Adaptive Policymaking under Deep Uncertainty: Optimal Preparedness for the next pandemic

Caner Hamarat, Jan Kwakkel, Erik Pruyt
Delft University of Technology
Faculty of Technology, Policy and Management
Jaffalaan 5, 2628 BX, Delft, The Netherlands
+31 15 2788080
c.hamarat@tudelft.nl

Abstract

The recent flu pandemic in 2009 caused a panic about the possible consequences due to deep uncertainty about an unknown virus. Overstock of vaccines or unnecessary social measures to be taken were all due to uncertainty. However, what should be the necessary actions to take in such deeply uncertain situation where there is no or very little information available? For uncertain and complex future, adaptivity and flexibility should be the main aim for designing robust policies. Here, we propose an iterative approach for designing adaptive and robust policies in the presence of deep uncertainty. A crucial part of this approach is the use of monitoring systems that provide the adaptivity and flexibility of the policy design. In the monitoring system, signposts to track specific information are defined. Specific values of these signposts are called triggers and they are triggered when pre-specified conditions occur in the system. The specification of trigger values is crucial for the policy performance but has not been studied in depth. Here, we use robust optimization to optimize the trigger values. This paper shows that our proposed approach with robust optimization improves policy design in deeply uncertain and complex situations where very little information is available.

Keywords: adaptive policymaking, robust optimization, influenza, Exploratory Modeling and Analysis, deep uncertainty.

1. Introduction

With the first reports of a new flu virus in March-April 2009, the panic started to spread all around the world. The very little amount of what was known about this new flu variant caused the panic to grow more. Most European governments ordered vaccines double the amount of the total population. For instance, Netherlands with 35 million doses ordered for a population of approximately 16 million people (NRC 2009). Such extremely cautious actions were caused by the deep uncertainty of the Mexican Flu. In the initial stages of the pandemic, the information was very limited and even not available sometimes. Therefore, there was a need for policies that could help overcome the deep uncertainty and be ready for surprises in the future.

Adaptive Policymaking

Uncertainty is quite impossible to avoid in policymaking. It is possible to face unforeseen events caused by uncertainty. Therefore, uncertainty should be taken into account in the policy design. An approach in policymaking is to aim for a static policy based on a best estimate future, which is doomed to perform poorly in an uncertain and complex future (Walker, Marchau, and Swanson 2010). For an uncertain and complex future, adaptivity and flexibility should be the main aim for designing robust policies (Neufville and Scholtes 2011; Lempert, Popper, and Bankes 2003; Neufville and Odoni 2003; Swanson et al. 2010; Walker, Rahman, and Cave 2001; Schwartz and Trigeorgis 2004).

There have been various studies for designing adaptive policies. An early study by Dewey (1927) puts forth an argument proposing that policies be treated as experiments, with the aim of promoting continual learning and adaptation in response to experience over time (Busenberg 2001). Policies are designed from the outset to test clearly formulated hypotheses about the behavior of an ecosystem being changed by human use (Lee 1993). A similar attitude is also advocated by Collingridge (1980) with respect to the development of new technologies. Given ignorance about the possible side effects of technologies under development, he argues that one should strive for correctability of decisions, extensive monitoring of effects, and flexibility. In a recent study by Walker et al. (2001), a structured, stepwise approach for dynamic adaptation is proposed. In this approach, plans should be adaptive, where only those actions that are non-regret and time-urgent are taken and others are postponed to a later stage. This part of the plan is called the basic policy. The postponed actions are taken with the help of a monitoring system and pre-specification of responses when specific trigger values are reached and this part is called as the contingency planning. The aim is to achieve a robust plan that is flexible and adaptive to plausible futures.

A common characteristic of these approaches is the combination of time urgent actions to be taken immediately with pre-specified action taken in response to how the future unfolds. In order to achieve a robust and adaptive policy design, it is important to correctly specify when to respond with these pre-specified actions. To this purpose, signposts to track specific information can be defined for monitoring the system. Specific values of these signposts are called triggers and they are triggered when pre-specified conditions occur in the system (Kwakkel, Walker, and Marchau 2010). However, the literature remains silent on the monitoring system and the specification of trigger values. A common approach is to consult for expert opinions or to estimate values based on historical trends. These approaches are open to surprises caused by uncertainty and can lead to poor policy performances. To this purpose, trigger specification should be done by using more intelligent and robust methods. A solution approach can be to use optimization.

Robust Optimization

Optimization is very popular in almost every aspect of decision making and also used in policymaking. Optimization can be defined as trying to find the optimum solution among a set of possible alternatives without violating certain constraints. It is mostly employed for predictive purposes, where the aim is for a single best estimate solution. However, under deep uncertainty, this predictive approach might be misleading for policymaking, where often an optimum single goal is not the main aim (Bankes 2011). Robust optimization aims to overcome this difficulty of uncertainty. Robust optimization methods aim at finding optimal outcomes in the presence of uncertainty about input parameters (Ben-Tal and Nemirovski 1998, 2000; Bertsimas and Sim 2004). To this purpose, robust optimization methods can be of great use for adaptive policymaking under uncertainty.

There is an enormous variety of different techniques and methods in the optimization literature. Among various optimization techniques, Genetic Algorithm (GA) is a commonly used heuristic method. GA is flexible and efficient in complex and irregular solution spaces. It mimics the evolution process and tries to find the fittest survivor. In GA, a candidate solution is represented as a chromosome where each allele of the chromosome is a decision variable. Each trigger value can be considered as a decision variable that these trigger values form a candidate policy. So, Genetic Algorithm can be used for the specification of trigger values in adaptive policymaking.

Organization of the paper

The rest of the paper is organized as follows: Section 2 introduces Exploratory Modeling and Analysis, our proposed iterative adaptive policymaking approach and Genetic Algorithm. In Section 3, the flu model is explained. Section 4 illustrates the results and Section 5 contains a discussion of the results and conclusions.

2. Methodology

Exploratory Modeling and Analysis

Exploratory Modeling and Analysis (EMA) is a research methodology that uses computational experiments to analyze complex and uncertain systems and support long-term strategic decision making under deep uncertainty (Agusdinata 2008; Bankes 1993). EMA can be contrasted with the use of models to predict system behavior, where models are built by consolidating known facts into a single package (Hodges and Dewar 1992). In predictive modeling, a single best estimate model is used as a surrogate for the actual system. Where applicable, this consolidative methodology is a powerful technique for understanding the behavior of complex systems. Unfortunately, for many systems of interest, the construction of a model that may be validly used as surrogate is simply not a possibility (Campbell et al. 1985; Hodges and Dewar 1992). For many systems, a methodology based on consolidating all known information into a single model and using it to make best estimate predictions can be highly misleading. However, models can be constructed that are consistent with the available information, but such models are not unique. Rather than specifying a single model and falsely treating it as a reliable image of the system of interest, the available information is consistent with a set of models, whose implications for potential decisions may be quite diverse. A single model run drawn from this potentially infinite set of plausible models is not a “prediction”; rather, it provides a computational experiment that reveals how the world would behave if the various guesses made in any particular model about the various unresolvable uncertainties were correct. By conducting many such experiments, EMA provides insights and understanding about the system functions and effectiveness/robustness of policies despite the presence of deep uncertainty. EMA is not a modeling technique by itself, but it is a methodology for building and using models under deep uncertainty.

