introduction, in high income countries the combination of insurance markets and nation states has allowed to turn the rational actor paradigm into remarkably successful systems of risk governance. New systemic risks like climate change call for new patterns of global risk governance. Whether the combination of the UN and IPCC will be sufficient for this purpose is an open question. And without actually developing effective global climate policies we will hardly find out what are key opportunities and pitfalls in developing long-term global risk governance.

Complementarity logic implies that by continuing to emit large and possibly increasing amounts of greenhouse gases we will be able to gather reliable knowledge about issues I and II, but not about issues III and IV. The converse is true if we will reduce emissions to near-zero levels in the coming decades. There even may be a feasible intermediate path where we can develop rough guesses, but no reliable knowledge about all four issues. In none of these cases will a comparison of

long-term GDP levels resulting from different strategies provide a reasonable basis for the design of global climate policy.

This is not to say that attempts to simulate the GDP resulting from the actions of a fictitious benevolent planner cannot be useful. They can give a sense of relevant orders of magnitude, of unavoidable uncertainties, even of ethical choices that will be made explicitly or implicitly. What they cannot do is to show the one best strategy for tackling climate change.

This does not mean that complementarity logic reduces the choice of climate policies to an irrational act, as would be the case if it was a pure matter of value judgements and if such judgements were never amenable to reasonable conversations. Points 1. to 7. are sufficient to lead many people to look for and support active climate policies and to bring their views into democratic debates and decision making. They are not sufficient to treat anybody who does not emphatically support such policies as ignorant and worse.

Moreover, points 1. to 6. definitely imply that research on how best to design and implement such policies is of the utmost importance. Such research cannot be satisfied with looking for answers in an imaginary world where all variables of interest would have well-defined values, it has to take into account the opportunities and limitations for research following from the actual developments resulting from a myriad of contingent decisions.

How then can an analysis of the climate challenge taking into account the insights of complementarity logic proceed? First of all, by starting with an actual agent – a government, a company, an association, an individual. As an example, consider a national government, say the one of Germany. Once a focal agent has been chosen, it can be situated in the broad setting of iterated games outlined above.

This requires at least roughly identifying the agents that the focal one interacts with in view of the problem at hand. Here, interaction means that the behavior of the focal agent matters to them and their behavior matters to the focal agent. The interaction may be direct or mediated by further agents. Some agents may be grouped into aggregate agents, e.g. the car industry, its customers, etc.

In the case of German climate policy, relevant agents may include the electorate, public opinion at home and abroad, politically salient countries from a German point of view (like France, Poland, the EU, U.S. and China), the scientific community, the German and European economy and last not least the global climate. How large and differentiated the set of players has to be will depend on the particular analysis to be performed.

For the purpose of analysis, the rest of the world is then treated as the environment. As explained in section 5, environment here is a technical term. In view of climate policy this means that the climate system is treated as an agent, albeit of course not a human one. One may object that Germany alone can hardly affect the global climate system, but the point of the analysis is to identify German strategies by which some networks of which Germany is part can indeed make a difference for the global climate of the future (Tabara et al., 2013). And of course, the global climate does affect Germany. So, for an analysis of German climate policy, global climate is an agent. On the other hand, the internet does affect Germany, too, but the influence of Germany on the internet may be quite small, and so for an analysis of German climate policy it may be part of the environment.

To develop a full analysis of climate policy for an agent like Germany lies beyond the scope of this paper. The key point here is that the kind of uncertainty highlighted by complementarity logic can be dealt with in a framework of social rationality. And even at the present stage of a very preliminary outline a fresh perspective can be fruitful. For a start, the goals of an agent like the German government can hardly be captured by the maximization of a utility function expressing the net present value

of future GDP, and even less by an immediate concern to avoid dangerous climate change. Rather, key goals may be to stay in power and perhaps to leave a legacy that will be recorded positively in the history books. From these, increasing GDP and trying to avoid dangerous climate change may then follow as secondary goals.

The first practical implication then is that German climate policy cannot and need not be justified by a cost-benefit calculus based on a comparison between short-term losses in GDP from climate policy and long-term damages to GDP from climate change. It cannot be so justified because the only way to find out actual GDP in, say, 2100 without climate policy, would be by not implementing climate policy and seeing what would happen. And the only way to find out GDP in 2100 under conditions of stringent climate policy would be to realize such a policy. These two possibilities correspond to incompatible propositions, they cannot be known together.

