

**AFRL-IF-RS-TR-2005-338**  
**Final Technical Report**  
**September 2005**



# **PRE-CONFLICT ANTICIPATION AND SHAPING (PCAS): MODELS-2-SHAPING INTEGRATION**

**Evolving Logic**

**Sponsored by**  
**Defense Advanced Research Projects Agency**  
**DARPA Order No. T430/00**

*APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.*

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

**AIR FORCE RESEARCH LABORATORY**  
**INFORMATION DIRECTORATE**  
**ROME RESEARCH SITE**  
**ROME, NEW YORK**

## STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2005-338 has been reviewed and is approved for publication

APPROVED: /s/

FRANCES A. ROSE  
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor  
Information Grid Division  
Information Directorate

# REPORT DOCUMENTATION PAGE

Form Approved  
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

<b>1. AGENCY USE ONLY (Leave blank)</b>	<b>2. REPORT DATE</b> SEPTEMBER 2005	<b>3. REPORT TYPE AND DATES COVERED</b> Final Oct 04 – Mar 05
---	---	--

<b>4. TITLE AND SUBTITLE</b> PRE-CONFLICT ANTICIPATION AND SHAPING (PCAS): MODELS-2-SHAPING INTEGRATION	<b>5. FUNDING NUMBERS</b> C - FA8750-04-C-0215 PE - 62303E PR - T438 TA - 00 WU - 01
--	---

<b>6. AUTHOR(S)</b> Steven Bankes, Robert J. Lempert and Steven W. Popper
--

<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Evolving Logic, Incorporated 21280 Colina Drive Topanga California 90290	<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>  N/A
---	--

<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Defense Advanced Research Projects Agency AFRL/IFGA 3701 North Fairfax Drive Arlington Virginia 22203-1714	<b>10. SPONSORING / MONITORING AGENCY REPORT NUMBER</b>  AFRL-IF-RS-TR-2005-338
--	---

<b>11. SUPPLEMENTARY NOTES</b>  AFRL Project Engineer: Frances A. Rose/IFGA/(315) 330-2038/ Frances.Rose@rl.af.mil
--

<b>12a. DISTRIBUTION / AVAILABILITY STATEMENT</b> APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.	<b>12b. DISTRIBUTION CODE</b>
---	-------------------------------

<b>13. ABSTRACT (Maximum 200 Words)</b> The project investigated techniques of computational support for devising shaping strategies and "fusing", or jointly using, multiple different models representing relevant information and knowledge. A demonstration problem involving actions the US might take regarding Pakistan's Northwestern Territories and utilizing an assembled suite of models was developed. The models were placed in Evolving Logic's Computer Assisted Reasoning system (CARs(tm)) software environment and provided the basis for compound computational experiments to demonstrate that the use of multiple models (both alternative model variants, and complementary models of different type) can provide more insight and better decisions than any one model could and that the suite of models together with the technology can provide a basis for devising shaping strategies. The analytical outcome suggests that a useful definition of a shaping strategy is one involving near-term actions that are robust to the uncertainties and satisfies a series of goals, in both the short and long term.
--

<b>14. SUBJECT TERMS</b> Uncertainty Management, Robust Decision Making, Strategic Planning, Hedging And Shaping Strategies, Exploratory Modeling	<b>15. NUMBER OF PAGES</b> 29	<b>16. PRICE CODE</b>
--	----------------------------------	-----------------------

<b>17. SECURITY CLASSIFICATION OF REPORT</b>  UNCLASSIFIED	<b>18. SECURITY CLASSIFICATION OF THIS PAGE</b>  UNCLASSIFIED	<b>19. SECURITY CLASSIFICATION OF ABSTRACT</b>  UNCLASSIFIED	<b>20. LIMITATION OF ABSTRACT</b>  UL
--	---	--	---

