Int. J. Disaster Risk Sci. 2010, 1 (1): 10–16
doi:10.3974/j.issn.2095-0055.2010.01.003
ARTICLE

Risk, Rationality, and Resilience

Carlo Jaeger*

Potsdam Institute for Climate Impact Research (PIK), 14473 Potsdam, Germany

Abstract Improving our ability to cope with large risks is one of the key challenges for humankind in this century. This article outlines a research program in this perspective. Starting with a concrete example of a relatively small disaster, it questions simplistic ideas of rationality. It then proposes a fresh look at the concepts of probability and utility in the context of socio-ecological systems. This leads first to an emphasis on the problem of equilibrium selection, and then to a distinction between three kinds of resilience that matter both for theory and practice of risk management. They can be investigated by paying attention to the transitions into and out of actual disasters.

Keywords integrated risk governance, rationality, resilience, risk

1 Introduction

The twenty-first century is likely to see widespread economic growth along with great successes in overcoming poverty all around the world. At the same time, the century will be marked by growing risks in many areas, including natural, technological, and institutional disasters. Improving our ability to cope with such risks is one of the key challenges for humankind in this century. We cannot afford to make many mistakes with large-scale disasters; therefore it is essential to learn from relatively small cases as much as possible in view of the large, sometimes global risks we are faced with.

The starting point of this article is a concrete example that highlights strengths and weaknesses of current practices in dealing with risks. It shows how important it is to move beyond an idea of rationality as essentially applying general principles to particular situations. Whatever general statements can be made need to be embedded in a sense of practical judgement that is renewed with each new risk to be mastered. This, however, is not enough. There is a need to investigate the role of probabilities and preferences in the context of socio-ecological systems. This leads to the problem of equilibrium selection as a major focus for research. By looking at the entry- and exit-transitions of emergencies, it is possible to gain new insights about three kinds of resilience that characterize socio-ecological systems in the face of risks. Along these lines, I suggest elements of a research program that can enable scholars and practitioners to jointly improve existing practices in view of the risks we will need to tackle in our common future.

2 A Tragic Event

In summer 1989, some 150 people met in Berlin for a little gathering: an electronic dance music festival and parade. The organizers announced the event as a political manifestation and called it "Love Parade". In the following years the event was repeated, attracting more and more people. After ten years, the Love Parade was an annual mega-event with about a million participants. In 2007, the event moved from Berlin to the Ruhr area, the former industrial core of Germany. Meanwhile, the series had become the greatest set of dance events worldwide. It was a remarkably peaceful affair, a huge advertisement opportunity mobilizing major commercial interests, a gathering of millions of young people enjoying themselves in the name of love.

In the run-up to the 2010 Love Parade, to be held in the city of Duisburg in July 2010, massive concern was expressed on internet chats by people who wanted to attend. They criticized the fact that the whole event was to take place in an enclosed area, and that the only way in and out would be through a tunnel-large perhaps, but with concrete walls and no alternative pathway in case problems should arise. "I can't believe it! I see people dying," wrote one blogger. The police and the fire department voiced similar doubts and suggested a different setting to avoid critical risks.

A safety concept developed with the participation of world-class risk researchers at Duisburg University dispelled these doubts, and the original arrangement was preserved. On a sunny Saturday afternoon, a crowd of hundreds of thousands of people began to move through the tunnel towards the single entrance with the intent of enjoying hours of music and dance.

Three hours later more than 300 participants had been injured, 19 were dead (The Economist 2010). They came from Germany, but also from Australia, Bosnia, China, Italy, and other countries. In a situation of huge stress, some had tried to bypass the bottleneck between the tunnel and the entrance of the enclosed area by using stairs that run along a

* E-mail: Carlo.Jaeger@pik-potsdam.de

concrete wall. The stairs were blocked by a little fence, but it was easily overcome. A sense of hope led thousands of people to move towards the wall. For those already close to the wall, the physical pressure of the masses of people was huge, creating injuries and killing people by suffocation. Exhaustion and panic set in amongst many of the tens of thousands of people in and around the tunnel, while the news of the tragedy slowly spread amongst the hundreds of thousands who were already listening to the music.

3 Reflective Practitioners

The Duisburg tragedy confronts us with some of the key challenges of contemporary risk management. It happened at a public event in a rich country with a highly developed safety culture and sophisticated technological infrastructure, and it happened despite a safety concept developed with the help of the most advanced scientific knowledge on pedestrian flows and mass panic. Learning from such experiences is vital if we are to deal responsibly with the large-scale risks of the future (Jaeger et al. 2001).

