

Risks and Games: Intelligent Actors and Fallible Systems

Elisabeth Paté-Cornell¹
Stanford University, Stanford, CA 94305

Copyright © 2009 by Elisabeth Paté-Cornell. Published and used by MIT ESD and CESUN
with permission.

Abstract

Failure risks that involve intelligent actors depend on their moves and on the effects of these moves on a system or situation. This paper examines two cases: the risks of failure of a technical system and that of a stabilizing policy in the face of an insurgency. Two models of games and risks are presented here to support a decision in each of these situations. The first case is that of a principal (manager) who sets incentives and penalties for a subordinate (agent), who in turn, manages the development of a technical system. What the agent does to meet the schedule constraints to which he is subjected can affect the performance of that technical system later in operation, especially if he decides to cut corners to satisfy the manager's requirements. The second case is that of a government trying to keep control of a country in the face of an insurgency whose leader's objective is to seize power. The government leader allocates his resources between short-term investments (trying to prevent immediate damage) and long-term ones meant to address the fundamental problems that fuel the insurgency. The model presented here is a simulation of that game as a set of alternate moves (government's and insurgents') over a number of time periods. The risk is that of being below a threshold of stability at the end of the simulation. In both cases, the objective is to present a tool in support of a decision. This paper describes the structure of the fundamental models and some illustrative results.

Key words: risk analysis, game, simulation, probabilistic risk analysis, engineering system

1. Games and risks of system failures

Many failure risks stem from interactions among actors and on their effects on the performance of the systems involved (technical, political, financial, etc.). The objective of this paper is to present two different approaches to mathematical models, based both on risk analysis and games analysis. These models are tools, designed to provide decision support for one of the actors (here: a manager or a government leader). The modeling problem is first to describe the interactions among actors: what do people know about their options and the consequences of their actions, what are they seeking and what are they rewarded for. The second is to

¹ Burt and Deedee McMurtry Professor and Chair, Management Science and Engineering, Stanford University, Terman 340, Stanford, CA 94305

characterize the influence of the various actors on the system's performance and resulting failure risk (here of a technical system and of a country stabilization policy).

Game theory offers an analytical framework that assumes, in general, that both sides act rationally; they may or may not have common knowledge of what the other side knows but they generally understand each other's preferences. Game theory relies on a fundamental principle of rationality, and the concept of Nash equilibrium (Von Neumann and Morgenstern, 1947; Nash, 1950; Harsanyi, 1959; Gibbons, 1992). Yet, the kind of game analysis that is useful in risk analysis does not necessarily assume that level of rationality and information. In some cases, the analysis can be done in a bounded-rationality framework, in which the main side (principal, government) simply assumes that the probability that the other side chooses an option is proportional to the expected utility of the adversary for that option among others (Paté-Cornell and Guikema, 2002). In what follows, it is assumed that both sides are rational in the von Neumann sense². Factors that are not necessarily known by both sides are the state of the system, the state of information of the adversary, and the options that the other side is considering.

Choosing an optimal course of action for the principal depends on his anticipation of the effects of this choice of incentives on the agent's actions and in turn, on the performance of the system (here: a technical system in one instance, the stability of a country in the other). There may not be statistical data sets that allow direct assessment of the failure risks. First the data may not exist in cases where failures are rare events. Second, existing statistics may not be relevant when the situation has changed and the existing data do not represent the current risk. Assessing the risk of rare events that result from conjunctions of adverse factors ("perfect storms") can be done by systems analysis including uncertainties and dependences.

Game theory has been a tool of policy analysis for a long time. Recently, there has been renewed interest in the use of game theory (or game analysis) in the context of risk analysis (Hausken, 2002; Bier et al., 2005). This paper presents two ways of assessing the risk and two different approaches to the analysis of a game in that perspective (Paté-Cornell, 2008).