## Table of Contents

<b>1. SUMMARY .....</b>	<b>1</b>
<b>2. INTRODUCTION.....</b>	<b>2</b>
<b>3. APPROACH.....</b>	<b>5</b>
<b>4. SUPERIOR DECISIONS FROM MULTIPLE ALTERNATIVE MODELS.....</b>	<b>8</b>
<b>5. RESULTS AND DISCUSSION .....</b>	<b>18</b>
<b>6. CONCLUDING REMARKS .....</b>	<b>22</b>
<b>7. LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS.....</b>	<b>24</b>

## List of Figures

<b>FIGURE 1: SOCIOECONOMIC MODELS CAN “PREDICT” ACCURATELY IN MANY CIRCUMSTANCES .....</b>	<b>3</b>
<b>FIGURE 2: CAUTION IS REQUIRED IN USING SUCH “PREDICATIONS” .....</b>	<b>4</b>
<b>FIGURE 3: CONSIDER MULTIPLE MODELS AND UNCERTAINTIES AFFECTING FUTURE SOCIOECONOMIC FACTORS.....</b>	<b>8</b>
<b>FIGURE 4: CARS™ SOFTWARE ASSISTS IN REASONING OVER MULTIPLE MODELS .....</b>	<b>9</b>
<b>FIGURE 5: RAP IDENTIFIES ROBUST STRATEGIES AND KEY UNCERTAINTIES WITH ITERATIVE PROCESS.....</b>	<b>10</b>
<b>FIGURE 6: FIRST STEP: IDENTIFY OPTIMAL STRATEGY FOR BASECASE FUTURE.....</b>	<b>11</b>
<b>FIGURE 7: OPTIMAL STRATEGIES FOR BASECASE XS FOR DIFFERENT COLLIER MODELS .....</b>	<b>12</b>
<b>FIGURE 8: WIDE RANGES OF XS CAPTURE UNCERTAINTY.....</b>	<b>13</b>
<b>FIGURE 9: SIMPLE EXAMPLE OF COMPARING REGRET UNDER TWO SCENARIOS AND UNDER UNCERTAINTY .....</b>	<b>14</b>
<b>FIGURE 10: OPTIMAL BASECASE STRATEGY IS VULNERABLE IN GRIEVANCE MODEL....</b>	<b>15</b>
<b>FIGURE 11: WE CAN DETECT CONDITIONS CAUSING OPTIMAL BASECASE STRATEGY FAILURE IN GRIEVANCE MODEL .....</b>	<b>16</b>
<b>FIGURE 12: STRATEGY EXISTS THAT IS ROBUST OVER ALL MODELS .....</b>	<b>17</b>
<b>FIGURE 13: PLACING BOTH RESULTS ON SAME SCALE .....</b>	<b>18</b>
<b>FIGURE 14: FUSION OF COMPLEMENTARY MODELS INTO A SUITE ALLOWS BETTER UTILIZATION OF THE DATA .....</b>	<b>19</b>
<b>FIGURE 15: SEEKING ACTIONS THAT MEET NEAR-TERM OBJECTIVES CONSISTENT WITH LONG-TERM GOALS .....</b>	<b>20</b>
<b>FIGURE 16: UNDERSTANDING TRADEOFFS AMONG BEST NEAR-TERM STRATEGIES.....</b>	<b>21</b>

## 1. SUMMARY

This technical report describes work performed by Evolving Logic for the Pre-Conflict Anticipation and Shaping (PCAS) seedling effort. This project was intended to investigate the application of techniques Evolving Logic has been developing to issue central to PCAS, in particular computational support for devising shaping strategies and the related problem of “fusing,” or jointly using, multiple different models representing relevant information and knowledge.

In pursuit of these goals, we selected a demonstration problem involving actions the US might take regarding Pakistan’s Northwestern Territories, and assembled a suite of models to help address this problem. These models included six variants of the models the World Bank’s Paul Collier and his colleagues created to model the occurrence of civil violence based on historical data; a re-implementation of the Fund for Peace CAST model that creates 12 indicators of state-failure from textual data such as news feeds; and a model we developed based on an elicitation with former and current government officials of possible strategies for intervention, and their possible effects on Pakistan.