Many important aspects of this particular accident are well-known from other disasters. They range from the financial pressure on the organizers to the mindless haste in which key preparations took place. It is necessary, however, to look at some underlying issues that relate to the idea of rationality and require fundamentally new research to be addressed.

The study of pedestrian flows, including the dynamics of mass panic, is one of the most interesting examples of how mathematical methods known from physics can be applied to human behavior with the help of computer simulations (Helbing et al. 2005; Schreckenberg and Selten 2004). These methods were used in the Duisburg case, and the experts involved did think about possible disasters ranging from panic triggered by misinformation to the possibility of a bomb being placed in the ominous tunnel. But the kind of accident that actually happened was not recognized as a possibility. After the event, a leading expert was quoted as saying: "The accident happened because some participants didn't follow the rules," by which he meant that they climbed structures—containers and a pylon—they were not supposed to nor allowed to climb. He claimed that the accident had happened because they fell down from those structures, a claim soon disproved by witness accounts and video recordings. The same expert had not visited the site before giving his advice to the organizers—he was confident that his know-how and toolkit were sufficient to tackle the questions he was being asked.

Simulation models of pedestrian flows are a special instance of multiagent models, and their track record in reproducing empirical data is among the best in the field of social science modelling. Moreover, it is evident that the management of pedestrian flows involving hundreds of thousands of people has much to gain from computer models that can help to assess the density, direction, and speed of these flows, both at the aggregate level and at the scale of single doors, stairs, and so on. But using such tools well requires more than technical expertise in the domain of mathematical models and their implementation on computers. It requires the kind of know-how that Schön (1983) described in his widely acclaimed, but still too widely neglected, work on reflective practitioners.

This is a know-how that treats the theoretical principles and analytical techniques so prominent in Western science as useful and indispensable, but also as meaningless without a context of practical problem solving that cannot be reduced to those principles and techniques. By this token, it is a know-how that transcends the idea of absolute rationality that Perrow (1984) identified as a root cause of increasing risks in our global society. The insights of contemporary science are not the bedrock of knowledge; they are sediments of professional practice. Each time they are used to solve a new problem, they need to be embedded in fresh conversations among reflective practitioners and people familiar with the specific problem.

Story-telling is a fundamental resource to be used in this spirit (Wilkinson 2009). Narratives provide the opportunity to blend generalized insights with unique events, and they can give a much richer sense of the possibilities generated by a concrete situation than the computer models currently available. In many situations such models are tremendously useful, often indispensable. However, applying simulation models to risky situations with which one does not have a familiarity that is much richer than those models, is a dangerous exercise—the pedestrian flows leading to the Duisburg tragedy are one example, the financial flows that led to the global financial crisis of 2008 another (Colander et al. 2009). How can we develop models and user interfaces that help professionals dealing with risks to become reflective practitioners, and never to forget this lesson, even in the face of the best conceivable models? This is the initial question for a research program to deal with the risks of the twenty-first century.

4 Probability and Utility

The modern approach to risk management began with insurance contracts for Italian ships in the fourteenth century. The insurance business gradually developed and was greatly expanded in the wake of the Great Fire of London of 1666. Since then, the practice of risk management, and especially of insurance, has been reflectively improved again and again with the help of mathematical probability theory.

Probability theory emerged out of reflections about the practice of gambling. They were triggered by the Chevalier de Méré, a French nobleman, in 1654, and were developed in an exchange of letters between famous French mathematicians Blaise Pascal and Pierre de Fermat (Devlin 2008). Their basic problem remains relevant today and can be framed as follows. Suppose you celebrate your birthday in a gambling salon and somebody offers you a choice between three gambles as a gift. In the first game, a coin is tossed three times, and if this yields heads three times, you get 100 Euros, otherwise nothing. In game two, a pair of dice is thrown two times, and if in at least one throw the sum of the two dice is 7, you get 100 Euros as well. And in the third game you are blindfolded and supposed to pick two balls out of an urn with 10 balls, 9 black and 1 red; if by chance you pick two balls of different colors, you again get 100 Euros. If you want to have the biggest chance of getting the 100 Euro prize, which game should you choose?