The first case is a classic principal-agent problem (Milgrom and Roberts, 1992), in which the principal sets incentives and information flow for the agent so as to reach his objectives³. The agent is in charge of a system's development and his actions, in turn, can affect the performance of the physical system in the operation phase. In the problem considered below, the (externally set) schedule constraints have led the project to being late. The agent is then tempted to take various shortcuts. His choice depends on his risk attitude, the incentives set by his manager, how late he is, and his understanding of the effects of these shortcuts on the system's capacity. The problem considered here is that of the manager who chooses whether or not to monitor the actions of the agent and sets the penalties to the agent for being late, caught cutting corners, or causing a failure of the system in operation (Paté-Cornell and Garber, 2006; Garber, 2007). The model shows that if the manager does not consider carefully the implications of the incentives that he sets, the agent may decide (rationally) to take shortcuts and to hide it, even though the

² This, of course, does not mean that they share common values, only internal consistency of their preferences at a given time.

³ This part of the paper is based in the PhD thesis of Russ Garber (2007) under the supervision of the author.

manager might prefer to know that there is a schedule problem and how it is going to be resolved.

The second case presented here is that of a government facing an insurgency⁴. The decision of the government leader is how to allocate resources between short-term measures that target immediate safety problems, and long-term measures that address fundamental issues that fuel the insurgency (Kucik, 2007). This part of the paper is illustrated by a real case in the Philippines, based on government data. The risk is that in a given time frame, the country is still unstable. This work follows a study of the risks posed by terrorist groups in which the question was to assess the probabilities of different threats of terrorist attacks on the US (Paté-Cornell and Guikema, 2002).

There are several interesting features in the models presented here. The first case demonstrates the use of a probabilistic risk analysis for the technical system in order to assess the effects of various shortcuts on the failure risk. More importantly perhaps, the model allows a manager to understand the implications of accepting (and in some cases, setting) too tight a deadline for the completion of a system's development. The model allows assessing the "shadow price"⁵ of the schedule constraint through its effects on the agent decisions and consequently, the failure risk in operation. The second model illustrates the measurement of the risk through a simulation that yields the probability that the level of stability of a country (or the satisfaction level with a system's performance) is below a specified threshold after a certain time.

These features are illustrated in what follows in the context of the design of decision support tools and models, meant to be applied to specific cases. Therefore, the formulations presented here have to be adapted to each situation.

2. Engineered systems, shortcuts and failure risks: a principal-agent model

Consider the hypothetical case of a construction project in a seismic area. The workers are supposed to secure a shear wall by nailing its panels to a frame with a given number of nails. They are late and decide to reduce that number of nails. Years later, an earthquake occurs and the structure fails. The manager is responsible and liable for the damage and so are the workers. How potentially damaging (in retrospect) was the decision to take shortcuts in this case? What should have been the manager's decision (incentives, inspection) at the time of the building construction? The result of the workers action was an increase in its failure probability over its lifecycle, which could have been anticipated and accounted for in a decision support of the manager's choices.

⁴ This part of the paper is based on the PhD thesis of Paul Kucik (2007) under the supervision of the author.

⁵ The shadow price of a constraint is the increase of the objective function –here the decrease in the failure probability- that is associated with relaxing the constraint by one unit. Note that in some cases, the deadline for a project development may not be flexible; it could be determined for instance by the weather, the position of the planets or other unchangeable realities. In these cases the manager may or may not have the option to decline to undertake the project.

The classic tool to assess the potential effects of these types of shortcuts –and thus of some deadlines- is a probabilistic risk analysis (PRA) of the technical system, in this case, the building under uncertain seismic loads. The PRA method is based on the decomposition of a system into its subsystems and components, identification of the failure modes and computation of the overall failure probability as a function of those of these different elements (Henley and Kumamoto, 1996; Paté-Cornell, 2007). One objective is to best allocate scarce resources to reduce the failure risk, including money, time and attention span.

A source of failure that has received much attention in the risk analysis literature is human errors, and their dependence on management decisions that include incentives, information available to the agents and incentive structure (see for instance an application to the shuttle tiles by Paté-Cornell and Fischbeck, 1993). One general model designed to link system failure risk and human and management errors is the SAM model (system-actions-management). SAM starts with an analysis of the failure probability of the physical system as a function of the failure probabilities of its components through a functional analysis, a systematic identification of the failure scenarios and a probabilistic computation involving dependencies among events. For each variable of the model, SAM then requires identification of the potential decisions and actions taken by immediate operators or agents that could affect the probabilities of different realizations. Finally at the third level, SAM tracks back the sources of these actions in management decisions, including incentives, operational rules and resources available (Murphy and Paté-Cornell, 1996).

