Policy Analysis, 1962-2012: From Predict And Act To Monitor And Adapt

Farewell Lecture, delivered on Wednesday, 19 October 2011 by Prof.dr. W.E. (Warren) Walker Professor of Policy Analysis At the Delft University of Technology Faculty of Technology, Policy and Management

Walker, Warren

Publication date
2011

Document Version
Final published version

Citation (APA)
Walker, W. (2011). *Policy Analysis, 1962-2012: From Predict And Act To Monitor And Adapt: Farewell Lecture, delivered on Wednesday, 19 October 2011 by Prof.dr. W.E. (Warren) Walker Professor of Policy Analysis At the Delft University of Technology Faculty of Technology, Policy and Management*. Delft University of Technology, Faculteit Techniek, Bestuur en Management.

Prof.dr. W.E. (Warren) Walker

Policy Analysis, 1962-2012: From Predict And Act To Monitor And Adapt

Farewell Lecture 19 October 2011

Faculteit Techniek,
Bestuur en Management

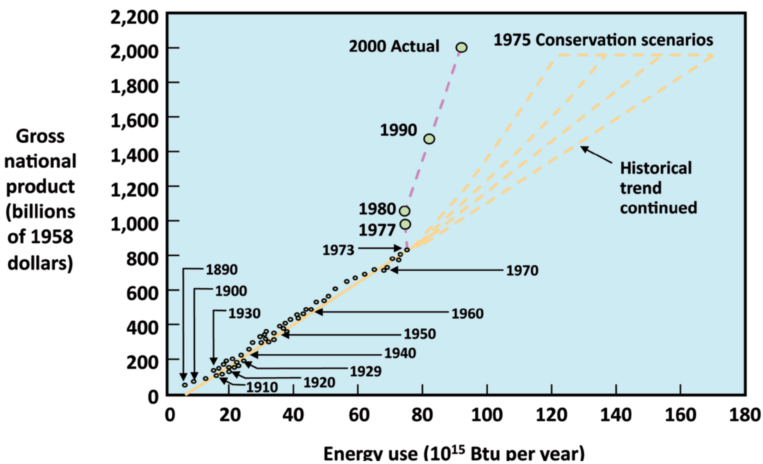

Delft University of Technology

Challenge the future

Policy Analysis, 1962-2012: From Predict And Act To Monitor And Adapt

Farewell Lecture, delivered on Wednesday, 19 October 2011
by Prof.dr. W.E. (Warren) Walker
Professor of Policy Analysis

At the Delft University of Technology
Faculty of Technology, Policy and Management

"[Under deep uncertainty] there is no scientific basis on which to form any calculable probability whatever. We simply do not know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact"
(John Maynard Keynes, 1937)

"It is not the strongest species that survive, nor the most intelligent, but the ones most responsive to change."
(Charles Darwin, Origin of Species, 1859)

"You can't control the wind, but you can adjust your sails."
(Yiddish proverb)

Mijnheer de Rector Magnificus, leden van het College van Bestuur,
Collegae hoogleraren en andere leden van de universitaire gemeenschap,
familie en vrienden, dames en heren.

1. INTRODUCTION

Eleven years ago, on the 29th of November 2000, I gave my Inaugural speech as a Professor of Policy Analysis at TU Delft. The speech was entitled "Uncertainty: The Challenge for Policy Analysis in the 21st Century" (Walker, 2000b). Since that time, I have been working with colleagues at TU Delft, RAND, and other organizations to develop ways to meet that challenge. Today, I want to share with you the fruits of this effort. But, to show you how far we have come, I will review the way that uncertainty as dealt with by operations researchers and policy analysts has evolved over time. I will explain the several dimensions and levels of uncertainty, and that different approaches are needed to deal with the different types. I will suggest that the deep uncertainties about the future that we are now facing require additional approaches. And I will describe one such new approach that seems sensible to apply to dealing with problems involving deep uncertainty.

2. DEFINING UNCERTAINTY

After laying down the challenge in my Inaugural speech, one of my first steps was to define what was meant by uncertainty in the context of policy analysis. That uncertainties exist in practically all policymaking situations was generally understood by most policymakers, as well as by the policy analysts providing decision support. But there was little appreciation for the fact that there are many different dimensions of uncertainty, and there was neither a commonly shared terminology nor an agreement on a generic typology of uncertainties. So, in 2003, six colleagues and I published a paper in which we distinguished three dimensions of uncertainty (Walker, et al., 2003):

1. the *location* of uncertainty - where the uncertainty manifests itself within the policy analysis framework - in the system's external environment, in the model of the system being studied, in the outcomes from the system model, or in the weights the policymakers and other stakeholders would apply to those outcomes. (I will provide more details on the policy analysis framework in a few minutes);
2. the *level* of uncertainty - where the uncertainty manifests itself along the spectrum between complete certainty and total ignorance;
3. the *nature* of uncertainty - whether the uncertainty is due to the imperfection of our knowledge or is due to the inherent variability of the phenomena being described.

In what follows, I will focus on the 'level' dimension and will describe how operations research (OR) and policy analysis have developed tools to deal with each of the levels. Why do I mention both OR and policy analysis? Well, 1962 is the year I began to study OR, and so explains the 1962 in the title of my talk. My Ph.D. is in OR, my work is anchored in the tools of the profession, and I continue to present papers at OR conferences and publish in OR journals. But, I have applied the tools of operations research in my work as a policy analyst at the RAND Corporation and TU Delft. OR offers quantitative tools to help solve real world problems - tools such as linear programming, queueing theory, Monte Carlo simulation, and Decision Analysis. Policy analysis is a systematic approach to helping policymakers find good solutions to real world public policy problems. Although not the same, they have a somewhat symbiotic relationship, which I will try to clarify later.