Computer Aided Dynamic Adaptive Policy Design

EMA could be used to develop dynamic adaptive policies. EMA allows for the explicit representation and exploration of a multiplicity of plausible futures under deep uncertainty. Thus, EMA can be used to identify the vulnerabilities and opportunities that this ensemble of futures holds, paving the way for designing targeted policies that address vulnerabilities or seize opportunities. The efficacy of these policy designs can then be tested against the ensemble of futures. Moreover, EMA can be used to identify conditions under which changes in a policy are required. That is, it can help in developing the monitoring system and its associated actions. It thus appears that EMA can be of use in all the steps of the design phase of a dynamic adaptive policy.

An iterative approach that is called as Computer Aided Dynamic Adaptive Policy Design (CADAPD) has been proposed (Hamarat, Kwakkel, and Pruyt Forthcoming):

1. conceptualization of the problem,
2. identification of the uncertainties,
3. development of an ensemble of models for exploring the uncertainties,
4. running the computer model(s) without introducing any policies in order to generate the ensemble of futures,
5. analysis of the results obtained from Step 4 in order to identify the vulnerabilities and opportunities,
6. design of candidate policies for addressing vulnerabilities and seizing opportunities,
7. testing of candidate policies across the ensemble of futures,
8. iteration through Steps 5-7 until a satisfying policy emerges.

Identification of vulnerabilities and opportunities is done by using the Patient Rule Induction Method (PRIM) (Friedman and Fisher 1999; Lempert et al. 2006; Groves and Lempert 2007). PRIM can be used for data analytic questions where the analyst tries to find combinations of values for input variables that result in similar characteristic values for the outcome variables. Specifically, one seeks a set of subspaces of the input variable space within which the values of output variables are considerably different from the average value over the entire output domain.

In addition, it is usually desired that these regions can be described in the form of easily interpretable rules. In the context of this paper, the input space is the uncertainty space. Thus, we use PRIM to identify the combinations of the uncertainties in the global uncertainty space that result in highly desirable or undesirable outcomes of interest.

Genetic Algorithm

Genetic Algorithms (GA) are optimization methods based on natural selection as can be observed in biological systems (Fraser and Burnell 1970; Holland 1975). This approach requires constructing an initial population composed of chromosomes, where each chromosome represents a candidate solution. Next, the fitness of each population member is assessed using a user specified objective function. In light of the fitness scores of the current population members, the next generation is created. For creating the next generation, the new members are reproduced from those selected through evolutionary processes such as crossover and mutation. Once the next generation is created, the fitness calculations are computed again for the new population members. This process of fitness evaluation and reproduction of new generation is repeated until a pre-specified termination criterion is met. Possible termination criteria include reaching a desired solution, a fixed number of iterations, and convergence of the fitness scores.

GA are commonly used for solving decision making problems due to their flexibility and efficiency in complex and irregular solution spaces (Chambers 1999). We argue that GA can be efficiently used in CADAPD for optimizing trigger values. The chromosome structure for representing a candidate solution can be easily used for representing a set of trigger values as a candidate policy setting. In this case, each genome of a chromosome will be a trigger value and each chromosome will represent is a complete representation of the monitoring system. So, GA can be employed for optimizing the set of trigger values.

The trigger values for the various actions in a monitoring system should be robust across the ensemble of plausible futures. The criterion used for performance calculation in robust optimization is quite important. There are different criteria such as minimizing the maximum regret (minimax), maximizing the minimum gain (maximin) or maximizing the maximum gain (maximax) (Winston & Goldberg, 2004). GA is often used for robust optimization (Herrmann, 1999; Li et al., 2005; Maruyama & Igarashi, 2008). In this study, a cardinality criterion, which is the number of cases above a certain threshold, is utilized. We start by generating a population of trigger values. Each population member is a set of trigger values for the actions in the monitoring system. The performance of each population member is evaluated according to the cardinality criterion over a fixed number of plausible futures.

3. The Flu Model

In this study, an exploratory System Dynamics model (Pruyt and Hamarat 2010) about the recent flu pandemic, which is known as the Mexican Flu or A(H1N1)v. The modeled world is divided in three regions: the Western World, the densely populated Developing World, and the scarcely populated Developing World. Only the two first regions are included in the model because it is assumed that the scarcely populated regions are causally less important for dynamics of flu pandemics. The basic stock-flow structure of one of the two regional sub-models of the ESD simulation model is displayed in Figure 1.