2.A The tool: a principal-agent model

The development of a complex system considered here is an application of the SAM model, in which the constraints have been set (often externally) and agreed upon by the manager, and the development of a project is late by L time units⁶. The manager may not know it. The agent in charge of the project is facing a penalty if he does not meet the deadline and may be contemplating shortcuts. His decision is thus to reveal that he is late and incur the cost of it, or to try to catch up by cutting corners. This choice depends on his risk attitude, the chances that he might be caught if the principal monitors his work, and the probability –as he estimates it-- that the system fails because of the shortcuts that he might take. If he cuts corners, the question becomes: in which components and how, for example, by substituting material or skipping some tests⁷. It is assumed here that the agent knows the system well enough to estimate the corresponding reduction of system capacity and the increase of the failure probability in operation. He can thus choose shortcuts that allow him to minimize the corresponding increase of failure risk while catching up (Garber, *ibid.*).

The game is thus a one-step game. Decision analysis allows identification of the optimal choice on each side. The agent minimizes his expected disutility for the costs involved (admitting being late, being caught cutting corners and causing a failure) given the delay.

⁶ In some cases, the deadline may be set externally, for example, by a change of the weather, the positions of the planets or other unchangeable realities. In these cases, the manager has the option to decline to undertake the project. In others, he may have some influence on the constraint itself.

⁷ In the design phase, shortcuts may actually increase the system's capacity at the cost of greater expense, for example, because design engineers decide to take a large safety margin rather than analyzing the effects of smaller ones.

The manager faces resource constraints in terms of money, time and attention span. He decides at the beginning of the development phase, whether or not to monitor the agent to make sure that all planned steps are taken (no short cuts), and what penalties to impose on the agent for being late, being caught cutting corners and causing failures later. But he does not necessarily know during the development phase that the project is late.

The interaction between the manager and the employee in charge of the project can be represented by a principal-agent model. Both sides identify their best option given the possibility of a delay for the principal, of a known delay for the agent. The model can then be used for a more general purpose, which is to assess the shadow cost of the schedule constraint, leading the manager to ask for less stringent deadlines or to simply decline to manage the project if the failure risk associated with the constraints is too high.

2.B The model

The components of the technical system are identified by the double index (k,i) of the subsystem k to which they belong and the component i inside the subsystem. The agent knows whether or not he is being monitored and thinks that there is a probability $p(C|\{(k,i)\})$ that he will be caught if he is and if he takes a shortcut in component (k,i). $G(k,i)$ is the amount of time that the agent can regain by cutting corners in component (k,i). \underline{U}_A is the disutility function of the agent, CC the cost to him of being caught cutting corners (event C), CL of admitting being late, and CF' the (delayed) cost to the agent of a failure caused by his short cuts. α is the discount factor of the agent. The cost of failure to him is thus the failure penalty set by the principal, discounted the agent's rate to account for the delay until operation. The following set of equations represents the simultaneous optimization for the principal and the agent.

For the agent, the problem is to minimize his expected disutility over all possibilities including shortcuts given that he knows whether or not he is being monitored. First, he assesses the potential effects of cutting corners (while minimizing the probability of failure and getting back on time); second, he compares the result to his expected disutility of the cost of admitting lateness; he then makes a decision of what to do. If he knows that he being monitored:

$$\text{Min}_{\{(k,i), (0,0)\}} p(C|\{(k,i)\}) \times \underline{U}_A(CC+CL) + (1-p(C|\{(k,i)\})) \times p[F'(\{p_{k,i}\}, \{p'_{k,i}\})] \times \underline{U}_A(\alpha CF')$$

$$\text{subject to: } \sum_{\{(k,i) \neq (0,0)\}} g(k,i) \geq L$$

If he is not being monitored:

$$\text{Min}_{\{(k,i), (0,0)\}} p[F'(\{p_{k,i}\}, \{p'_{k,i}\})] \times \underline{U}_A(\alpha CF')$$

$$\text{subject to: } \sum_{\{(k,i) \neq (0,0)\}} g(k,i) \geq L$$

In both cases, one obtains first, a local optimum $\{(k,i)\}_{SC}^{**}$ (optimum choice of shortcuts if the agent decides cut corners), then compares it to admitting lateness. He then chooses the option that corresponds to $\text{Min} [\underline{U}_A(\{(k,i)\}_{SC}^{**}), \underline{U}_A(CL)]$. Therefore:

If $\underline{U}_A(\{(k,i)\}_{SC}) < \underline{U}_A(CL)$, the global optimum is $\{(k,i)\}^* = \{(k,i)\}_{SC}^{**}$ (selective shortcuts)

If $\underline{U}_A(\{(k,i)\}_{SC}) > \underline{U}_A(CL)$, the global optimum is $\{(k,i)\}^* = \{0\}$ (no short cut; admit being late).

For the principal, the problem is to minimize his expected disutility \underline{U}_P for the decision to monitor (M) or not, the possibility that the agent is taking shortcuts, and for the costs and penalties that he sets, knowing that they may discourage hidden shortcuts but also affect his ability to hire someone in the future⁸. The objective is thus:

$$\text{Min}_{M,CL,CC,CF} \text{EU}_P[\text{conseq. of scenarios} \mid \text{set of options } \{M, CL, CC, \text{ and } CF\}]$$

The scenarios of this equation involve whether the project is late, the agent is caught if he takes shortcuts and he is monitored, and if not, whether undetected shortcuts cause a failure in operations. Solving this set of two equation yields the optimal joint decision of both principal and agent. The algorithm is to solve the problem first for the agent as a function of the principal's decision, then for the principal, based on his assessment of the probability that the project could be late, and anticipating what the agent would do in each case⁹.




Figure 1: A toy model system illustration: functional diagram

2.C A toy model illustration

Consider the system represented in Figure 1. Assume that the agent's subsystem development is late by 10 days and that he knows that he can gain time by taking shortcuts in each of the components, for example:

$$g_{11} = 5 \quad g_{21} = 3, \quad g_{22} = 3 \quad g_{31} = 4$$

He also knows the probabilities of failure (e.g., in the first year of operations) of each of the components (conditional probabilities for the redundant components) with and without shortcuts.

Without shortcuts:	$p_{11} = 0.01$	$p_{21} = 0.001$	$p_{22 21} = 0.9$	$p_{31} = 0.01$
With shortcuts:	$p'_{11} = 0.1$	$p'_{21} = 0.01$	$p'_{22 21} = 0.95$	$p'_{31} = 0.1$

⁸ Note again that he does not know whether the project is/will be late when he makes these decisions.

⁹ This paper presents a simplified version of a more complex model (Garber, 2007). In particular it is assumed that all decisions are made at the same time. Operation occurs later and spans a longer time interval. To simplify the model it is assumed here that the system if it is to fail will fail when operation starts. Discounting thus applies only to the costs that the agent may incur for causing a failure at the beginning of operations by taking shortcuts in development.

Based on these illustrative data, the probability of system's failure without shortcuts is 0.0209 or about 2%. If the agent decides to cut corners, the best that he can do to minimize the overall failure probability is:

$\{(1,1), (2,1), (2,2)\}$ $p(F') = 0.11$ $p(F) = 0.12$ (F can also be caused by failure of 3.1)
 $\{(2,1), (2,2), (3,1)\}$ $p(F') = 0.11$ $p(F) = 0.12$ (F can also be caused by failure of 1.1)

The probability of failure has increased by a factor of about five¹⁰; but it could increase even more if he chose:

$\{(1,1), (2.1), (3.1)\}$ or
 $\{(1,1), (2.2), (3.1)\}$,
 in which cases $p(F)$ would become 0.209 and 0.19 respectively, roughly 20%.