Before discussing how uncertainty is dealt with in OR and policy analysis, I will begin by presenting some basic background information on policy analysis and the differences between OR and policy analysis. Some of you may be quite familiar with both. But many - especially my friends and relatives - will get lost in what follows if they do not have this as a foundation.

A policy analysis study generates information on the consequences that would follow the adoption of alternative policies. It uses a variety of tools to identify these consequences and to present the consequences to the parties involved in the policymaking process in a manner that helps them come to a decision.

The traditional policy analysis approach (Walker, 2000a) is built around an integral system description of a policy field (see Figure 1). Central to this view is the *system* comprising the policy domain, defined by distinguishing its component elements (or subsystems) and their mutual interrelationships. The system model represents the cause-effect relationships that characterize the system (identified as R in this figure).

The results of these interactions (the system outputs) are called *outcomes of interest* (O) and refer to the characteristics of the system that are considered relevant criteria for the evaluation of policies. The *preferences* refer to the (relative) importance given to the outcomes by crucial stakeholders and the policymakers, reflecting their goals and objectives, and are often represented by giving weights (W) to the outcomes of interest. In case there is a gap between (some of) the system outcomes and the goals, *policies* (P) are implemented to influence the behavior of the system in order to help achieve the goals. If policies were the only forces affecting the system we would have a 'closed loop' system, based upon which the policymakers and stakeholders could fully control the system in order to try to reach their desired goals. However, in reality, there are also *external forces* (X) influencing the system. External forces refer to forces that are not controllable by the policymakers but influence the system significantly (e.g. technological developments, demographic developments, and economic developments). As such, both policies and external forces are developments outside the system that affect the structure of the system and, hence, affect the outcomes of interest to policymakers and other stakeholders.

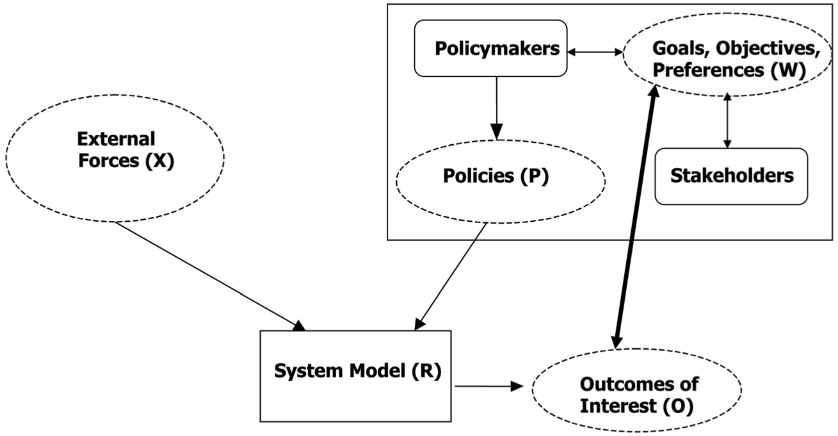

Figure 1 - Framework for model-based policy analysis

The locations of uncertainty we identified in our 2003 paper are labeled X, R, O, and W in this figure. The system model (labeled R) is the focus of most OR efforts. But to a policy analyst, models are merely tools, much as brushes are an artist's tools - they are a means to an end, not the end in itself. In policy analysis, formulating the problem is the key to a project's success. Russell Ackoff, one of the fathers of operations research, who later became somewhat disillusioned with its ability to solve real world problems, once said: "We fail more often because we solve the wrong problem than because we get the wrong solution to the right problem" (Ackoff, 1974).

So, what is the relationship between policy analysis and operations research? Any specific policy analysis study can make use of OR tools. Some operations researchers divide the members of the profession into two nearly disjoint groups - those who focus on the mathematics of operations research - solving what some call 'technical' or 'tame' problems - and those who apply OR to the real social and political problems of the day - solving what some call 'wicked' or 'practical' problems (see Rosenhead, 1989, pp. 10-11). The latter sometimes call what they do 'community operations research' or 'soft OR' (Johnson, 2012). At an OR conference, I had occasion to discuss this conflict with Jonathan Rosenhead - a charter member of the community OR club. He wanted to find a name for the members of this club that would make the distinction clearer to the outside world. I suggested that we call ourselves 'policy analysts', because what policy analysts do is to use technical OR tools (and other tools) in order to solve wicked real-world problems. But, our differences aside, what is critical is that we still consider ourselves members of the OR profession. And, the professional society to which we belong - the Institute for Operations Research and the Management Sciences (INFORMS) - considers us to be valued members of the profession.

How did this distinction come to be made? I believe it has to do with the expansion of the scope of the problems that operations researchers and policy analysts have had to deal with. In the beginning, OR techniques were applied to problems in which there were few parameters and a clearly defined single objective function to be optimized (e.g., the design of a new weapon system or the placement of radar installations). Gradually, the problems being analyzed became broader and the systems more complex. Health, housing, transportation, and criminal justice policies were being analyzed. Single objectives (e.g., cost minimization or single variable performance maximization) were replaced by the need to consider tradeoffs among multiple (and conflicting) objectives (e.g., the impacts on health, the economy, and the environment), and the distributional impacts on different social or economic groups. Non-quantifiable and subjective considerations had to be considered in the analysis. Optimization was replaced by satisficing, which means finding an acceptable or satisfactory solution to a problem, instead of an optimal solution. And, most important to today's topic, uncertainty became a more important element in the analysis.