The problem can be described as follows: you have three possible actions; for each action, you are interested in a particular possibility, namely that you win the prize; and you want to maximize your chances to win that prize. The problem then can be written as:

$$\text{Max}_{[x]}\ 100 \text{ Euro} * \text{prob}\,(\,\textit{winning} \mid \mathbf{x}\,) \qquad \text{Eq. 1}$$
$$\text{s.t: } \mathbf{x} \in [1, 2, 3]$$

The expression "Max $_{[x]}$" indicates that the variable x shall be chosen so as to maximize the term following it. The expression "*winning* | **x**" indicates the situation where you win having chosen game x; "prob" is a function that indicates the probability of each possible outcome of each game. The challenge is to define that function.

In the first game, there are eight possible patterns of heads and tails. Writing 1 for heads and 0 for tails, they can be written as the binary numbers from zero to seven: 000, 001, . . . , 111. As these are all the possible patterns of heads or tails, the possibility that at least one of them will occur (that is, one of the set of all eight patterns) is assigned a probability of 1. If moreover one assumes that these patterns have the same probability—a key assumption that is far from trivial—then each one must have a probability of 1/8. This is the probability of winning. By the same kind of reasoning, the probability of winning in game two can be computed as 1/6 and in game three as 1/5. Now the "prob" function is defined for all three cases; clearly, the third game is the best and the solution to problem 1 is x = 3.

Since the days of Fermat and Pascal, probability theory has been greatly refined, as has the practice of risk management. A major step at the interface of these two developments occurred in 1738, when the Swiss mathematician Daniel Bernoulli published (in the journal of the St. Petersburg Academy of Sciences) his solution to a problem posed 25 years earlier by his cousin Nicholas Bernoulli. The problem, known as the St. Petersburg paradox, was how to evaluate a game where a coin is tossed until it yields tails, and where the prize is 2n ducats, with n the number of tosses so far. The idea captured in problem 1, where money is multiplied with probability, leads into difficulties here. The value of the prize increases more rapidly than the probability of the prize decreases, and this happens over an unlimited number of steps. As a result, the game seems to have an infinite value, but nobody would be willing to pay a large amount of money to be allowed to play it. Daniel Bernoulli addressed this difficulty by generalizing problems like in Eq. 1 into problems of the form:

$$\text{Max}_{[X]}\ \Sigma_{[C]}\ u\,(\,\mathbf{c}\,) * \text{prob}\,(\,\mathbf{c} \mid \mathbf{x}\,) \qquad \text{Eq. 2}$$
$$\text{s.t: } \mathbf{x} \in X,$$
$$c \in C$$

Here, C is the set of possible consequences of actions, X the set of possible actions, and u is a utility function representing the preferences over those consequences. Daniel Bernoulli argued that the utility function has to be strictly concave, displaying diminishing marginal utility and implying risk averse behavior. He saw that his framework was not only relevant to deal with the abstract possibility of infinite gains, but also in view of the very real possibility of large finite losses, including the risk of bankruptcy. As he wrote to Nicholas: "If only the Bernoullis, who lost so much when the Müllers went bankrupt, had paid attention to the very principles that I have established, they would probably not have lost as much" (see Jallais, Pradier, and Teira 2008, 49).

Daniel Bernoulli's idea of diminishing marginal utility became a cornerstone of economic theory as we know it, and his combination of probability and utility has deeply shaped the tradition of risk management. However, while probability theory made tremendous progress in the times following his breakthrough, the combination of probability and utility lay dormant until 1947. In that year, as an appendix to the second edition of their seminal work on game theory, von Neumann and Morgenstern (1947) published the first axiomatic structure combining probability and utility.

With that step, professional risk management had come of age. Dealing with risks by assessing—even if provisionally—the probability of different events as well as their relevance for the goals a decision-maker is pursuing became common practice. Rough and ready methods were used in this spirit, as well as highly sophisticated models, including those of pedestrian flows, in the risk assessment for the Duisburg love parade.