The agent thinks that the principal is monitoring his actions and that, for each of these options, the probability of being caught is 0.3. The principal, if he decides to monitor operations, has set the penalties so that the disutilities to the agent are¹¹.

$$\underline{U}_A(CL) = 0.5, \quad \underline{U}_A(CC + CL) = 0.8, \quad \underline{U}_A(\alpha CF') = 0.9.$$

The disutilities for the principal are¹²:

$$\begin{aligned} \underline{U}_P(CM) &= 0.2, & \underline{U}_P(CL) &= 0.8, & \underline{U}_P(CM+CL) &= 0.85, \\ \underline{U}_P(CF) &= 0.9, & \underline{U}_P(CF+CM) &= 0.95, & \underline{U}_P(CM+CL+CF) &= 1 \end{aligned}$$

In this particular case, because of the high cost to the agent of admitting that he is late, his best option is to cut corners whether he is being monitored or not, either in components $\{(1,1), (2,1), (2,2)\}$ or $\{(2.1), (2.2), (3.1)\}$ (he is indifferent). In both cases, the portion of the failure risk that *cannot* be attributed to shortcuts, i.e., is the portion that is due to the failure of (1,1) or (3,1) (whichever is not affected by shortcuts). For the principal, the best option given his preferences and what he knows of the agent is to not monitor the project closely.

This is an example where the optimum is not what one would expect nor perhaps, want to see. The manager is better off "not wanting to know" and the agent is better off taking shortcuts and hiding it in order not to be punished for being late. The manager may re-consider (and decrease) the penalty to the agent for admitting being late, find ways to monitor the project at a lower cost, and increase the penalty to the agent for cutting corners to the extent that it is feasible.

The model's results can also be useful to whoever, at or above the manager's level, is setting the constraints assuming that there is any flexibility. Assume that the deadlines have no effect on the agent's diligence and thus the lateness of the project. If the time constraint were relaxed by 1, 2, or 3 days, the project would then be late only by 9, 8, or 7 days instead of 10. Yet, the optimum in the short cuts would not change as long as the agent wants to catch up for the full

¹⁰ Note that in the design phase, shortcuts may actually increase the system's capacity at the cost of greater expense.

¹¹ The costs are the inverses of the disutilities.

¹² The principal's disutility for failure does not depend on whether shortcuts are involved.

delay. But there would be a step decrease in the risk if the time constraint were relaxed by 4 days because the project would only be late by 6 days. In that case, if the agent chooses to take shortcuts, he can do so in components (2,1) and (2,2) and the probability of failure would then increase only to 0.03 (from 0.021 without shortcuts). That is because the probability of failure of sub-system 2 is low compared to that of subsystems 1 and 3, and would increase from 10^{-3} to 10^{-2} . The risk would thus be much smaller than in the case where the project is late by 10 days (0.03 instead of 0.11), because no shortcuts in components (1,1) or (3,1) would have to be considered¹³.

In this illustrative case, the model shows that the benefit of relaxing the time constraint does not materialize until the deadline is extended by four days. But if that happens, the failure probability is then reduced by a factor of about four.

3. Counter-insurgency policy analysis: a game simulation model

One alternative to the principal-agent model in which two expected utility functions are optimized in parallel is a simulation of the game. This is what was done to estimate counter-insurgency policies and support a government's sequence of decisions over time (Kucik, 2007). As mentioned above, the problem for a government facing such a situation, is to allocate its resources between the short-term needs for immediate security and long-term expenditures to address the basic problems that fuel the insurgency. Allocations at each time period are assumed here to be irreversible¹⁴ and cumulative. Their optimization implicitly relies on a discount rate of immediate versus long-term expenditures and benefits.

The approach is dynamic and based on the simulation of an alternate game, one step at a time for a number of periods (in the illustration presented further, 36 months). Each side at each step observes what the other side has done in the last time period, and the effect of that past move. The next player then considers the expected result of different possible options and optimizes his choice for the next period. The model therefore involves a sequential decision analysis (Raiffa, 1968), in which the (rational) decisions on each side depend on the opponent's last move, and are represented by two inter-dependent influence diagrams (Hong and Apostolakis, 1992). In this simple model, there is no assumption of memory (in the sense of learning and updating), nor explicit consideration of the effect of one's actions on the future moves of the adversary¹⁵. The preferences of each side are represented by a multi-attribute utility function understanding rationality does not imply shared values. It is also assumed that the government knows the insurgents' state of information and preferences. The insurgents, however, do not know how the government has allocated its resources and can only observe the outcomes.

¹³ Note that the constraint set at the beginning of the system's development, can also increase the effort that the agent puts into the project development, thus reducing the potential delay (Garber, 2007).