I would now like to take you through the four levels of uncertainty that we defined in our 2003 paper, and link them with the tools that policy analysts and operations researchers have used to deal with them.

3. LEVEL 1 - LITTLE UNCERTAINTY (WE KNOW IT ALL)

Location	**Level 1**
Context (X)	A clear enough future 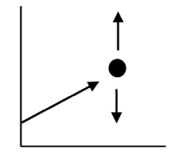 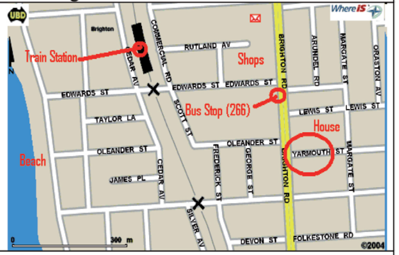
System Model (R)	A single (deterministic) system model
System Outcomes (O)	A point estimate for each outcome
Weights on Outcomes (W)	A single set of weights

In dealing with problems associated with *Level 1 uncertainty*, it is assumed that uncertainty is not an important issue. There are many situations in which this assumption is reasonable and optimization tools are appropriate. These are generally situations involving short-term planning in which the system of interest is well defined and it is reasonable to assume that historical data can be used as predictors of the future.

One of the fundamental methods of OR - linear programming - is an example of a way to deal with Level 1 uncertainty. In this case, a model is used to find the 'optimal' policy, and sensitivity analysis on the model's parameters is used to explore how sensitive the policy results are to the assumptions about the future, the model, and the objectives.

This is sometimes called a 'predict-and-act' approach. (A colleague of mine calls it the 'plan and pray' approach.) In this approach, the (often implicit) assumption is that the future will continue to look significantly like the past; the future world will be structurally more or less the same as the current world - perhaps more populated, richer, more polluted - but, essentially the same. Unfortunately, there is no particular reason why the future should look like the past. By making this assumption, we do not solve the uncertainty problem, we merely sweep it under the rug, often with serious consequences. (The 'pray' part of 'plan and pray' means that you pray that the assumptions on the basis of which you made the plan actually get realized.) The resulting policy may be called 'optimal', but its optimality is dependent on the correctness of the underlying assumptions.

4. LEVEL 2 - WE KNOW THE PROBABILITIES

Location	Level 2
Context (X)	Alternate futures (with probabilities) 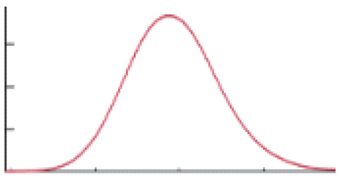 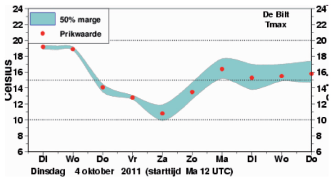
System Model (R)	A single (stochastic) system model
System Outcomes (O)	A confidence interval for each outcome
Weights on Outcomes (W)	Several sets of weights, with a probability attached to each set

In the case of Level 2 uncertainties, it is assumed that the system model or its inputs can be described probabilistically (e.g., travel demands for an airline, or river flows for a water management system) or that there are a few alternative futures that can be predicted well enough (and to which probabilities can be assigned). The system model includes parameters describing the stochastic - or probabilistic - properties of the underlying system. In this case, the model can be used to estimate the probability distributions of the outcomes of policies for these futures. A preferred policy can be chosen based on the outcomes and the associated probabilities of the futures (i.e., based on 'expected outcomes' and levels of acceptable risk). The tools of probability and statistics can be used to solve problems involving Level 2 uncertainties.

An example from my career can be used to illustrate how useful such tools can be in designing policies. Soon after I received my Ph.D.in Operations Research from Cornell University, I went to work for the RAND Corporation in New York City. Between 1970 and 1975, I worked on the RAND Fire Project (The RAND Fire Project, 1979). This project, which was carried out for the Fire Department of New York, was the most extensive research project ever carried out in the area of fire department deployment analysis. Using 1967-1969 data on fire alarms, we predicted alarm rates and the expected proportion of serious fires for each street alarm box in the Bronx. We then compared the predictions to actual 1970 data. For example, as shown in Table 1, we predicted for Box 2277 that less than 0.5 percent of all alarms would be structural fires, while for Box 2209 we predicted almost 32 percent. In 1970, both of these boxes had about

the same number of alarms. In both cases, the predictions were quite close to the actual results. Box 2277 had no structural fires and therefore no serious fires (major incidents requiring at least two ladder companies) while Box 2209 had 25 structural fires, 12 of which were serious. A striking aspect of this example is that both alarm boxes are on the same street, and are about three blocks apart. Under the traditional dispatching policy, if an alarm had been received from Box 2277, all of the closest companies (3 engine companies and 2 ladder companies) would have been dispatched to that alarm. If, in the next few minutes, another alarm had been received from Box 2209, none of the closest units would have been available to respond to the alarm, and companies from further away would have had to be sent. This type of probabilistic analysis led to a new dispatching policy, called adaptive response. The essence of adaptive response was to send fewer units to boxes like 2277, which have a small chance of signaling a serious fire, and more units to boxes like 2209, which have a high chance.