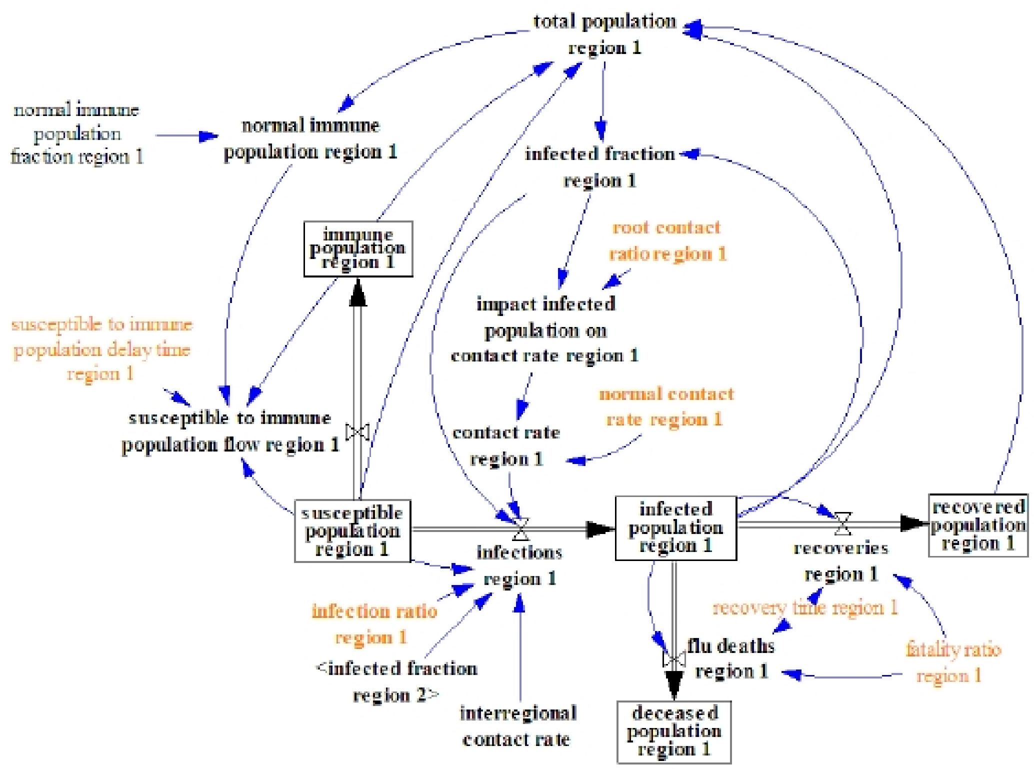

Figure 1: The stock-flow diagram of the region 1.

The figure represents only the Western World but the stock-flow diagram of the region 2 is very similar to the region 1, with some minor differences. For instance, less average normal contact rate, higher infection rate are assumed for region 2. The link between two regions is provided via the interregional contact rate. A more detailed information about the model and its specifications can be found in (Pruyt and Hamarat 2010).

4. Analysis

In this study, an exploratory System Dynamics model about the recent flu pandemic is used for illustrating how our proposed approach can be effectively used for adaptive policymaking. It is a small, high-level and simplistic model. However, it is still quite useful for explorative purposes. The model includes a variety of uncertainties to be explored over about the flu pandemic (See Table 1). The main outcomes of interest are the number of deceased people in region 1 and the fraction of infected people over the total population in region 1. The model is executed over a time horizon of 48 months. The details about a quick exploration of basic behaviors and a sensitivity analysis of important variables can be found in (Pruyt and Hamarat 2010).

Table 1: Uncertainties and their upper and lower limits to be explored

Parameter	Lower Limit	Upper Limit
additional seasonal immune population fraction region 1	0.0	0.5
additional seasonal immune population fraction region 2	0.0	0.5
fatality ratio region 1	0.0001	0.1
fatality ratio region 2	0.0001	0.1
initial immune fraction of the population of region 1	0.0	0.5
initial immune fraction of the population of region 2	0.0	0.5
normal interregional contact rate	0.0	0.9
vaccination decision action delay	4	9
permanent immune population fraction region 1	0.0	0.5
permanent immune population fraction region 2	0.0	0.5
recovery time region 1	0.2	0.8
recovery time region 2	0.2	0.8
root contact rate region 1	1.0	10.0
root contact rate region 2	1.0	10.0
infection ratio region 1	0.0	0.1
infection ratio region 2	0.0	0.1
normal contact rate region 1	10	200
normal contact rate region 2	10	200

No Policy

The initial analysis starts without introducing any policy to see the behavior of the system without any external influence. Exploring over the uncertainty space, 10,000 simulations are executed. Figure 2 shows the behavior of the following outcomes: the number of deceased people in region 1 and the fraction of infected people over the total population in region 1. Due to the illustrative constraints, only 1,000 randomly selected simulations out of 10,000 are visualized. The blue shaded areas show the upper limits of the outcomes over 10,000 simulations. The number of deceased people has an upper limit around 50 million deaths. However, only few cases have such catastrophic results. More than half of the simulations result in a casualty number of less than one million people. For the infected fraction of region 1, the observed behavior is a sharp early peak followed by a mild peak and afterwards gradually decreasing. The maximum peak observed is around a level of 40-45%, which means 40-45% of the population of region 1 is infected. Similarly, such catastrophic cases are very unlikely but still plausible.

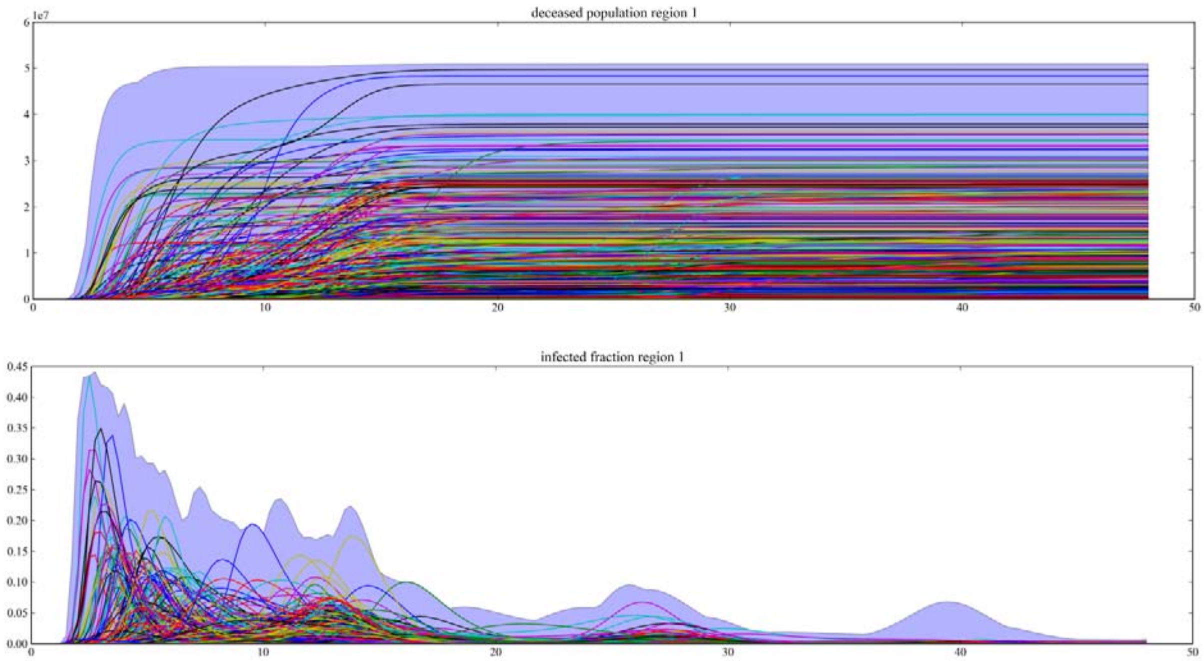

Figure 2: Deceased population and infected fraction for region 1 without any policy

Since such catastrophic cases where dramatic number of deaths and high levels of infection happen are not very likely but plausible, it is necessary to understand the underlying reasons behind such behaviors. If the underlying mechanisms can be revealed, it can be possible to design action(s) to prevent the system to face catastrophe. To this purpose, PRIM is used for finding the combination of uncertainties that has common characteristics. We looked for the cases where the number of deceased people is above one million. We classified those cases as 1 and others as 0 and PRIM is applied to find subspaces in the uncertainty space where the average of the cases in a subspace is above 0.9. PRIM results show four different subspaces, where each colored line in

Figure 3 represents a different subspace. The uncertainties shown on the figure are the relevant uncertainties for the classification criterion used for PRIM. Gray shaded area ranges between 0 and 1 where each uncertainty range is also normalized correspondingly. For instance, the red line shows that a combination of higher recovery time for region 1, lower root contact rate for region 1, higher infection rate and normal contact rate for region 1 result in cases where the number of deceased people is higher than one million. The other colored lines can be interpreted in a similar manner. A common observation for each of the PRIM subspaces is that higher infection rate and normal contact rate for region 1 cause undesirable behaviors.