¹⁴ In a more sophisticated model, one may want to consider the possibility that the government changes its resource allocation as events unfold, and may not have to follow up on earlier commitments.

¹⁵ One can consider that players on each side implicitly accounts for the longer-term effects of their moves by considering for example in their utility function the long-term implication of indiscriminate killing and its effect on civilian population support.

The two linked influence diagrams at the core of the model represent the decision problem for each side at each time period, based on the outcomes and observations of the previous period. The probabilities and preferences of the insurgents' side are assessed by the government based on intelligence information¹⁶. On the government's side, the options include first, counter-insurgency operations, such as increasing the number of troops on the ground and improving their equipment, and second, long-term investments for example, in food supply, medical care, education and infrastructure. The preferences are encoded in a multi-attribute utility function that includes five attributes:

- The economic situation measured as the percentage of the population below the poverty line
- The political, social and religious situations measured as the percentage of the population that is dissatisfied with non-economic conditions
- The population's security level: it was decided to measure it as the number of actionable human intelligence tips received by the government in the last three months divided by the number of insurgents attack initiatives
- Counter-insurgency combat operations measured as the number of combat engagements initiated by government forces
- Precision in counter-insurgency combat operations measured by the fraction of people killed by government forces who actually are insurgents.

On the insurgents' side, the options at each time period are to launch an attack, organize civil disobedience or seek a lasting peace. The insurgents' preferences are characterized by a multi-attribute utility function similar to that of the government. Both utilities are a simple linear function of a measure of each of the attributes. This formulation assumes independence among the attributes, which is only a first approximation and can be refined by introducing cross-terms. At each time period, each side (in turn) observes the past move of the other, assesses the probabilities of the different scenarios, and makes a decision based on utility (including a discount rate). The two influence diagrams are linked. The outcome of the insurgents' actions becomes input for the government's decision, and the outcome of the government's actions then becomes an input for the insurgents.

In the end, the result of the simulation is the probability of meeting a threshold of political and economic stability given the strategy of the government. This risk estimation allows, for instance, comparison of policies aimed primarily at nation building and policies focused essentially on immediate control of the country, and mixed policies.

3.A Illustration

This model was used to assess after the fact the probability that the policy that was followed in the Philippines to combat an Islamic insurgency in its Southern islands would lead to a certain threshold of stability for the local populations. The time frame was the 36 months period between 2000 and 2003, and the model was run for a specific allocation of government funds. Figure 2 shows a few paths in that simulation to describe how the results were used to compute a probability distribution for the country's stability level in the end.

¹⁶ Preferences can change, for example, with a change of leadership.

The results of this illustration show that simulation would have predicted that the policy would not be particularly successful and the government would end up pretty much where it had started, but not lose ground. The starting point (observed) was characterized by a utility value of 0.203, and the simulation yielded a distribution of utilities with a median of 0.206, and 5th percentile of 0.171, and a 95th percentile of 0.234. The reality as observed in 2003 was actually close to the computed median.