The models were placed in Evolving Logic’s Computer Assisted Reasoning® system (CARs™) software environment and provided the basis for compound computational experiments to demonstrate that the use of multiple models (both alternative model variants, and complementary models of different type) can provide more insight and better decisions than any one model could and that the suite of models together with the technology can provide a basis for devising shaping strategies.

Using the single Collier model that had the maximal likelihood as reported in the original literature, we determined the “optimal” strategy. Testing this strategy on alternative versions of the Collier model revealed potential failure modes of policies recommended by the best model. As the likelihood of the alternative models are only slightly less than that of the “best” model, these failure modes should be considered before committing to a course of action. By using all six alternative models, we can discover (through computer search) strategies that have no failure modes on any of the six models. Thus, the use of all six helps us to discover more robust strategy options. This suggests that a useful definition of a shaping strategy is one involving near-term actions that are robust to the uncertainties and satisfies a series of goals, in both the short and long term. When multiple models are used, more information is brought into the analysis as each model incorporates a difference representation of what is known or a filtering of the data through different assumptions. The approach of discovering strategy options that are robust can be useful in incorporating a suite of models into the analysis, by seeking strategies that are robust to modeling choices. This insight suggests an approach to model fusion, though much work remains to be done on this topic.

Any questions regarding this report should be directed to Steve Bankes at Evolving Logic (solutions@[evolvinglogic.com](mailto:solutions@evolvinglogic.com)).

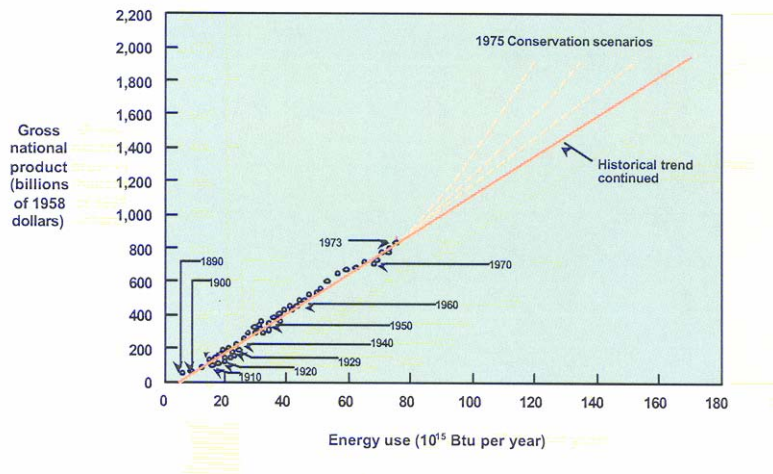
## **2. INTRODUCTION**

Models-2-Shaping was based on the promise of PCAS and the creation of technology that allows the diversity of data, information, and knowledge related to the social and political situation in countries of interest to be brought to bear to inform the decisions of military commanders, and civilian officials. By illuminating the implications of changing circumstance, such technology could provide earlier warning of untoward events in such countries, potentially helping to avoid surprises. Models can illuminate the reasons for conflict and instability and in doing so provide a means for evaluating the consequences of candidate preventative or mitigating shaping strategies.

Models alone will not suffice to achieve the promise of PCAS. Means are needed to use models capturing our understanding to anticipate possible future contingencies such as violent conflicts or social-political instability and to devise strategies that can help shape the future in a more positive direction. This project initiated research into better understanding how this might be accomplished.

Multiple knowledge sources exist that have the potential to usefully inform decisions regarding a country of interest. They include social science models and up-to-the-minute data that provide a dynamic picture of the evolving situation. Better models and better access to data can undoubtedly produce superior insights. Yet, the future will always be significantly uncertain as there are many events that are possible and beyond our control that can profoundly alter the situation. Thus, even the forecasts of the best models will always be vulnerable to surprise.

In seeking means to exploit all available sources of knowledge, we must maintain a balance between knowing as much as possible without allowing this information to blind us to the possibility of surprise. Instead of making decisions tuned to the predictions of the best models, we must seek decisions that, while based on all available information, are designed to be robust to the widest possible range of possible exogenous surprises.