5 The Problem of Equilibrium Selection

There is a fundamental problem with this approach, however, a problem that led to the development of game theory by von Neumann and Morgenstern in the first place. In most practical situations the environment of a human agent includes other human agents—be they physical persons or various kinds of collective actors. Therefore, no agent can assess his or her own risks without assessing how the other agents assess theirs, and those other agents cannot do that without assessing how the first agent assesses his or her own risks. Clearly, there is a kind of circularity here. If agents frame their situations according to problem 2, they neglect the interdependencies between them. Therefore, rational as the approach may seem, it is bound to lead to major failures whenever these interdependencies matter. Major examples of this kind of failure are financial crises, as these usually are promoted by behavior that is quite advantageous for an agent as long as those interdependencies can be neglected, but becomes self-destructive otherwise.

von Neumann and Morgenstern addressed this problem by looking for situations where each agent knows the risk assessments by all agents, and where the resulting actions are optimal for each agent in the given situation. This pattern was analyzed in general terms by Nash (1950) and is now known as a Nash equilibrium. It has turned out to be tremendously useful in a large variety of situations, for example, the so-called prisoners' dilemma elaborated by Flood, Dresher, and Tucker in 1950 (see Flood 1952 and the discussion in Poundstone 1992). Basic aspects of mass panics can be analysed with the model of the prisoner's dilemma, too.

However, only very simple games have obvious Nash equilibria, and even then there may be many of them. Systems of interdependent agents with shared environments—socio-ecological systems (Young et al. 2006)—are characterized by dynamics that are much more complex than simple convergence towards an equilibrium state or persistence in such a state. Therefore, the question of how systems of interdependent agents evolve, whether they find equilibrium positions, and if so, how, becomes crucial for research on risks. This is especially important in view of global risks, as here the interdependence between agents plays a central role. Two important answers to this problem deserve our attention here. They center on the phenomena of conventions and of social norms.

A paradigmatic example of a convention is given by patterns of right-hand-side or left-hand-side driving. Research on the evolution of conventions has shown how conventions can emerge out of stochastic processes in socio-ecological systems, how such systems can persist in Nash equilibria established by convention, and how random events can lead to switches between such equilibria (Young 1993). The relevance of conventions for risk management is well illustrated by their role in financial markets (Wyart and Bouchaud 2003).

Once a convention is established, it tends to persist simply because it establishes a Nash equilibrium: in general, it is in the best interest of agents to follow the convention, as the example of right-hand-driving illustrates. However, consider the problem of precedence at intersections. The convention that the driver coming from the right (or the one from the left) has precedence only works if there is little traffic. Without conventions, the situation becomes what is known as a chicken game: if one driver dares to move while the others wait, the first one has an advantage; if all wait, they are all in trouble; and if all dare to move simultaneously, the trouble is greater still. A widely used solution to this coordination problem is the installation of traffic lights. They allow for a different kind of equilibria, so-called correlated equilibria (Aumann 1974). They have two important advantages over Nash equilibria: they are much easier to find, and they often allow for outcomes that are better than those from Nash equilibria. The latter point is well illustrated by the difference between a traffic light and a chicken game at intersections. Gintis (2009) calls the correlating device a choreographer and argues that this is a key role of social norms. In this perspective, norms develop through similar processes as conventions. However, they work in much more sophisticated ways, including the influence of moral arguments and sanctions.

For conventions and norms to work, each agent must know that the other agents orient their actions according to those conventions and sanctions. This is especially relevant in situations of crisis, where the validity of specific conventions and norms is in doubt. As far as an individual agent is concerned, a key tool for reflecting on his or her possibilities to navigate the risks that come with such situations is an algorithm updating probabilities based on Bayes's theorem.

That theorem was discovered around 1755 by Thomas Bayes (Swinburne 2002). It establishes a simple relation between two conditional probabilities. Consider probabilities over a set of events C (where not only the elements of C, but also its subsets as well as C itself are considered as events). Let A and B be two subsets of a set C and define conditional probabilities as follows:

$$\text{prob}(A \mid B) = \text{prob}(A \cap B) / \text{prob}(B) \qquad \text{Eq. 3}$$

Then it immediately follows that:

$$\text{prob}(A \mid B) = \text{prob}(B \mid A) * [\text{prob}(A) / \text{prob}(B)] \qquad \text{Eq. 4}$$

This is such a simple result that one may wonder why it has become one of the most influential concepts in statistics, decision theory, programs for machine learning, and more. And in fact it is not the theorem as such, but its relation to much more complex situations that makes it so important.

Consider, for example, a system that can assume several states, say Achilles who can be angry or satisfied, together with observable behaviors of that system, say Achilles shouting or smiling (the following analysis owes much to Feyerabend 1992). Achilles can shout and smile in both moods, but the probability of doing so changes with his mood. Moreover, his state can be a mixture of anger and satisfaction, again affecting his probability of shouting or smiling. Imagine you are Ulysses, who has the task to go into Achilles's tent and negotiate with him on behalf of the king. You have an idea about how angry he is, and you start talking with him, noticing his rare smiling as well as his frequent shouting. And as he starts shouting less and smiling more, you realize that his anger is fading away and that your chances of reaching an agreement are on the rise.