Alarm Box Number	Predicted % Structural, from 1967-1969 Data	Total Alarms (1970)	Structural Fires (1970)
2277	0.4	96	0
2209	31.8	94	25

Table 1: Using Historical Data to Predict the Probability of a Structural Fire

Note that this policy is based on an assumption of stationarity in the alarm patterns, which held very well between 1967 and 1970. However, such a policy would not work well if the pattern of alarms were changing from year to year.

5. LEVEL 3 - FROM PROBABILITIES TO PLAUSIBILITIES

Location	Level 3
Context (X)	A multiplicity of plausible futures 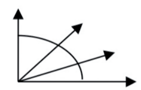 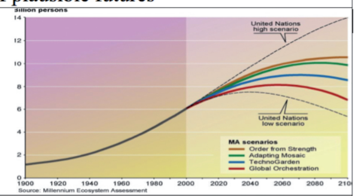
System Model (R)	Several plausible system models, with different structures
System Outcomes (O)	A known range of outcomes
Weights on Outcomes (W)	A known range of weights

Level 3 uncertainties involve situations in which there are a multiplicity of plausible futures, system models, outcomes, or weights, and probabilities cannot be assigned to them - so the tools of neither Level 1 nor Level 2 are appropriate. When faced with a level of uncertainty in which the predict-and-act approach visibly fails, policy analysts usually opt for scenario analysis, or 'what-if' policy analysis (see, for example, (van der Heijden, 1996)). The core of this approach is that the future can be predicted well enough to identify policies that will produce favorable outcomes in one or more specific plausible future worlds. The future worlds are called *scenarios*. Policy analysts use best-estimate models (based on the most up-to-date scientific knowledge) to examine the consequences that would follow from the implementation of each of several possible policies in each scenario. The 'best' policy is the one that produces the most favorable outcomes across the scenarios. (Such a policy is called a *robust* policy.) A scenario does not predict what will happen in the future; rather it is a plausible description of what can happen. The scenario approach assumes that, although the likelihood of the future worlds is unknown, the range of plausible futures can be specified well enough to identify a (static) policy that will produce acceptable outcomes in most of them.

The benefits from using scenarios in policy analysis are threefold. First, it helps us to deal with situations in which there are many sources of uncertainty. Second, it allows us to examine the "what ifs" related to scenario uncertainties. It suggests ways in which the system could change in the future, and allows us to examine the implications of these changes. Finally, scenarios provide a way to explore the implications of Level 3 uncertainties for policymaking by identifying possible future problems and identifying (static) robust policies for dealing with the problems.

From an analytic perspective, however, the scenario approach has several problems. The first problem is deciding which external forces to include in the scenarios. Typically, these forces are decided upon by experts. However, in the face of uncertainty, no one is in a position to make this judgment. A second problem is that we have no way to know whether the range of futures provided by the scenarios covers all, 95%, or some other percentage of the possible futures. Thus, even if we choose a policy that performs well in our scenarios, we have no idea whether this policy will perform well in the real future or not. We can only say that the policy will perform well in the future if the future turns out to resemble one of the futures we have included in our scenarios. A third problem with this approach has to do with the large range (and often even contradictory directions) in the performance estimates generated by the scenarios. If the uncertainty included in this range is large, policymakers often fall back

on the do-nothing approach, with the following sort of reasoning - we do not have sufficient information to make a decision at this time. This is probably the worst possible outcome - when the level of uncertainty is high, and the potential consequences are large, it is imperative that policymakers act (in a wise way) rather than wait.

6. LEVEL 4 - THE FUTURE IS UNKNOWABLE

Location	Level 4
Context (X)	Unknown future 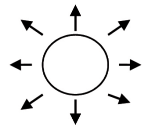 
System Model (R)	Unknown system model; know we don't know
System Outcomes (O)	Unknown outcomes; know we don't know
Weights on Outcomes (W)	Unknown weights; know we don't know

Level 4 uncertainty represents the deepest level of recognized uncertainty; in this case, what is known is only that we do not know. This type of uncertainty is increasingly becoming a common feature of life, because catastrophic, unpredicted, surprising, but painful events seem to be occurring more often. Nassim Nicholas Taleb (2007) calls these events "Black Swans". He defines a Black Swan event as one that lies outside the realm of regular expectations (i.e., "nothing in the past can convincingly point to its possibility"), carries an extreme impact, and is explainable only after the fact (i.e., through retrospective, not prospective, predictability). One of the most dramatic recent Black Swans is the concatenation of events following the 2007 subprime mortgage crisis in the United States. The mortgage crisis (which some had forecast) led to a credit crunch, which led to bank failures, which led to a deep global recession in 2009, which was outside the realm of most expectations. Another recent Black Swan was the level 9.0 earthquake in Japan in 2011, which led to a tsunami and a nuclear catastrophe, which led to supply chain disruptions (e.g., for automobile parts) around the world.

Some of my colleagues at RAND have defined deep uncertainty as "the condition in which analysts do not know or the parties to a decision cannot agree upon

(1) the appropriate models to describe interactions among a system's variables, (2) the probability distributions to represent uncertainty about key parameters in the models, and/or (3) how to value the desirability of alternative outcomes" (Lempert, et al., 2003).

An example of why acknowledging uncertainty and dealing with it is of great importance - but, not with predictive models - is the experience of the financial crisis that gripped the world in 2008-2009. The speed and the severity of the decline in world economies was unprecedented, but policymakers did not see it coming and were unprepared to deal with it. As Alan Greenspan admitted in October 2008: "I found a flaw in the model that I perceived is the critical functioning structure that defines how the world works ... I was shocked, because I had been going for 40 years or more with very considerable evidence that it was working exceptionally well" (Hearing Before the Committee on Oversight and Government Reform, 2008, p. 46). What he was saying is that he had been using Level 2 approaches to deal with Level 4 uncertainty. But, using predictive models based strictly on past statistics and trends is like driving a car while looking only into the rear view mirror.