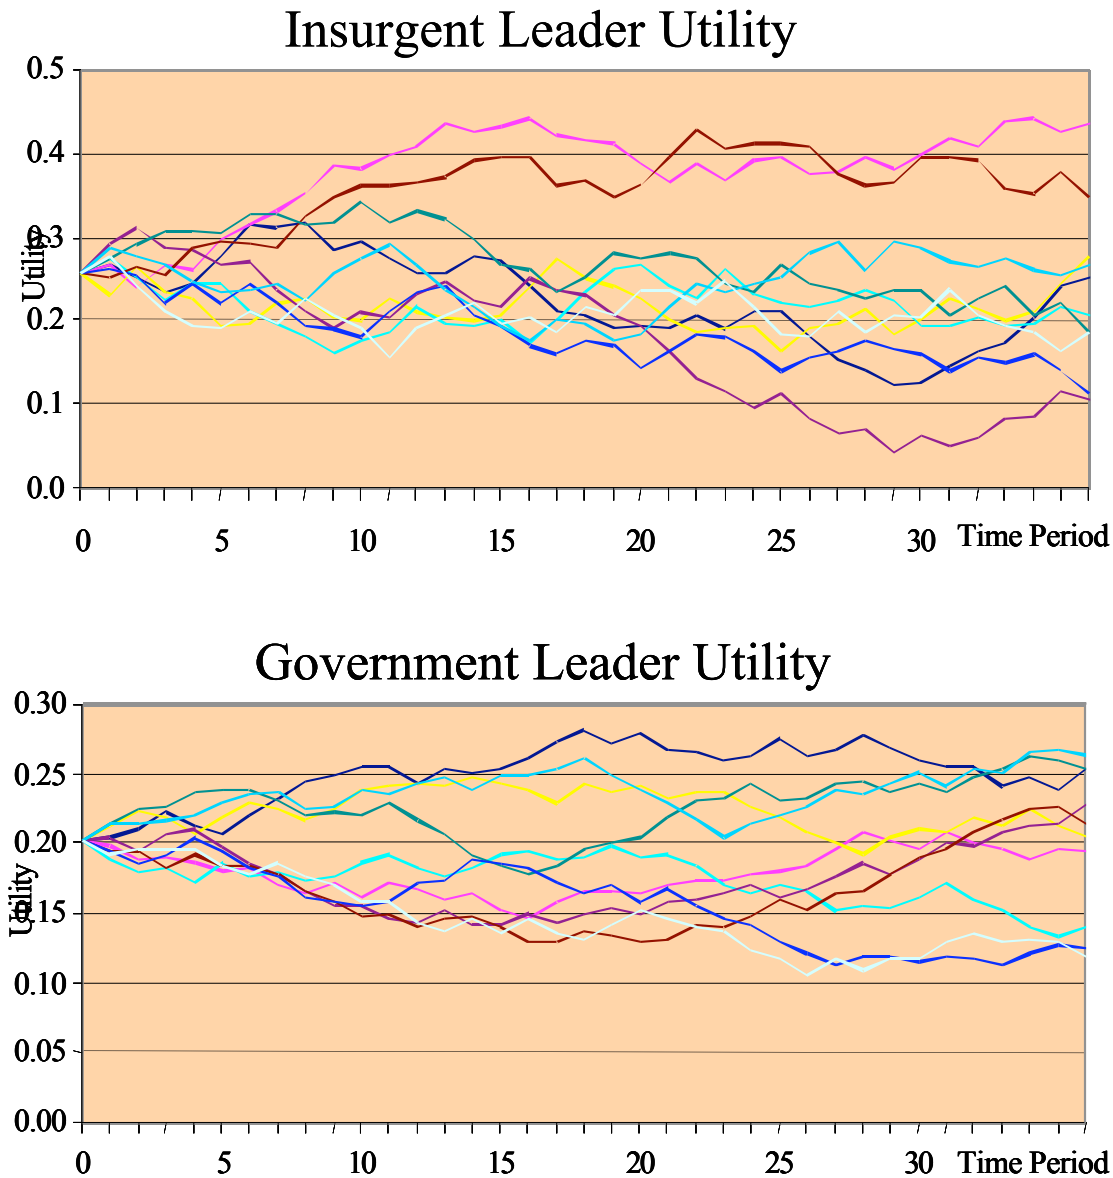


Figure 2: Illustration of the simulation results of the case of an insurgency in the Philippines (Source: Kucik, 2007)

The benefit of this model and its results is that it can help a government decision maker weigh the costs and the benefits of different types of options. In particular, it provides a balance to a

natural inclination to focus on short-term benefits, and can help consider the relative advantages of long-term goals. Also, measuring the risk as the probability of being at a given time under a specific level of satisfaction is useful in complex cases, in which the risk can be measured in terms of a desirable feature (here the stability of the country). This model, as usual, is only an approximation and a simplification of reality. For example, the linear utility functions used in this model assume independence of preference attributes and may have to be modified. Also the assumption that investment commitments once made, are irreversible may also have to be revisited to introduce more flexibility.