Let $\mathbf{o}_t$ be the observation made at time t and $\mathbf{s}_j$ the j-th possible state (with "S" the set of possible states). Writing "p" for probability one can then formulate the following equation for Bayesian updating of probabilities:

$$p(\mathbf{s}_j \mid \mathbf{o}_t) = p(\mathbf{s}_j \mid \mathbf{o}_{t-1}) * [p(\mathbf{o}_t \mid \mathbf{s}_j) / \Sigma_{[S]}\, p(\mathbf{o}_t \mid \mathbf{s}_j) * p(\mathbf{s}_j \mid \mathbf{o}_{t-1})] \qquad \text{Eq. 5}$$

It is clearly analogous to Eq. 4, but now it describes a dynamical process in which the probability assigned to state $\mathbf{s}_j$ evolves depending on observations $\mathbf{o}_t$.

This looks quite promising as a description of how Ulysses might perceive Achilles in the course of their conversation, but can the formula be proved? Of course not, unless much more assumptions are given. On the other hand, for systems that can be described with this kind of dynamics, more might be provable. For example, different observers who start with different assessments of the probabilities for the initial states—the so-called priors—may or may not converge on similar assessments of the actual states by sharing a sequence of observations (Hawthorne 1994). This is important, as the relevant systems might include not only Achilles but also other human beings, or quantum systems with their states and observables, or the earth crust in California with its inner tensions and intermittent earthquakes.

What happens if interdependent agents try to coordinate themselves by updating their probabilities according to an algorithm like Eq. 4? This depends on all sorts of specific conditions. But as Gintis (2009) has shown, they have a fair chance of stabilizing norms and conventions if they start from common priors, while they will find themselves in serious trouble without common priors. Therefore, solving coordination problems in socio-ecological systems requires processes by which common priors are established and maintained as part of common sense—not once and for all, but whenever a new coordination problem arises.

In the Duisburg case, common priors were taken for granted by the organisers, and for a while this assumption turned out to be reasonable. However, the assumption was misleading in the face of people who lost their trust that they would happily reach the concert area by following the instructions of the organisers. Not having foreseen this possibility remains the responsibility of the modellers. In a similar way, the possibility of trust breaking down on financial markets had not been foreseen by economic modellers before the global financial crisis of 2008.

6 Three Kinds of Resilience

Looking at risks in the context of socio-ecological systems offers a way out of the definitional maze that has grown around the concept of risk. Probability, utility, and game-theoretic equilibria—be they Nash or correlated—will all have their role to play, but the starting point is the difference between normal states of affairs and emergencies. This difference need not be introduced by an external observer. Rather, socio-ecological systems have grammars that enable the people involved in them to make that distinction. For somebody driving a car on the road, an accident usually is an emergency; for a car insurer, it is an event in the normal course of affairs. A child cutting a sibling with a knife at dinner is likely to trigger an emergency; a surgeon cutting the skin of a patient is performing a normal operation. If a global nuclear war should start, it would be an emergency for humankind as a whole.

Many emergencies are challenges that can be met without disrupting the system in question. These are familiar risks that can be represented in the probability-utility framework: events one prefers to avoid but also knows how to handle. In many cultures, floods, droughts, even earthquakes belong to this category—those cultures have learned to deal with these emergencies as one can learn to deal with various accidents, illnesses, and death. This may be called first-order resilience.

An important way in which socio-ecological systems develop first-order resilience is by developing subsystems that have the capability to treat as normal business what for the rest of the system is an emergency. The insurance industry is a major example of such a subsystem, but so are hospitals, armies, fire brigades, and many others. Governments assume this role in various circumstances, for example, when they bail out banks they consider too big to fail. In the face of global risks that exceed current coping capacities, developing new institutions that maintain first-order resilience in the face of the corresponding emergencies is a major task (see Shiller 2003, for a discussion of global financial risks).

By distinguishing between normal states of affairs and emergencies, the grammars of socio-ecological systems define entry transitions into and exit transitions out of emergencies. When first-order resilience is established via a specialized subsystem, then the entry transition moves this subsystem into action, while with the exit transition it goes into standby mode again. Improving the speed of those transitions often is an important way of improving first-order resilience. More generally speaking, studying the entry and exit transitions into emergencies is a key task, if we are to enhance our capability to deal with future risks (Integrated Risk Governance Project 2010, building on Kasperson et al. 1988).