Another example of uncertainty and the inadequacy of predictive models is climate change. Climate change research is plagued by imperfect and incomplete understanding about the functioning of natural (environmental) phenomena and processes, about how changes in these phenomena and processes translate into increases in global temperatures, and the economic and social consequences of such an increase in temperature. For a long time, the presence of these uncertainties allowed the very existence of global climate change to be denied. Now, the uncertainty as to whether climate change is taking place has been largely removed (Stern, 2006). There is still, however, considerable uncertainty about:

- The magnitude of climate change (there are a whole range of future scenarios that describe very different increases in average temperatures);
- The increased extremes in weather behavior that climate change might cause (e.g., this year's record heat in the Southwest United States, record numbers of tornadoes in the Midwest, driest spring and the wettest summer in the Netherlands);
- The speed of climate change (which determines how quickly policy actions need to be taken);
- What this means for specific areas and regions (the effects of climate change are potentially larger for low-lying countries like Bangladesh and the Netherlands);

- What should be done to mitigate climate change and its adverse consequences (because there is a lack of knowledge about the costs and benefits of different alternatives for protecting ourselves from the adverse conse quences of climate change).

How can policymakers develop policies to protect lives and property in the face of such uncertainties?

The first part of the answer is to not ignore the uncertainty. Ignoring uncertainty could lead to large adverse consequences for people, countries, and the earth, and policymakers have an interest in minimizing the possibility of such adverse consequences happening. The challenge for enlightened policymaking is to develop other, innovative approaches to handle these uncertainties. An approach is needed that adapts to the future course of events and fully exploits knowledge that becomes available as time proceeds.

Over the last few years, I have worked with several colleagues and Ph.D. students to operationalize an approach to making policies under deep uncertainty that we now call Dynamic Adaptive Policymaking (DAP). I first mentioned DAP as a possible approach to handle deep uncertainty in my Inaugural speech eleven years ago. The analysis and choice of an adaptive policy requires a new process for policymaking and policy implementation that explicitly takes into account the uncertainties and dynamics of the problem being addressed.

The basic concept of DAP is easy to explain. It is analogous to the approach used in guiding a ship through a long ocean voyage. The goal - the end point - is set at the beginning of the journey. But, along the way, unpredictable storms and other traffic - or even icebergs - may interfere with the original trajectory. So, the policy - the specific route - is changed along the way. It is understood before the ship leaves port that some changes are likely to take place - and contingency plans may have already been formulated for some of the unpredictable events. The important thing is that the ultimate goal remains unchanged, and the policy actions implemented over time remain directed toward that goal. An adaptive policy would include a systematic method for monitoring the environment, gathering information, implementing pieces of the policy over time, and adjusting and re-adjusting to new circumstances. The policies themselves would be designed to be incremental, adaptive, and conditional.

DAP can be divided into two phases: a policy design phase, and a policy implementation phase. As shown in Figure 2, the policy design phase consists of five steps - one step (Step I) that sets the stage for policymaking, three steps

(Steps II, III, and IV) for designing the portion of the adaptive policy that gets implemented initially (at time t = 0), and one step (Step V) that designs the portions of the adaptive policy that *may* be implemented in the future (at unspecified times t > 0). The implementation phase then consists of two parts - implementation of the portions of the policy that get implemented at time t=0, and adaptation of the initial policy whenever it is needed. But, this adaptation has already been planned for, and a monitoring system has been set up to warn of the need for adaptation.