4. Conclusion

There is a spectrum of models that blend risk analysis and game theory, e.g., financial or strategic. In this paper, I described two such models: a principal-agent model of management in an engineering world, and a strategic model in a world of conflicts and insurgency. The originality of the first is to show the relevance of the classic probabilistic risk analysis when human behaviors, influenced by management decisions, can affect the risk of failure of a technical system. In the second case, I show how the dynamics of the game in a counter-insurgency policy analysis can be described by a simulation over time to yield the probability that the system remains unstable after years of investment. The success depends on the nature of these investments (short-term versus long-term). The challenge is mostly in the description of the actors' behaviors and the descriptions of their reciprocal influences.

In all cases, the most important feature of these models is that they rely on systems analysis and probability. This kind of analysis could not be done based on statistics that do not exist, are insufficient because the sample size is too small, or simply irrelevant because the situation has changed.

References

- [1] Bier Vicki M., A. Nagaraj and V. Abhichandani, 2005: "Protection of Simple Series and Parallel Systems with Components of Different Values" Reliability Engineering and System Safety, Vol. 87, pp. 315-323.
- [2] Garber R.G., and M.E. Paté-Cornell, 2006: "Managing shortcuts in Engineering Systems and Their Safety Effects: A Management Science Perspective", Proceedings of PSAM8, Paper # 0036.
- [3] Garber R.G., 2007: "Corner Cutting in Complex Engineering Systems: A Game Theoretic and Probabilistic Modeling Approach", Doctoral Thesis, Department of Management Science and Engineering, Stanford University.
- [4] Gibbons R. 1992: Game Theory for Applied Economists, Princeton University Press, Princeton, NJ
- [5] Harasanyi, J. C. 1959: "A Bargaining Model for the Cooperative N-Person Game" in A.W. Tucker and R.D. Luce (eds), Contributions to the Theory of Games IV, Annals of Mathematical Studies, Vol 40, Princeton University Press, pp. 325-355.
- [6] Hausken Kjell, 2002: "Probabilistic Risk Analysis and Game Theory", Risk Analysis, Vol. 22, pp. 17-27.
- [7] Hong Y. and G. Apostolakis, 1992: "Conditional Influence Diagrams in Risk Management", Risk Analysis, Vol. 13, No. 6, pp. 625-636.

- [8] Kucik P. D., 2007: "Probabilistic Modeling of Insurgency", Doctoral Thesis, Department of Management Science and Engineering, Stanford University.
- [9] Kumamoto H. and E.J. Henley, 1996: Probabilistic Assessment and Management for Engineers and Scientists, IEEE Press, New York.
- [10] Milgrom P.R. and J. Roberts, 1992: Economics, Organization and Management, Prentice-Hall, New Jersey.
- [11] Murphy, D.M. and M.E. Paté-Cornell, 1996: "The SAM Framework: A Systems Analysis Approach to Modeling the Effects of Management on Human Behavior in Risk Analysis", Risk Analysis, Vol. 16, No. 4, pp.501-515, 1996.
- [12] Nash, J., 1950: "Equilibrium Points in N-Person Games", Proceedings of the National Academy of Sciences, 36 (1) 48-49.
- [13] Paté-Cornell, M.E. 2007: "Rationality and Imagination: The Engineering Risk Analysis Method and Some Applications" in Advances in Decision Analysis, Edward, Miles and von Winterfeldt Eds., Cambridge University Press, in press.
- [14] Paté-Cornell, M.E. and P.S. Fischbeck, 1993: "PRA as a Management Tool: Organizational Factors and Risk-based Priorities for the Maintenance of the Tiles of the Space Shuttle Orbiter," Reliability Engineering and System Safety, Vol. 40, No. 3, pp. 239-257.
- [15] Paté-Cornell, M.E. and S.D. Guikema, 2002 "Probabilistic Modeling of Terrorist Threats: a Systems Analysis Approach to Setting Priorities Among Countermeasures," Military Operations Research, Vol. 7, No 4.
- [16] Paté-Cornell, M.E. 2008: "Games and risk analysis: three examples of single and alternate moves", with the collaboration of Russ Garber, Seth Guikema and Paul Kucik in Game theory and Risk Analysis (Bier and Cox Eds.) in press.
- [17] Raiffa, H, 1968: Decision Analysis, Weiley, Cambridge Mass.
- [18] Von Neumann J. and O. Morgenstern, 1947: Theory of Games and Economic Behavior, Princeton University Press, Princeton, NJ.