The safety concept for the Duisburg Love Parade took first-order resilience for granted, as did the pedestrian flow models used to design and assess that concept. The financial regulations in place before the 2008 financial crisis took first-order resilience for granted in the much larger domain of global financial markets—as did the models used to design and assess those regulations.

Larger perturbations, however, can destroy the coordination capacity of a given socio-ecological system, leading to a very different kind of situation. This is where the capacity to manage the unexpected (Weick and Sutcliffe 2001) becomes critical. First-order resilience is based on patterns of conventions and norms that keep solving coordination problems in the face of perturbations. But no such pattern can be maintained in the face of arbitrary perturbations. If the conditions under which a socio-ecological system operates change—perhaps as a result of the dynamics of the system itself—existing patterns of coordination may break down. The capability to handle this breakdown until the system can switch back into its normal way of operation may be called second-order resilience.

If a socio-ecological system operates under stable and foreseeable conditions, first-order resilience is usually all that matters. If the conditions are more turbulent, however, second-order resilience becomes much more important. So-called high-reliability organizations are characterized by high second-order resilience, while they may actually have weaker first-order resilience then their average competitors.

Because in the entry transition to an unfamiliar emergency key coordination mechanisms have broken down, the shared environment becomes particularly important for the different agents, both because it presents new challenges and because it may offer opportunities to improvise new coordination patterns. In the Duisburg accident, the absence of alternatives to the tunnel as well as the presence of obstacles that could be surmounted on the way forward beyond the tunnel played key roles. In environmental disasters like floods or earthquakes, the shape of terrain, buildings, and debris plays a similar role. More generally, the point of view of socio-ecological systems is essential to understand the dynamics of norms and conventions in the face of risk (Walker et al. 2006). The environment here is not to be understood as some pristine natural landscape, but as the material conditions of human interaction, including the built environment and all sorts of technological devices (Zha 2006).

With regard to the exit transition, a fundamental distinction arises: it may lead back to the previous state of normality or to a new one. The latter case is especially interesting if the system in question has found ways to reduce its vulnerability to the kind of perturbation it went through. This may be called third-order resilience. While first-order resilience is mainly a matter of robust patterns of norms and conventions, second-order resilience depends critically on the ability to improvise; third-order resilience in turn depends on the capability of the socio-ecological system to find a creative answer to the disruption it has experienced.

7 Conclusion

Human beings trying to cope with the risk of emergencies as well as with their occurrence are always embedded into socio-ecological systems. In fact, the research program sketched here may be called one of risk governance in socio-ecological systems.

In this phrase, the proposition "in", rather than "of", is consciously chosen. The latter can too easily foster an attitude where management is seen as controlling the behavior of automata hardly recognizable as human beings. In the case of the Duisburg Love Parade, this meant that the autonomy and creativity of the participants was treated only as a possible disturbance and cause of disaster, not as a resource that might have helped to cope with the challenge of coordinating several hundred thousand people. Similarly, in the case of an earthquake, there is a fundamental difference between an approach that fosters the capability of people in the disaster area to take action and solve problems and an approach that treats them as objects of an intervention they are supposed to accept and welcome, but not to shape.

With this background, three elements of a research program on risk governance in socio-ecological systems can be identified:

(1) In the face of future risks, how can we use story-telling, computer models, and other tools to support and foster a culture of reflective practitioners?
(2) What features of entry-transitions into emergencies can be identified so as to enhance first-order resilience of socio-ecological systems?
(3) What features of entry- and exit-transitions can be identified so as to enhance higher-order resilience?

It is impossible and not necessary to know in advance with great detail where this research program will lead. But it is clear that the intertwined themes of risk, rationality, and resilience are so pervasive in today's global society that significant insights can be expected with regard to an unusually wide array of problems. And it is also clear that in a culture of reflective practitioners such a research program cannot evolve first in academia and then be transferred to a domain of "applications." What is required are patient inquiries structured by a continuous dialogue between researchers and practitioners, where advances in research alternate with improvements in practice.

This dialogue will refer to risks and emergencies at different scales, ranging from a fire in a village to sea level rise around the globe, from a technical failure in a small business to a breakdown of global financial markets. The issue of sustainability looms large in the relevant inquiries, in particular the tension between the need for economic growth and the need to avoid environmental disasters. The research program proposed in this article shall help to address these challenges.