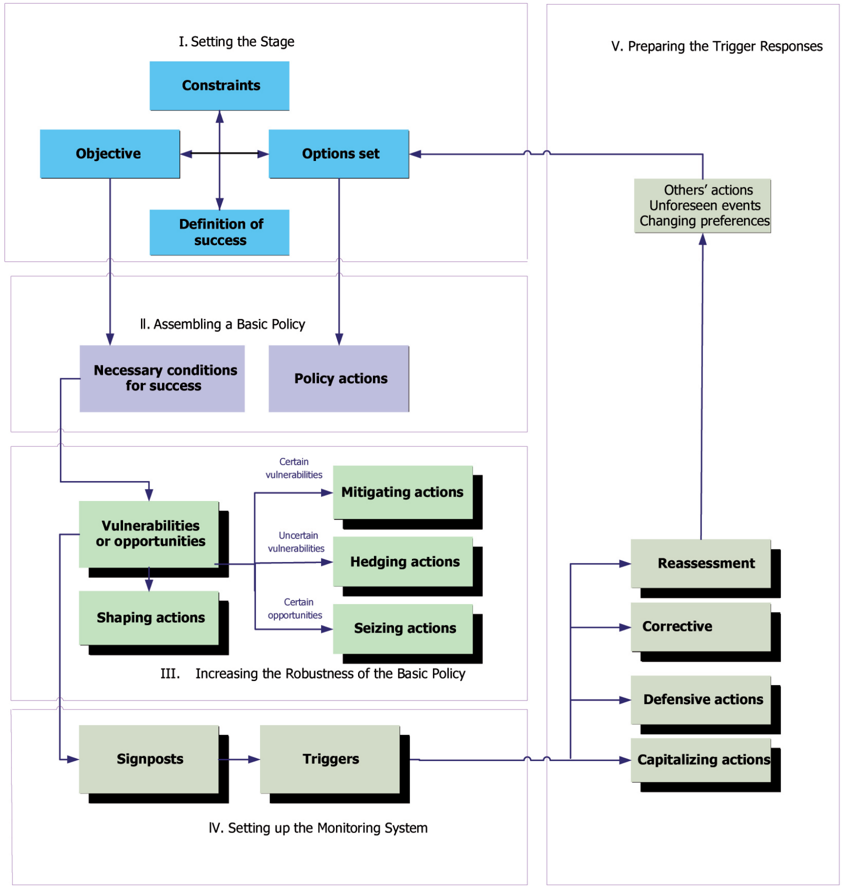

Figure 2 - Steps in designing a dynamic adaptive policy [based on W.E. Walker, S.A. Rahman, J. Cave (2001)].

The Design Phase: Steps in Designing a Dynamic Adaptive Policy

I will not go through all of the steps in detail. But, I will try to give you a feel for what they are. And, I will use a highly simplified example of planning for the possible expansion of a large airport close to a built-up area for the long-term future.

The first and the second steps are basically the same as those that are carried out in designing a static policy using the traditional policy analysis process. Step I involves the specification of the system boundary and the objectives, constraints, and available policy options. This specification should lead to a definition of success, i.e. the specification of the desired outcomes. This is a critical element of DAP, which is usually not specified in a traditional policy analysis - understanding what 'success' means enables the analysts to identify ways to keep the policy from failing.

In this simple case, the system is the airport, and the objective is to improve its capacity to handle increased demands (which are highly uncertain). The major constraints on the policy are costs and public acceptance. Success means having a good match between supply and demand - not too much capacity, which would mean a lot of unused capacity; but not too little capacity, which would lead to delays in take-offs and landings.

In Step II, a *basic policy* is assembled. This step involves (a) the specification of a promising basic policy and (b) the identification of the conditions needed for the basic policy to succeed. These conditions will be used in Step III to set up a monitoring system to provide advance warning in case conditions change and the policy might fail. For the airport case, assume that the basic policy is to build a new runway. Two conditions for the success of the new runway might be that demand continues to grow and that the extra aircraft noise generated does not bring strong protests.

In *Step III* of the DAP process, the actions to be taken immediately (i.e., at time t = 0) to enhance the chances of success of the policy are specified. This step is based on identifying in advance the vulnerabilities associated with the basic policy, and specifying actions to be taken in anticipation. *Vulnerabilities* are external developments that could degrade the performance of the policy so that it is no longer successful. In short, what we are doing is asking 'how can the basic policy fail?', and then designing ways to prevent it from failing.

Scenarios are used in this step and in Step IV; but they are used in a different way from the way they are used in dealing with Level 3 uncertainty. They are used to identify the ways in which the basic policy could go wrong (i.e., not lead to success). In DAP, since we are looking for changes in the world that can make the basic policy fail, the scenarios should differ from the present in major ways. For example, there should be some very negative scenarios. People tend to view very negative scenarios as implausible and reject them out of hand. Nevertheless, they are crucial to an adaptive policy; having thought about a situation (no matter how implausible) in advance allows contingency plans to be formulated so that they are ready to be implemented in the (however unlikely) event they are needed.[1] So, as many Black Swans as possible should be identified, in order to 'be prepared' in case one of them actually occurs. In the airport case, demand for air transport is one of the key scenario variables. There could be a sharp decrease in demand, for example due to a financial crisis. This would make the policy fail. But, there could be a sharp increase in demand, which could lead to unacceptable delays in takeoffs and landings, which would also make the policy fail. I will deal with this vulnerability when I discuss Step IV.

Another vulnerability of the basic policy is resistance from people living around the airport because of the noise from the anticipated additional flights. This vulnerability is fairly certain. So, at the same time as the new runway is agreed upon (at time t=0), it would be wise to offer financial compensation to residents in the high noise zone to enhance the chances of success of the basic policy. (This would be one of the 'mitigating actions' according to Fig. 2.)

Steps IV and V set up the monitoring system and prepare contingency plans for changing the policy in response to changes in the world. They prepare actions to be taken if needed to guarantee the policy's progress and success. In these steps, *signposts* are identified that specify information that should be tracked, and critical values of signpost variables (called *triggers*) are specified beyond which actions to change the policy should be implemented to ensure that the resulting policy keeps moving the system in the right direction and at a proper speed. The starting point for the identification of signposts is the set of vulnerabilities specified in Step III.

In the airport case, it is possible that the increases in demand are much greater than expected. This would lead to unacceptable delays and airlines might decide to shift flights (or even their hubs) to other airports, which would lead to failure of the plan. In preparation, as part of the policy design phase, plans could be made to shift specific types of flights to surrounding airports (e.g., all-cargo flights or flights by low cost carriers). Making these plans would not be expensive, and they may never be needed. But, if the conditions warranted them, the plans would be there and could be implemented quickly at the appropriate time (specified by the trigger), thus saving the basic policy.

[1] Thomas Schelling, in a Foreward to Wohlstetter's (1962) study Pearl Harbor: Warning and Decision, wrote "There is a tendency in our planning to confuse the unfamiliar with the improbable. The contingency we have not considered seriously looks strange; what looks strange is thought improbable; what is improbable need not be considered seriously."