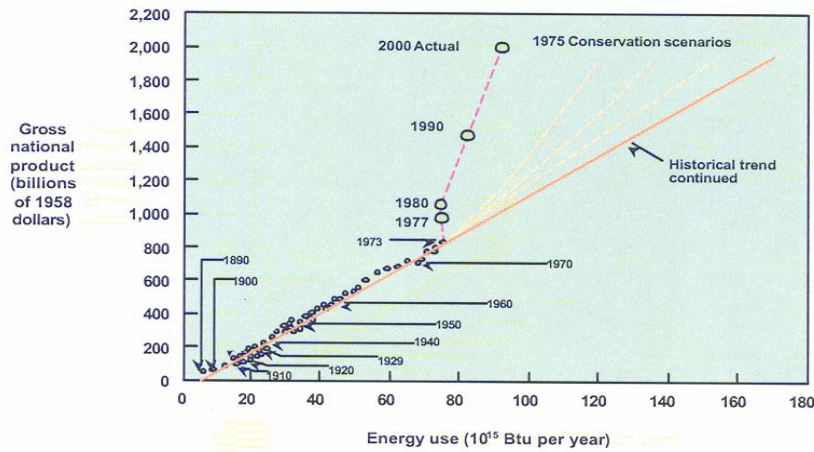


**Figure 1: Socioeconomic models can “predict” accurately in many circumstances**

This graphic, displaying data on US energy consumption, demonstrates the balance we must strive to achieve. Here we see the value models of social systems can bring to decisions. The regression line shows a very strong correlation between GNP and energy use during the 80 year period of 1890 through 1973. Thus, if in 1960 one had been asked to forecast US energy consumption for the next 10 years, this simple model would have allowed one to do so with high accuracy.

In spite of the clear predictive value this model delivers, caution is needed in its use. Suppose one was now in 1973, and again asked to forecast future energy consumption. This graphic shows both the extrapolated trend line, and a fan of what were conceived at the time as plausible forecasts, some of them quite radical. A risk averse utility company might choose to use these more conservative forecasts as a basis for making investments in new production facilities.

However, the next graphic shows the hazard of using even bold-seeming forecasts as the basis for such a decision.



**Figure 2: Caution is required in using such “predications”**

Here the previous graph is augmented with the actual data for years after 1973. As we can see, there was a significant change due at least in part to the oil shocks of the 1970’s, and the rate of increase in energy consumption seems to have been permanently altered. Thus, investments made by utilities using only very conservative forecasts, might well have led them into insolvency.

This then is the balance we must seek. To eschew the use of models altogether would deprive us of the advantage of learning about the regularities in the world that models can reveal. However, to use the forecasts of such models as a basis for decision-making can be hazardous, as such decisions may be vulnerable to the surprises that regularly occur.

What we seek is a means to utilize the insights such models produce without the fragility that relying on forecasts can produce.

The phrase we use for situations where we must balance the desire to use available knowledge with the need to avoid relying on predictions is “deep uncertainty”. Problems are characterized by deep uncertainty when we either do not know – or cannot agree upon -- the correct form of the true model for that system or the correct probability distributions to use to represent our uncertainty or the correct way to value and therefore assess outcomes. Problems involving fragile states and the processes that affect them, such as those of concern to PCAS, are frequently deeply uncertain in this sense. Our goal for this project was to demonstrate the ability to strike a balance between the need to exploit available knowledge with the hazard of believing in predictions made based on that knowledge. We also

demonstrate a different sort of balancing act, that of finding shaping actions that strike a balance between our long-term goals and the requirements of the near-term situation.

The first task of the seedling program was to choose a “vignette” to use as a basis for demonstrating our approach. We chose to look at US counter-terrorism policy towards Pakistan. This choice was motivated by the availability of models and data pertaining to this problem, and by its characteristics. This is a clear example of a deeply uncertain problem where there is debate about the correct models to use and where the choice among policy options is not clear.