German climate policy need not be justified by a trade-off between future damages and present losses in GDP, because the only realistic policies are ones with negligible or no GDP losses in the short term. Once the economy is understood as a metastable system with a variety of possible equilibria, however, it is perfectly reasonable to search for a low-carbon equilibrium whose GDP growth is at least as high as the one of the present growth path. In this search, Germany will take some action – declaring climate policy goals, pushing for renewables, phasing out nuclear, sticking to coal, experimenting with electric vehicles, etc. – and observe what other relevant agents do. On the basis of experiences observed and made, next steps will then be taken in the kind of sequential decision-making described in section 5.

For other agents, be they governments or not, similar analyses can be made, with results that may or may not be similar to the one sketchily outlined above (Jaeger et al, 2012). Climate risks form but one of many areas of risk research where a conscious effort to move beyond the rational actor paradigm is warranted. The

overarching task is to embed the individual rationality we are familiar with into a social rationality yet in the making.

8. Conclusion

The successes of the rational actor paradigm and its application to risk management through the expected utility approach are undeniable. However, research testing the paradigm has discovered serious anomalies that call for a broader view of risk and rationality. Moreover, the paradigm itself has contributed to the emergence of new systemic risks that require such a broader view, too.

From the research side, three elements of a more comprehensive paradigm can be identified: the introduction of a risk function different from and equally relevant as utility functions for risk aversion, the specification of iterated games among heterogeneous agents where optimization alternates with adaptive learning, and the acknowledgement of complementary descriptions as a key feature of many risk situations. They are linked by a more pragmatic understanding of rationality than the one that has shaped the present world economy and the mechanisms of risk governance on which it relies.

Not surprisingly, research can only contribute to such a pragmatic view of rationality by engaging with actual practitioners. This is especially appropriate for risk research. The three generalizations discussed in the present paper, therefore, can best be brought to fruition in research about new systemic risks. In the previous section I have sketched how this might be done with regard to climate risks. Risk research along similar lines will increasingly be feasible and necessary with regard to other kinds of new systemic risks, including those of global financial instability, of pandemics, of new technologies leading to unintended, but massive harm, of perhaps unlikely, but by no means impossible large scale military conflict, and more. The closer the interaction between what at first sight might look like basic research

and the very practical efforts to deal with new systemic risks, the sooner can the coming breakthrough in risk research be achieved.

References

Belnap, N., Perloff, M., Xu, M., Bartha, P., Green, M., Horty, J. (2001) *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford, Oxford UP.

Buchak, L. (2013) *Risk and Rationality*, Oxford, Oxford University Press.

Collins, R. (2004) *Interaction Ritual Chains*. Princeton, Princeton University Press.

Dowker, F., Kakas, A., Toni, F. (2015) *Argumentation logic as an anhomomorphic logic*. Paper presented at The 10th Panhellenic Logic Symposium, 2015, [www.doc.ic.ac.uk/~ft/PAPERS/AL-Quantum-9\(CR\).pdf](http://www.doc.ic.ac.uk/~ft/PAPERS/AL-Quantum-9(CR).pdf).

Eagle, A. (2010) *Philosophy of Probability: Contemporary Readings*, London, Routledge.

Eggers, D. (2011) *Hobbes and Game Theory Revisited: Zero-Sum Games in the State of Nature*. *Southern Journal of Philosophy*. 49, 193–226.

Flyvbjerg, B., Bruzelius, N., Rothengatter, W. (2003) *Megaprojects and Risk: An Anatomy of Ambition*. Cambridge, Cambridge University Press.

Gintis, H. (2010) *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences* Herbert Gintis. Princeton, Princeton University Press.

Green J. (ed.) (2003) *Albert Einstein*. Melbourne, Ocean Press.

Hallier, M. (2014) *Formalisation and Metastability Analysis of Agent-Based Evolutionary Models*. PhD thesis in Mathematics and Computer Science at the Free University of Berlin.

Ingle, S., Dagworth, V. (2013) *Teaching and Training Vocational Learners*. London, Sage.

IPCC (2013-2014) 5th Assessment report. <http://www.ipcc.ch/index.htm> [retrieved May 1, 2016]

Jaeger, C., Hasselmann, K., Leipold, G., Mangalagiu, D., Tabara, J.D. (2012) Reframing the Problem of Climate Change. From Zero Sum Game to Win-Win Solutions. New York, Earthscan.

Jaeger, C., Renn, O., Rosa, E.A., Webler, T. (2001) Risk, Uncertainty, and Rational Action. Earthscan, New York.