The Implementation Phase

Once the basic policy and additional actions are agreed upon, the entire adaptive policy is implemented. In this phase, the actions to be taken immediately (from Step II and Step III) are implemented and a monitoring system (from Step IV) is established. Then time starts running, signpost information related to the triggers is collected, and, when triggered, policy adaptation actions (from Step V) are implemented.

To sum up, DAP helps to develop robust plans by accepting uncertainty and acknowledging that we cannot know the future (even probabilistically). The approach calls for implementing a basic policy based on what we know today, and constructing a system for monitoring the (unpredictable) developments that could potentially affect the effectiveness of the chosen policy. The resulting policy is dynamic; the element of time and the possibility of learning are explicitly taken into account by the policy. Whereas other approaches are based on the notion that policymaking is a discrete one-time event and that the resulting policy is static, dynamic adaptation is explicitly defined as a continuous process in time that involves monitoring and making pre-specified changes to existing policy in response to unpredictable developments.

The DAP framework offers several advantages over other approaches. Most important of these are (1) it does not ignore uncertainty; it acknowledges that we cannot know the future and bases policy on this assumption, and (2) it institutionalizes the process of ex-post policy evaluation and monitoring. As Taleb (2007) has written: "It is often said that 'is wise he who can see things coming.' Perhaps the wise one is the one who knows that he cannot see things far away."

7. FINAL WORDS

		LEVEL					
		Complete Certainty	Level 1	Level 2	Level 3	Level 4	Total ignorance
LOCATION	Context (X)		A clear enough future 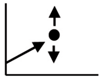	Alternate futures (with probabilities) 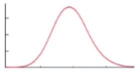	A multiplicity of plausible futures 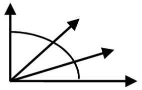	Unknown future 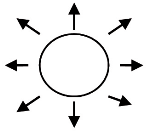	
	System Model (R)		A single (deterministic) system model	A single (stochastic) system model	Several system models, with different structures	Unknown system model; know we don't know	
	System Outcomes (O)		A point estimate for each outcome	A confidence interval for each outcome	A known range of outcomes	Unknown outcomes; know we don't know	
	Weights on outcomes (W)		A single set of weights	Several sets of weights, with a probability attached to each set	A known range of weights	Unknown weights; know we don't know	

To expand on the simple airport planning example given above, several more complex, real-world DAP examples have been developed. Kwakkel et al. (2010) present a more realistic case of airport strategic planning, based on the current challenges of Amsterdam Airport Schiphol. Various other areas of application of DAP have also been explored, including flood risk management in the Netherlands in light of climate change (Rahman et al., 2008) and policies with respect to the implementation of innovative urban transport infrastructures (Marchau et al. 2008), congestion road pricing (Marchau et al., 2010), intelligent speed adaptation (Agusdinata, et al., 2007), and 'magnetically levitated' (Maglev) rail transport (Marchau et al., 2010).

However, dynamic adaptation has a long way to go before it becomes commonplace in public policymaking. More research is required before this will happen.

First, its validity and efficacy need to be established. This will be difficult to do since, as Dewar et al. (1993) have pointed out, "nothing done in the short term can 'prove' the efficacy of a planning methodology; nor can the monitoring, over time, of a single instance of a plan generated by that methodology, unless there is a competing parallel plan." Nevertheless, at TBM we are beginning to gather evidence through a variety of methods, including gaming and computational experiments (see, for example, Kwakkel, et al., forthcoming). Also, the costs and benefits of dynamic adaptation measures compared to traditional policymaking approaches need to be studied. We are working on this, too. Finally, the implementation of dynamic adaptation will require significant institutional and governance changes, since some aspects of these policies are currently not supported by laws and regulations (e.g., the implementation of a policy triggered by an external event).

New approaches to dealing with Level 4 uncertainties, however, are gradually being accepted as valid - and, indeed, necessary. For example:

- In the financial area, as the example of Allan Greenspan indicates, financial planners have already seen that their standard models based on statistics and probabilities are insufficient to deal with the recent Black Swans that have caused huge swings in stock prices - such as the subprime mortgage and the debt ceiling debacles in the United States and the Greek debt crisis in Europe.
- Defense planners are beginning to understand that current defense planning methodologies need to be changed. I recently received a draft report from a respected defense planning organization that says "our current defence planning methodologies which still focus primarily on ... trends and drivers that we presume to 'know'... insufficiently take into account the true nature of today's deep uncertainty."
- And, in the area of water management and flood safety, a prepublication copy of a report from the National Research Council of the U.S National Academy of Sciences notes that water management systems have traditionally been designed based on the assumption of stationarity (which means that the variability in their statistical patterns does not change over time, so that flood protection norms can be confidently based on past statistics). But, it concludes that "continuing to use the assumption of stationarity in designing water management systems is no longer practical or defensible" (National Research Council, Committee on Hydrologic Science (2011), p. 8).

I am amazed to see how much the views on uncertainty have changed in the eleven years since my Inaugural speech. In the year 2000, the idea of adaptive policies was seen as quite revolutionary. Now, it is coming to be accepted as necessary. But, as I have already indicated, further research is needed before it can become the norm in policymaking. The ball is now in motion - my students, and others will have to pick up the ball and solve the many outstanding problems of developing and implementing adaptive policies. However, I am fairly certain that some elements of dynamic adaptive policies will become more commonplace in the future.

8. THANKS

I wish I had time to thank by name the literally hundreds of persons who have helped me to achieve the successes I have achieved in my life. I truly consider myself a dwarf who has been lucky enough to stand on the shoulders of giants.