Using this vignette as a demonstration problem, we created a suite of models to demonstrate how these can be used to support decision making. We used this suite to discover and assess near-term shaping actions that can satisfy multiple objectives involving multiple time scales, assessed across multiple models.

We went on to demonstrate that this approach is demonstrably superior to the decisions that would be made if we limited ourselves to use any single model from the suite.

### **3. APPROACH**

The main themes of this project were to investigate the use of multiple, possibly contradictory models (Model Fusion) and the use of a suite of models to find shaping strategies that are robust to uncertainty. That is, to use our models to find strategies that have satisfactory performance on long term goals while also satisfying near term constraints across the broad range of alternative futures implied by the uncertainties present.

Our use of multiple models is predicated on two arguments:

— in confronting deep uncertainty, we do not want to make any assumptions that might leave us more vulnerable to surprise. To choose a single model when numerous models have some support and cannot be rejected by our present level of information would not be prudent.

— models addressing different aspects of the problem, or based on different modeling formalisms, can bring different information to bear. In using multiple models we both avoid assuming away uncertainty and allow ourselves access to a greater range of information potentially relevant to the problem.

In devising robust shaping strategies, we must test options across all of these models, and also against the full diversity of objectives, including both near term and long term goals.

“Model fusion” is the use of multiple models to solve a problem. (This is perhaps not the ideal terminology, as we will discuss below.)

Multiple models are useful for a variety of reasons. There may be alternative models that provide differing “opinions”. In the absence of evidence that proves that one dominates the others, all should be consulted in arriving at a decision. Models may exist that represent differing parts of the problem. All such models will need to be used together to get a complete picture. And models can be constructed using a variety of formalisms or with different choices of representation. Typically, each model will be able to utilize a different set of information. Employing models from a range of formalisms can give us access to insights from a wider spread of data and disciplines than could be possible with any single model.

Our approach to the use of multiple models is based on using computation to derive inferences based on the properties of the suite of models. This is in contrast to trying to directly wire all the models together to create a single consensus model. While directly wiring models together can work in engineering applications, for deeply uncertain problems there are typically multiple plausible ways to do the wiring, and so the wiring process will require making unwarranted assumptions, or will result in further increasing the dimensionality of the uncertainty. Neither of these options is helpful. Instead, our technology allows us to treat each of the models as a platform for doing computational experiments and to perform joint experiments consisting of coupled or correlated experiments across multiple models. The specification for the joint experiment and hence each of the individual ones is driven by the problem being solved, which can avoid the problems created by “wiring together” the various models.

Our approach to seeking robust shaping strategies is based on a methodology Evolving Logic has been developing called RAP™ - Robust Adaptive Planning™. The foundation of this approach is the realization that humans deal with deeply uncertain problems all the time. Their approach to dealing with uncertainty is very different from that advocated by traditional decision theory. Instead of specifying the uncertainty precisely and then choosing the option with maximal expected utility, decision makers seek options that are insensitive to the uncertainty. They seek naturally hedged positions where values are not optimized for one particular set of circumstances but are instead intended to do “well enough” across a wide range of plausible futures. We seek to do something similar computationally by defining robust strategies as ones whose performance is acceptable compared to other possible decisions across the widest possible range of possible futures.

We have been experimenting with a variety of methods to seek robust strategies. Important components are measuring the performance of strategies across all the uncertainties, determining which uncertainties are most important in choosing between strategies, and where no completely robust strategy can be found, identifying the tradeoff between alternative strategies in terms of these uncertainties.

Our approach is based on four central principles:

— We utilize ensembles of alternatives and avoid reducing the diversity of options down to a single choice until as late as possible. This means in particular maintaining ensembles of alternative possible models and correspondingly ensembles of alternative possible scenarios.

— We use robustness as opposed to optimality as our means for option selection. This means seeking strategies that satisfice, to use the term coined by the decision theorist, Herbert Simon, across the broadest possible range of plausible scenarios.