References

Aumann, R. 1974. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics* 1 (1): 67–96.

Colander, D., M. Goldberg, A. Haas, K. Juselius, A. Kirman, T. Lux, and B. Sloth. 2009. The Financial Crisis and the Systemic Failure of Academic Economics. *Critical Review* 21 (2&3): 249–67.

Devlin, K. 2008. *The Unfinished Game: Pascal, Fermat, and the 17th-Century Letter that Made the World Modern.* New York: Basic Books.

The Economist. 2010. *Death in Duisburg*. http://www.economist.com/blogs/newsbook/2010/07/love_parade.

Feyerabend, P. 1992. Erkenntnis ohne Theorie. Vom Nutzen der Abstraktion und vom Recht des Besonderen. *Lettre International* (L 16): 66–71.

Flood, M. M. 1952. *Some Experimental Games*. Santa Monica, CA: Rand Corporation.

Gintis, H. 2009. *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton: Princeton University Press.

Hawthorne, J. 1994. On the Nature of Bayesian Convergence. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, vol 1, 241–49.

Helbing, D., L. Buzna, A. Johansson, and T. Werner. 2005. Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions. *Transportation Science* 39 (1): 1–24.

Integrated Risk Governance Project. 2010. *IRG-Project Science Plan*. Potsdam-Beijing. http://www.irg-project.org/index.php?eID=tx_nawsecuredl&u=0&file=fileadmin/publications/PUBLIC/ABOUT_IRG-Project/irg-project_SP_Mar10_final__wf_small.pdf&t=1282858339&hash=66b9951889b248eefd3322694520bbb7.

Jaeger, C., O. Renn, E. A. Rosa, and T. Webler. 2001. *Risk, Uncertainty, and Rational Action*. London: Earthscan.

Jallais, S., P.-C. Pradier, and D. Teira. 2008. Facts, Norms and Expected Utility Functions. *History of the Human Sciences* 21 (2): 45–62.

Kasperson, R. E., O. Renn, P. Slovic, H. S. Brown, J. Emel, R. Goble, J. X. Kasperson, and S. Ratick. 1988. The Social Amplification of Risk. A Conceptual Framework. *Risk Analysis* 8 (2): 177–87.

Nash, J. 1950. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences* 36 (1): 48–49.

von Neumann, J., and O. Morgenstern. 1947. *Theory of Games and Economic Behavior (2nd edition)*. Princeton: Princeton University Press.

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

Poundstone, W. 1992. *Prisoner's Dilemma: John von Neumann, Game Theory and the Puzzle of the Bomb*. New York: Doubleday.

Schön, D. 1983. *The Reflective Practitioner: How Professionals Think in Action*. London: Temple Smith.

Schreckenberg, M., and R. Selten, eds. 2004. *Human Behaviour and Traffic Networks*. Berlin: Springer.

Shiller, R. J. 2003. *The New Financial Order: Risk in the 21st Century*. Princeton: Princeton University Press.

Swinburne, R., ed. 2002. *Bayes's Theorem*. Oxford: Oxford University Press.

Walker, B. H., J. M. Anderies, A. P. Kinzig, P. Ryan, eds. 2006. Exploring Resilience in Socio-Ecological Systems. *Ecology and Society* 11 (1): 12.

Weick, K. E., and K. M. Sutcliffe. 2001. *Managing the Unexpected—Assuring High Performance in an Age of Complexity*. San Francisco: Jossey-Bass.

Wilkinson, A. 2009. In Search of Theory. *Journal of Futures Studies* 13 (3): 107–14.

Wyart, M., and J.-P. Bouchaud. 2003. *Self Referential Behaviour, Overreaction and Conventions in Financial Markets*. http://www.cfm.fr/papers/Wyart2.pdf.

Young, H. P. 1993. The Evolution of Conventions. *Econometrica* 61 (1): 57–84.

Young, O. R., F. Berhhout, G. C. Gallopin, M. A. Janssen, E. Ostrom, and S. van der Leeuw. 2006. The Globalization of Socio-Ecological Systems: An Agenda for Scientific Research. *Global Environmental Change* 16 (3): 304–16.

Zha, H. 2006. Interactive Technologies and Sociotechnical Systems. In *Proceedings of the 12th International Conference, International Society on Virtual Systems and Multimedia*. October 18–20, 2006, Xi'an, China.