My father and mother, who put me on this path, supported me, and helped me to appreciate the importance of intellectual curiosity and education.

My children, who may have suffered some fatherly deficits while I was devoting an inordinate amount of time at the beginning of my career, but who have grown into strong, loving, successful human beings in spite of this neglect (or maybe because of it).

The teachers, who guided me and set me on the path that led to this day and who appeared in my life to provide direction at the most critical points - when I had no idea what to do next.

My colleagues during my more than 37 years at RAND, who challenged me, inspired me, and collaborated with me on research. Without them, I would have accomplished much less, would have made less of a difference in the world, and would have been significantly less fulfilled as a person.

And, last, but not least, my colleagues at TU Delft. I have been associated in one way or another with the Faculty of Technology, Policy and Management since its conception in the late 1980s. I am grateful to those, like Wil Thissen and Cees Boerman, who trusted me to give them advice about what a school of system engineering, policy analysis, and management might look like, and why it might be important. I am grateful to those, like Henk Sol and Hugo Priemus, who gave me important roles and responsibilities in the faculty. And I am grateful to my colleagues and students who over the years have challenged me, broadened my perspectives, and helped me to stay young, intellectually alive, and productive.

Ik heb gezegd.

REFERENCES

Ackoff, R.L. (1974). Redesigning the Future: a Systems Approach to Societal Problems, John Wiley & Sons, New York.

Agusdinata, D.B., V.A.W.J. Marchau, and W.E. Walker (2007). "Adaptive Policy Approach to Implementing Intelligent Speed Adaptation", IET Intelligent Transport Systems (ITS), Vol. 1, No. 3, pp. 186-198.

Dewar, J. A., C.H. Builder, W.M. Hix, and M.H. Levin (1993). Assumption-Based Planning: A Planning Tool for Very Uncertain Times, Report MR-114-A, RAND, Santa Monica, CA.

Hearing Before the Committee on Oversight and Government Reform (2008). The Financial Crisis and the Role of Federal Regulators, Serial No. 110-209, U.S. Government Printing Office, Washington, D.C.

Kwakkel, J.H., W.E. Walker, and V.A.W.J. Marchau (2010). "Adaptive Airport Strategic Planning", European Journal of Transport and Infrastructure Research, Vol. 10, No. 3, pp. 249-273.

Johnson, M.P. (ed.) (2012). Community Based Operations Research: Decision Modeling for Local Impact and Diverse Populations, Springer, New York.

Kwakkel, J. H., W. E. Walker, and V.A.W.J. Marchau (forthcoming). "Assessing the Efficacy of Adaptive Planning of Infrastructure: Results From Computational Experiments", Environment and Planning B (Planning and Design).

Lempert, R.J., S.W. Popper, and S.C. Bankes (2003). Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis, Report MR-1626-RPC, RAND, Santa Monica, California.

Marchau, V.A.W.J., W.E. Walker, and R. van Duin (2008). "An Adaptive Approach to Implementing Innovative Urban Transport Solutions", Transport Policy, Vol. 15, Issue 6, pp. 405-412.

Marchau, V.A.W.J., W.E. Walker, and G.P. van Wee (2010). "Dynamic Adaptive Transport Policies for Handling Deep Uncertainty", Technological Forecasting & Social Change, Vol. 77, No. 6, pp. 940-950.

National Research Council, Committee on Hydrologic Science (2011). Global Change and Extreme Hydrology: Testing Conventional Wisdom, The National Academies Press, Washington, D.C.

Rahman, S.A., W.E. Walker and V.A.W.J. Marchau (2008). Coping with Uncertainties About Climate Change in Infrastructure Planning - An Adaptive Policymaking Approach, Report for Raad voor Verkeer en Waterstaat, Ecorys Research and Consulting, Rotterdam.

Stern, N., (2006). The Economics of Climate Change, the Stern Review, Cabinet Office, HM Treasury, London.

Taleb, N.N. (2007). The Black Swan: The Impact of the Highly Improbable, Random House, New York.

The RAND Fire Project (W.E. Walker, J.M. Chaiken, E.J. Ignall (eds.))(1979). Fire Department Deployment Analysis: A Public Policy Analysis Case Study, Elsevier North-Holland, Inc., New York.

Van der Heijden, K. (1996). Scenarios: The Art of Strategic Conversation, John Wiley & Sons, Ltd, Chichester, England.

Walker, W.E. (2000a). "Policy Analysis: A Systematic Approach to Supporting Policymaking in the Public Sector", Journal of Multicriteria Decision Analysis, Volume 9, No. 1-3, pp. 11-27.

Walker, Warren E. (2000b). Uncertainty: The Challenge for Policy Analysis in the 21st Century, TU Delft (ISBN: 90-5638-074-5), Delft.

Walker, W.E., S.A. Rahman, and J. Cave (2001). "Adaptive Policies, Policy Analysis, and Policymaking", European Journal of Operational Research, Vol. 128, pp. 282-289.

Walker, W.E., P. Harremoës, J. Rotmans, J.P. van der Sluijs, M.B.A. van Asselt, P. Janssen, and M.P. Krayer von Krauss (2003). "Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support", Integrated Assessment, Vol. 4, No. 1, pp. 5-17.

Wohlstetter, R. (1962). Pearl Harbor: Warning and Decision, Stanford University Press, Palo Alto, CA.

Faculteit Techniek, Bestuur en Management
Jaffalaan 5
2628 BX Delft

Tel: +31 (0)15 27 83408

www.tudelft.nl