Kaploun, V. (2013) From Geertz to Ryle: The Thick Description Concept and Institutional Analysis of Cultures. Working Paper WP20/2013/01. Moscow, The National Research University Higher School of Economics.

Kellenberg, D.K., Mobarak, A.M. (2007) Does rising income increase or decrease damage risk from natural disasters? *Journal of Urban Economics*, 63, 788-802.

Kochen, S., Specker, E. (1967) The Problem of Hidden Variables in Quantum Mechanics. *Journal of Mathematics and Mechanics*, 17, 59-87.

Lehmann, B., Steiger, U., Weber, M. (2007) *Landschaften und Lebensräume der Alpen. Zwischen Wertschöpfung und Wertschätzung*. Zurich, vdf.

Liang, Y.-C., Spekkens, R.W., Wiseman, H.M. (2011) Specker's Parable of the Overprotective Seer: A Road to Contextually, Nonlocality and Complementarity. *Physics Reports*, 506, 1-39.

van der Leeuw, Sander (2012) Global systems dynamics and policy: Lessons from the distant past. *Complexity Economics*, 1, 33-60.

van der Leeuw, Sander (2008) Climate and society: lessons from the past 10,000 years. *Ambio*. 37(14), 476-82.

Machina, M., Viscusi, K. (2014) *Handbook of the Economics of Risk and Uncertainty*, Volume 1. Amsterdam, North-Holland.

Matthews, H.D., Solomon, S., Pierrehumbert, R. (2012) Cumulative carbon as a policy framework for achieving climate stabilization. *Philos Trans A Math Phys Eng Sci.* 1974, 4365–4379.

Munich Re (2013) NatCatService. Downloadcenter for statistics on natural catastrophes. www.munichre.com/natcatservice (retrieved 23 August 2015).

Nordhaus, W.D. (2013) *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*. New Haven, CT, Yale University Press.

OECD (2011) *Future Global Shocks: Improving Risk Governance*. OECD, Paris.

Perrow, C. (1984) *Normal Accidents. Living with High-Risk Technologies*. New York, Basic Books.

Pindyck, R.S. (2011) Fat Tails, Thin Tails, and Climate Change Policy. *Review of Environmental Economics and Policy*, 5, 258-274.

Read, D. (2004) *Utility theory from Jeremy Bentham to Daniel Kahneman*. Working Paper No: LSEOR 04-64. London, The London School of Economics and Political Science.

Renn, O. (2008) *Risk governance: coping with uncertainty in a complex world*. London, Earthscan.

Renn O, Klinke A. (2004) Systemic risks: a new challenge for risk management. *EMBO Reports*. 2004;5(Suppl 1):S41-S46. doi:10.1038/sj.embor.7400227.

Renn, O. and Walker, K. (eds.) (2008) *Global Risk Governance. Concept and Practice Using the IRGC Framework*. International Risk Governance Council Book Series 1. Berlin and Heidelberg, Springer.

Savage, L. (1954) *The Foundations of Statistics*. Hoboken, N.J., Wiley.

Schneeweiss, H. (1999) Resolving the Ellsberg Paradox by Assuming that People Evaluate Repetitive Sampling. In: Leopold-Wildburger, U., Feichtinger, G., Kistner, K.-P. (eds.) *Modelling and Decisions in Economics: Essays in Honor of Franz Ferschl*. Heidelberg, Springer.

Shi, P., Jaeger, C., Ye, Q. (2013) Integrated Risk Governance. Science Plan and Case Studies of Large-Scale Disasters. Beijing and Berlin, Springer.

Specker, E. (1960) The Logic of Propositions which are Not Simultaneously Decidable (original in German). In: Hooker, C.A. (ed.) (1975) The Logico-Algebraic Approach to Quantum Mechanics, Vol. I, Dordrecht, Riedel.

https://www.researchgate.net/publication/50599843_E_Specker_The_logic_of_non-simultaneously_decidable_propositions1960

Stewart, A., Plotkin, J. (2012) Extortion and Cooperation in the Prisoner's Dilemma, Proceedings of the National Academy of Sciences, 109, 10134–10135.

Tabara, J.D., Mangalagiu, D., Kupers, R., Jaeger, C., Mandel, A., Paroussos, L. (2013) Transformative targets in sustainability policy making: the case of the 30% EU mitigation goal. Journal of Environmental Planning and Management, 56, 1180-91.

Young, H.P. (1998) Individual Strategy and Social Structure. Princeton, Princeton U.P.

Vetterlein T. (2011) Partial Quantum Logic Revisited. Int. J. General Systems 40(1), 23-38.