— We support the creation of explicitly adaptive strategies that can achieve robustness in part by adjusting their behavior in the future as new information becomes available. This means as a practical matter identifying through analysis the signposts and triggers that would signal such a future adaptive change.

— Finally, we are developing technology that supports highly interactive software embodying these approaches allowing users to collaborate with computers in seeking robust strategies. We want to build systems that participate in human deliberations, and avoid providing “answers” that users must either accept or reject without further interaction.

Evolving Logic has developed exploratory modeling software (CARs™) that provides a powerful foundation for building systems based on the RAP™ principles. This software encapsulates models in a framework that facilitates treating them as platforms for performing large numbers of computational experiments. CARs™ supports interactive search and graphical visualization to allow the user to create, explore, compare, and understand large ensembles of plausible scenarios. We use this technology to create systems that use RAP™ methodology to confront problems with deep uncertainty.

While CARs™ itself does not constitute a decision support system, it provides a host of services that can be useful in creating decision support systems utilizing RAP™ principles. In particular, it provides substantial capabilities for interactive visualization of the behavior of models that are connected to it, supports generation of computational experiments driven by the users reasoning goals, and provides means for navigating among alternative views of a problem.

We use CARs™ capabilities to support RAP-based analysis. CARs™ facilitates searches across models and model inputs seeking strategies with desired properties, as well as searches to discover the vulnerabilities of a given strategy. Where no single strategy that meets all needs can be discovered, CARs™ can also facilitate trade-off analysis to illuminate the choice between alternative candidates.

#### 4. SUPERIOR DECISIONS FROM MULTIPLE ALTERNATIVE MODELS

For this analysis we use six regression models created by Paul Collier of the World Bank and his colleagues. Collier fit a wide variety of such models to a data set that contained the historical data regarding internal conflict across 120 countries over the past forty years. Six of these models had high measures of goodness of fit. Though there is a maximally likely model, differences in likelihood across the six do not provide an unambiguously strong statistical basis for rejecting any of the other five. These six models include four models whose predictors are consistent with a theory that conflict results from opportunity for material gain from overthrowing the government, also known as “greed”, and one model whose predictors are consistent with the view that conflict results from “grievance”. The most likely model used combinations of these predictors.

We use these six models to forecast the probability of conflict in Pakistan in 2015. These forecasts involve assumptions regarding the continuation of observed trends and possible US policy over the intervening years. We can use these models to test the impact of alternative US policies on the probability of conflict in Pakistan.

We demonstrate that a strategy devised to be robust over all six models has better performance than the “optimal” strategy based upon the “best estimate” model.