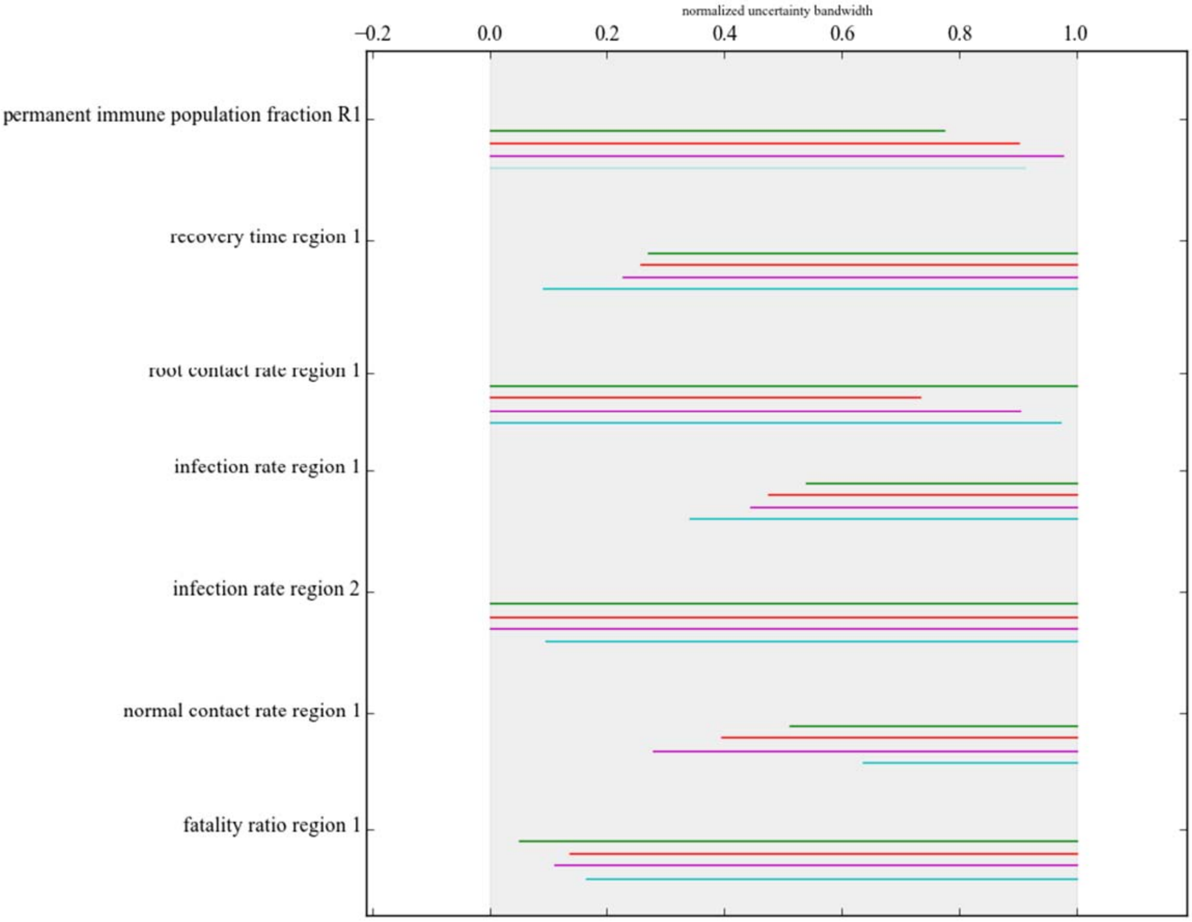

Figure 3: PRIM Results for without any policy

Basic Policy

In the light of PRIM results, a basic policy is designed to overcome the undesirable results focusing on reducing the infection and social contact. The basic policy consists of two actions. The first action is to vaccinate 40% of the region 1 against the influenza to reduce the infection level. The second action of the basic policy is for the reduction of social contact. An orchestrated contact rate reduction is applied by monitoring the infection level. An S-shaped lookup function that monitors the infected fraction is designed to reduce the normal contact rate accordingly. This lookup function connects the following points:

[(0,0),(0.05,0.05),(0.1,0.2),(0.2,0.75),(0.3,0.85),(0.4,0.9),(1,1)]

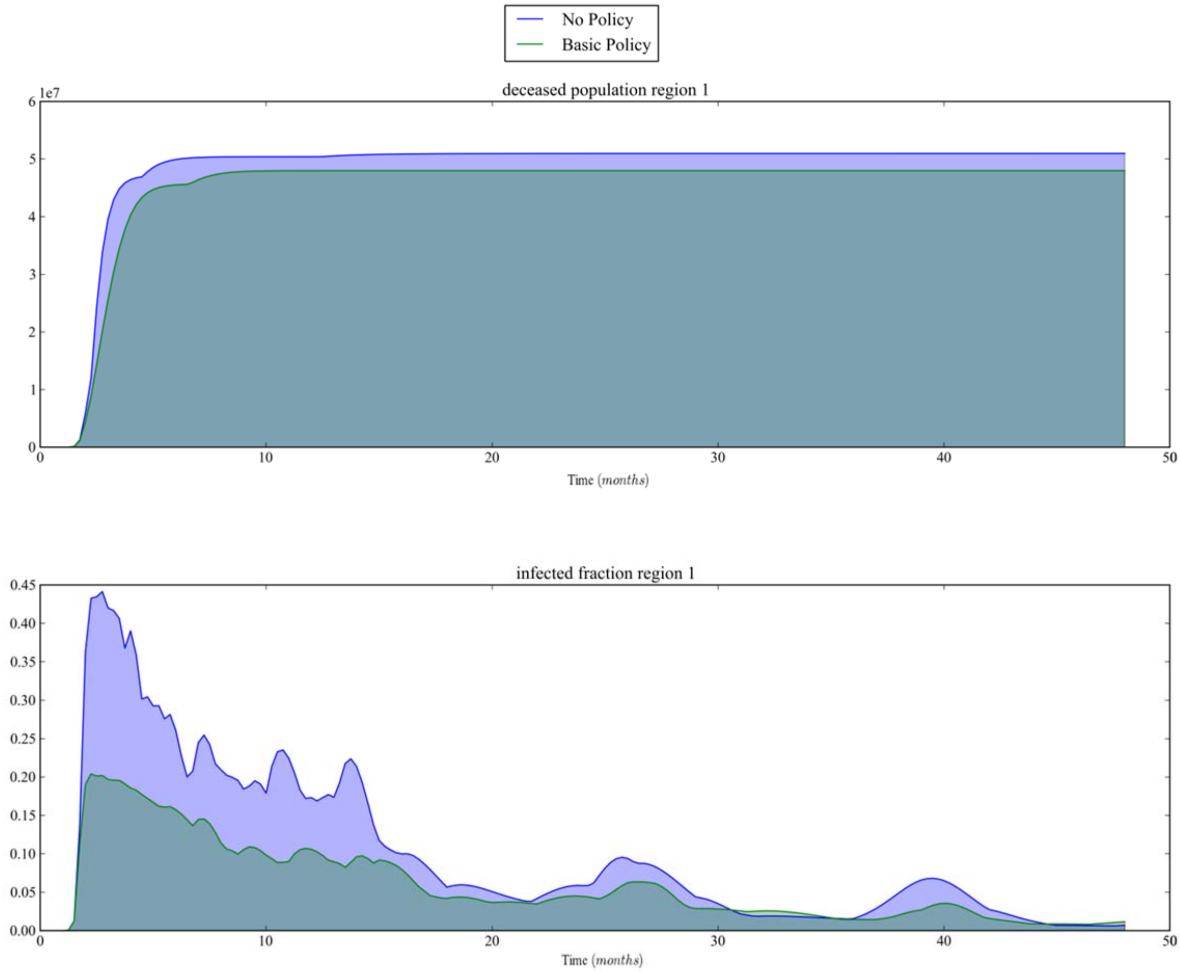

Figure 4: Comparison of No Policy and Basic Policy

This basic policy again tested by exploring the uncertainty space with 10,000 simulations. In order to illustrate the effectiveness of the basic policy, the results for no policy and basic policy are compared in Figure 4. The blue envelope showing the upper and lower limits over 10,000 simulations belongs to the no policy and the green envelope is for the basic policy. As can be seen both in terms of the deceased population and infected fraction, there is a reasonable reduction with the introduction of the basic policy. The maximum infected fraction peak is reduced from 40-45% to a level around 20%. Similarly, the maximum casualty level decreased from 50 million to 45 million people. Although there is a considerable improvement with the basic policy, it is still plausible to face dramatic undesirable results.

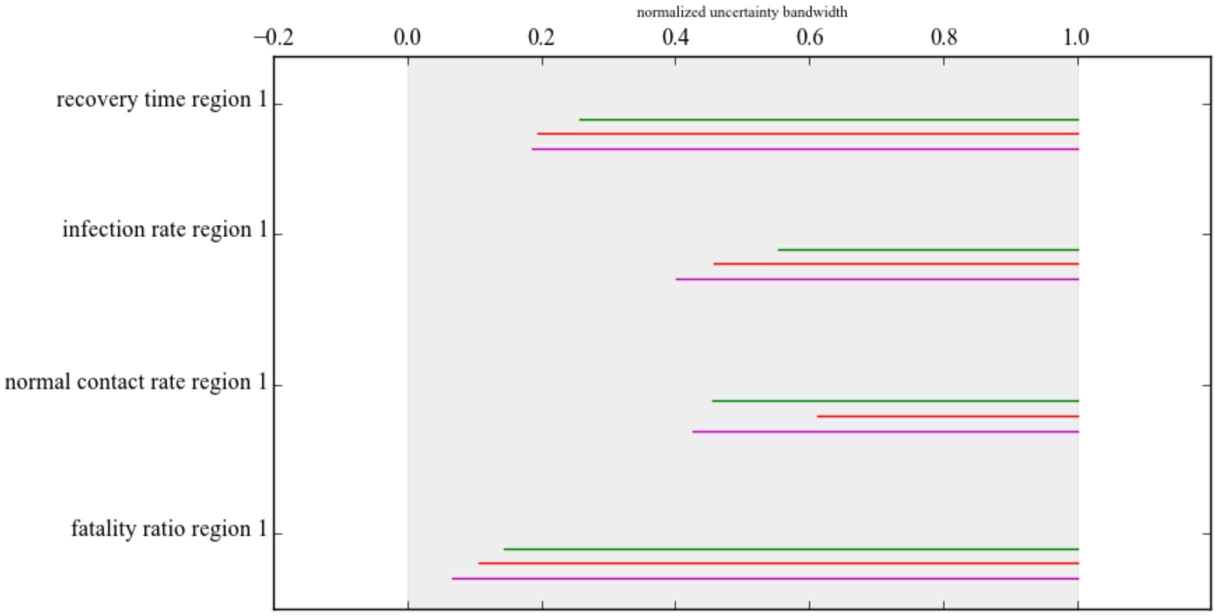

Figure 5: PRIM Results for Basic Policy

In order to find the vulnerabilities of our basic policy, we applied PRIM with the same criterion used before on the results of the basic policy. Figure 5 shows the relevant uncertainties which are the recovery time, the infection rate, the normal contact rate and the fatality ratio of region 1. The most obvious observation, again, is that it is still needed to reduce the social contact and the infection level.

Adaptive Policy

Under deep uncertainty, adaptivity should be the main aim for a robust policy design. To this purpose, we modified our basic policy in the light of PRIM results and adaptive policymaking design. First of all, the vaccination action is revised to make it more adaptive and flexible. 40% minimum level of vaccination is kept but in addition to that, vaccinated level is increased gradually according to the observed case fatality ratio. For an observed case fatality ratio (cfr) of 0.1%, the vaccination level is increased to 60%. If the observed cfr is 1%, then vaccination level is 80% and for cfr of 10%, vaccination level is at the maximum 100% level. The aim of this action is to be effective for reducing the infection level. Another crucial issue is the reduction in the social contact. For that, we designed an alert activation monitoring system. This system checks on a weekly basis whether the rate of increase in the infected fraction level is positive or not. If the rate of increase for the infected fraction is positive for three consecutive weeks, then the alert is activated and an extra 50% emergency contact rate reduction is applied. For this action, alert is triggered if the rate is positive for three consecutive weeks, which is our trigger value. These two actions that are added on our basic policy forms our adaptive policy design. The adaptive policy is tested for 10,000 simulations for exploring the uncertainty space and a comparison of the results of the basic and the adaptive policy is shown on Figure 6.

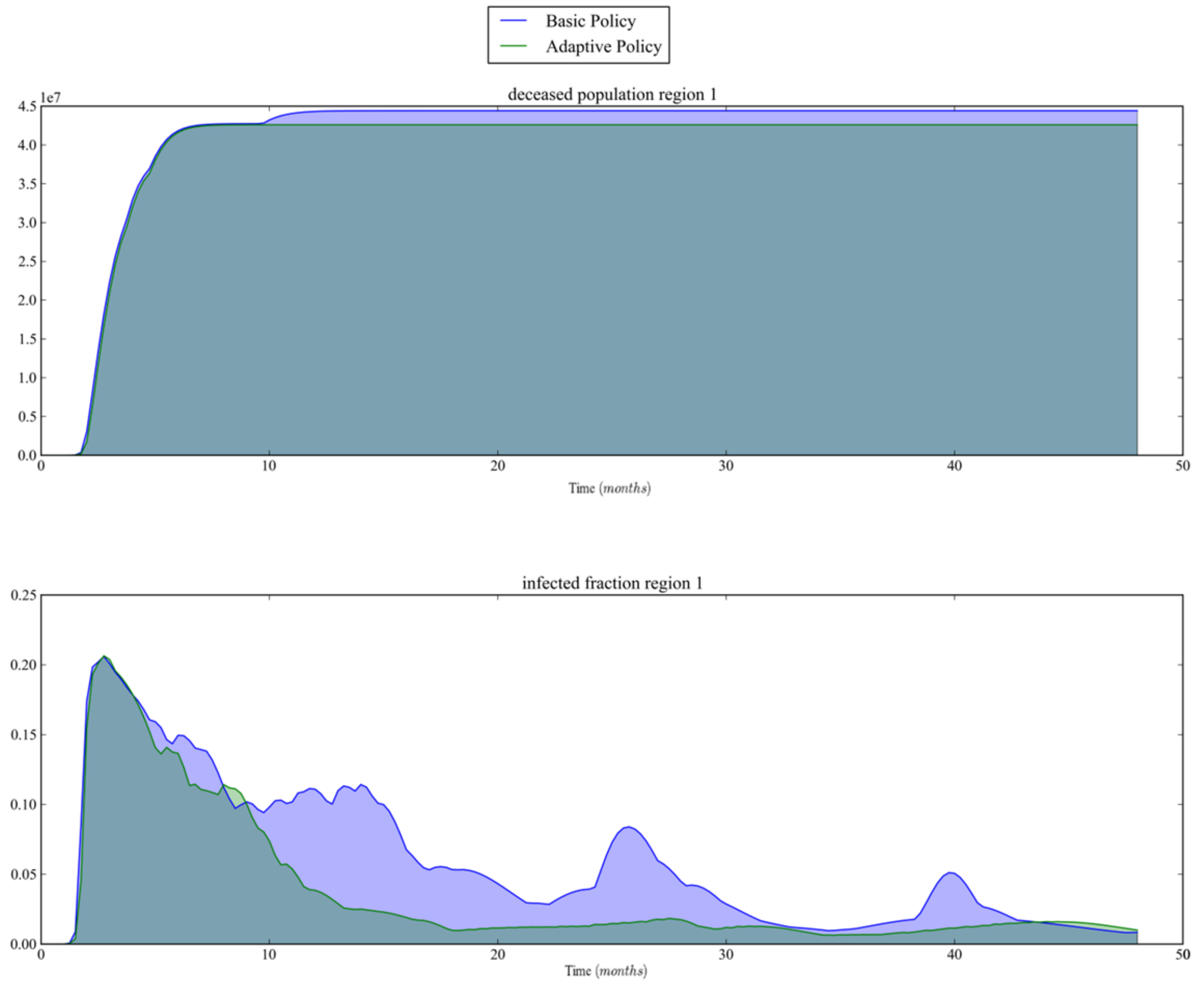

Figure 6: Comparison of Basic Policy and Adaptive Policy

The results illustrate that there is only a little improvement with the introduction of the adaptive policy. The upper limit of the number of deceased people is reduced to a level around 40 million deaths. In terms of infected fraction, the maximum peak of the infection is around the same level of 0.20 but the adaptive policy seems to be effective for preventing the later smaller peaks. In our adaptive policy design, two triggers are used for designing the alert activation system: one of the triggers is the check period and the other is the number of consecutive positive checks. The specification of the trigger values is done based on experience and estimates. However, by using robust optimization, it is possible to specify these triggers more intelligently.

Optimized Adaptive Policy

In order to optimize the triggers used in the adaptive policy, we use Genetic Algorithm (GA) in our robust optimization. A set of these two triggers is a candidate solution and for calculating the fitness of a candidate solution, we utilize a cardinality based criterion. Each candidate solution is executed over 200 cases and the fraction of the number of cases that are below 600,000 deceased people over 200 cases is used as the fitness score. The reason for using the threshold of 600,000 is that it is 0.1% (Guidance 2007) of 600 million which is the total population of region 1. With a population size of 50 and over 50 generations, robust optimization is executed by using GA with a crossover rate of 0.7 and a mutation rate of 0.01. The optimized values of the triggers are as follows: 0.4057 month for check period and 1.013 for the number of consecutive checks. This means that the optimized policy for the alert activation system should be to check every 12 days and if the rate of increase in the infected fraction level is positive for 12 days, then the alert should be activated. Current alert design only activates an extra social contact reduction but it does not do anything about vaccination. To this purpose, the optimized policy is redesigned to start additional vaccination when the alert is activated. The comparison of the resulting optimized policy with the other policies can be found in Figure 7.

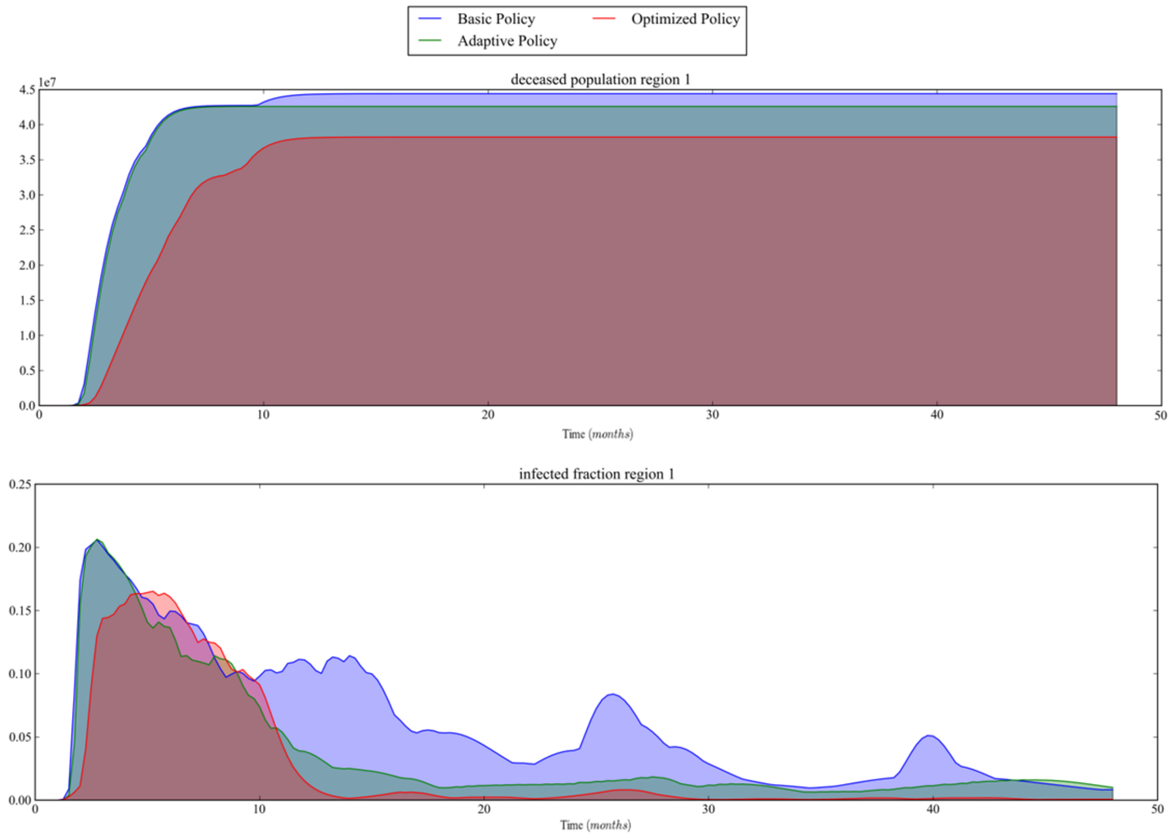

Figure 7: Comparison of Basic Policy, Adaptive Policy and Optimized Policy

The figure above shows that the optimized policy is effective both in terms of reducing the upper limit of the deceased population and the maximum peak of the infected fraction. By the optimized policy, the upper limit of the deceased population is at a level around 35 million and the maximum peak of infected fraction is around 15%. Additionally, the later smaller peaks are almost gone with the introduction of the optimized policy.

In order to illustrate the performances of the policies more clearly, Figure 8 represents the comparison of the policies. In this figure, each graph includes 10,000 blue dots where each blue dot represents a single simulation. Each single dot shows the corresponding numbers of the deceased population and the maximum level of infected fraction reached for the associated simulation. For the basic policy, there are more cases where the number of deceased population and the maximum infected fraction level is high. Adaptive policy helps to reduce the number of these catastrophic cases. Furthermore, the optimized policy is more effective because most of the cases are kept at very low levels of deaths and infection.

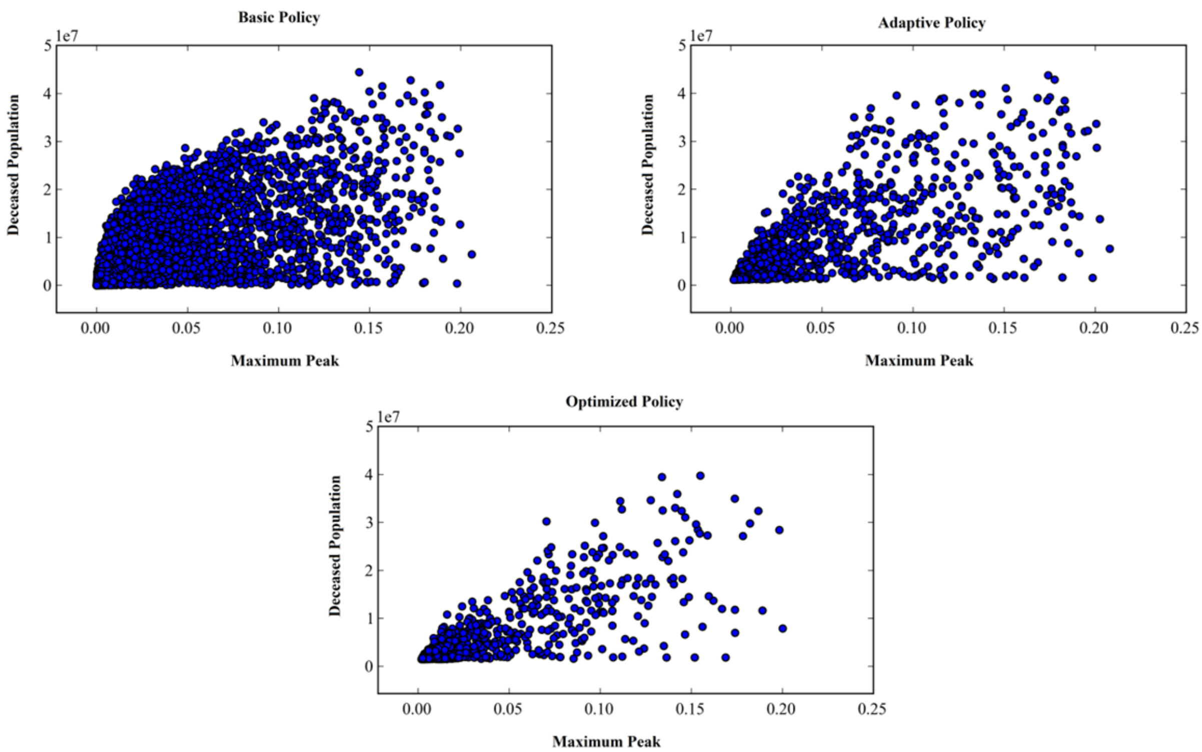

Figure 8: Comparison of the policies according to the maximum peak and the number of deceased people

In our analysis, each policy was executed over 10,000 simulations. Table 2 shows the number of cases where the number of deceased population is above 100,000 and 1,000,000 for each policy. It is observed that at each step, there is a clear decrease in the number of deaths. This shows that our iterative approach is effective for designing policies that are robust and adaptive under deep uncertainty.

Table 2: Number of cases above certain number of deceased population for each policy

	Number of cases (out of 10000) where deceased population is	
	>100.000	> 1.000.000
Basic Policy	4953	3589
Adaptive Policy	2316	1044
Optimized Policy	1532	544

5. Discussion & Conclusions

The flu pandemic of 2009 has shown us that when the information is very limited or even not available, it is quite difficult to design a robust strategy. Under deep uncertainty, adaptivity and flexibility should be aimed in order to be ready for surprises and unforeseen events. This study illustrates that our proposed approach can be effectively used for developing robust, adaptive and flexible policy under conditions of deep uncertainty.

In this study, an iterative model-based approach for adaptive policymaking under uncertainty and the use of robust optimization in policymaking have been illustrated through a case about the 2009 flu pandemic. Starting without introducing any policy, we identified the vulnerabilities of the model and designed a basic policy according to the PRIM analysis. The results of the basic policy showed that there was still room for improvement in the number of deaths and infection levels. To this purpose, an adaptive policy was designed in the light of the PRIM analysis of the basic policy. Since the triggers used for adaptive design were based on estimates, new values were specified for these triggers by using robust optimization.

Robust optimization is an optimization methodology where robustness is aimed in the presence of uncertainty. A crucial concept in robust optimization is the choice of robustness calculation. In this study, we also employ a robust optimization method for calculating the fitness of a candidate solution according to a cardinality criterion. However, there are other criteria such as minimax, maximax or maximin. Maximin is a common approach used in robust optimization where we look for the worst cases for each candidate solution and try to minimize the worst cases. Since the choice of optimization criterion can have a great importance on the solution, other possible criteria should also be considered and tested.

6. Acknowledgements

This study is part of the project “Dealing with Uncertainties in Infrastructure Planning and Design: Exploratory Modeling, Real Options analysis and Policy Design” which is supported by the Next Generation Infrastructures (NGI) Foundation.

7. References

Agusdinata, D. B. 2008. Exploratory modeling and analysis: a promising method to deal with deep uncertainty, Technology, Policy and Management, Delft University of Technology, Delft.

Bankes, S. 1993. Exploratory Modeling for Policy Analysis. *Operations Research* 41 (3):435-449.

Bankes, Steven. 2011. The Use of Complexity for Policy Exploration. In *The SAGE Handbook of Complexity and Management*, edited by P. Allen, S. Maguire and B. McKelvey. London: SAGE Publications Ltd.

Ben-Tal, A., and A. Nemirovski. 1998. Robust convex optimization. *Mathematics of Operations Research*:769-805.

Ben-Tal, A., and A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* 88 (3):411-424.

Bertsimas, Dimitris, and Melvyn Sim. 2004. The Price of Robustness. *Operations Research* 52 (1):35-53.

Busenberg, G. J. 2001. Learning in Organizations and Public Policy. *Journal of Public Policy* 21:173-189.

Campbell, David, Doyne Farmer, Jim Crutchfield, and Erica Jen. 1985. Experimental mathematics: the role of computation in nonlinear science. *Commun. ACM* 28 (4):374-384.

Chambers, L. 1999. *The Practical Handbook of Genetic Algorithms: Complex coding systems*. Vol. 3: CRC.

Collingridge, D. 1980. *The Social Control of Technology*. London, UK: Frances Pinter Publisher.

Dewey, J. 1927. *The Public and its Problems*. New York: Holt and Company.

Fraser, A., and D. Burnell. 1970. *Computer models in genetics*: McGraw-Hill New York.

Friedman, J. H. , and N. I. Fisher. 1999. Bump Hunting in high-dimensional data. *Statistics and Computing* 9:123-143.

Groves, David G., and Robert J. Lempert. 2007. A new analytic method for finding policy-relevant scenarios. *Global Environmental Change* 17 (1):73-85.

Guidance, IPP. 2007. Community Strategy for Pandemic Influenza Mitigation in the United States—Early, Targeted, Layered Use of Non-pharmaceutical Interventions. *Atlanta, GA: Centers for Disease Control and Prevention, Department of Health and Human Services*.

Hamarat, Caner, Jan Kwakkel, and Erik Pruyt. Forthcoming. Energy Transitions: Adaptive Policymaking Under Deep Uncertainty. In *Manuscript submitted to Technological Forecasting and Social Change*.

Hodges, J. S., and J. A. Dewar. 1992. Is it You or Your Model Talking? A Framework for Model Validation. Santa Monica: RAND.

Holland, J.H. 1975. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*: University of Michigan Press.

Kwakkel, J. H., W. E. Walker, and V. A. W. J. Marchau. 2010. Adaptive airport strategic planning. *European Journal of Transportation and Infrastructure Research* 10 (3):249-273.

Lee, K. 1993. *Compass and Gyroscope: Integrating Science and Politics for the Environment*. Washington: Island Press.

Lempert, R. J., S. Popper, and S. Bankes. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Santa Monica, California: RAND.

Lempert, Robert J., David G. Groves, Steven W. Popper, and Steve C. Bankes. 2006. A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios. *Management Science* 52 (4):514-528.

Neufville, Richard de, and A. Odoni. 2003. *Airport Systems: Planning, Design, and Management*. New York: McGraw-Hill.

Neufville, Richard de, and Stefan Scholtes. 2011. *Flexibility in Engineering Design*. Massachusetts, USA: The MIT Press.

NRC. *Mexican flu to cost Netherlands 700 million euros* 2009. Available from http://vorige.nrc.nl/article2276325.ece.

Pruyt, Erik, and Caner Hamarat. 2010. The Influenza A (H1N1) v Pandemic: An Exploratory System Dynamics Approach. Paper read at The 28th International Conference of the System Dynamics Society, at Seoul, South Korea.

Schwartz, Eduardo S., and Lenos Trigeorgis. 2004. *Real Options and Investment under Uncertainty: Classical Readings and Recent Contributions*: The MIT Press.

Swanson, Darren, Stephan Barg, Stephen Tyler, Henry Venema, Sanjay Tomar, Suruchi Bhadwal, Sreeja Nair, Dimple Roy, and John Drexhage. 2010. Seven tools for creating adaptive policies. *Technological Forecasting and Social Change* 77 (6):924-939.

Walker, W. E., S. A. Rahman, and J. Cave. 2001. Adaptive policies, policy analysis, and policy-making. *European Journal of Operational Research* 128 (2):282-289.

Walker, Warren E., Vincent A. W. J. Marchau, and Darren Swanson. 2010. Addressing deep uncertainty using adaptive policies: Introduction to section 2. *Technological Forecasting and Social Change* 77 (6):917-923.