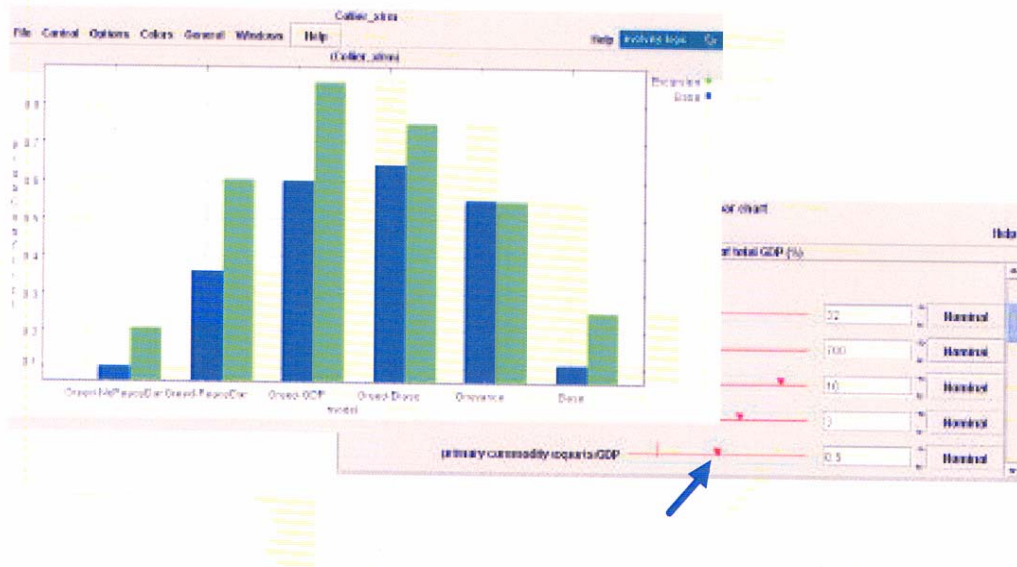
<p><b>Exogenous uncertainties</b></p> <ul style="list-style-type: none"> <li>• Future economic growth rate</li> <li>• Future population growth rate</li> <li>• Change in schooling</li> <li>• Change in social fractionalization</li> <li>• Change in democracy</li> <li>• Year of last civil war (proxy for quality of insurgent network)</li> </ul>	<p><b>policy Levers</b></p> <ul style="list-style-type: none"> <li>• Economic assistance</li> <li>• Cut external funds</li> <li>• Control borders</li> <li>• Military assistance</li> <li>• Fight corruption</li> </ul>
<p><b>Relationships</b></p> <ul style="list-style-type: none"> <li>• Six greed and/or grievance models, fit to full dataset, with statistically indistinguishable predictive power</li> </ul>	<p><b>Measures of merit</b></p> <ul style="list-style-type: none"> <li>• Probability of conflict in 2015</li> </ul>

Figure 3: Consider multiple models and uncertainties affecting future socioeconomic factors

We have found it useful to analyze decision problems into the four quadrants displayed here (XLRM). These four aspects of the decision problem are:

- **eX**ogenous factors- the uncertainties with which any decision must contend,
- **P**olicy Levers - the components (levers) of possible decisions,
- **M**easures - the measures we will use to judge the performance of any given strategy, and
- **R**elationships - the relationships that connect a vector of assumptions about the uncertainties and a strategy (a vector of levers) with measures of outcome.

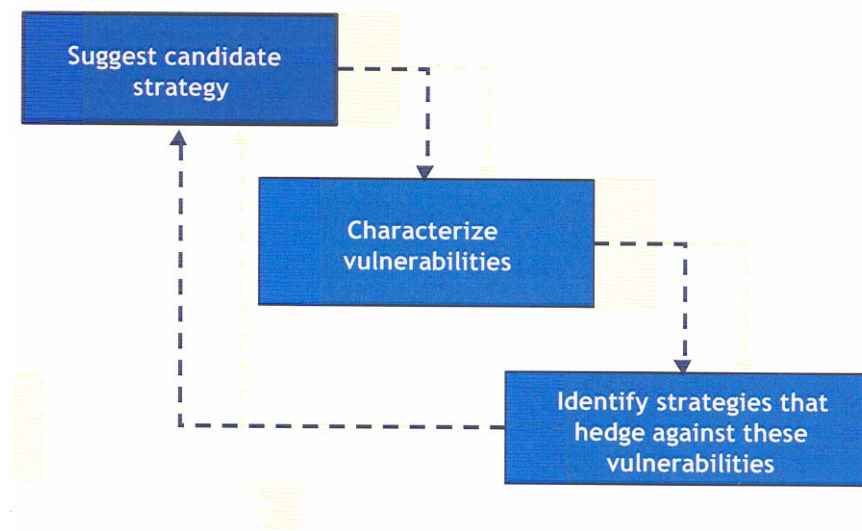
For the analysis we are about the present, the XLRM factors are as shown here. Note that “Year of last civil war” is the name of a predictor in Collier’s original work. We retain this name here for consistency. However, the “year of the last civil war” is not an uncertainty when considering Pakistan in 2005. Instead, this variable can be thought of as a surrogate for the capacity of the insurgent network, where recent civil wars correspond to well-armed and organized opposition, and many years of peace suggests that insurgents are less well organized.



**Figure 4: CARs™ software assists in reasoning over multiple models**

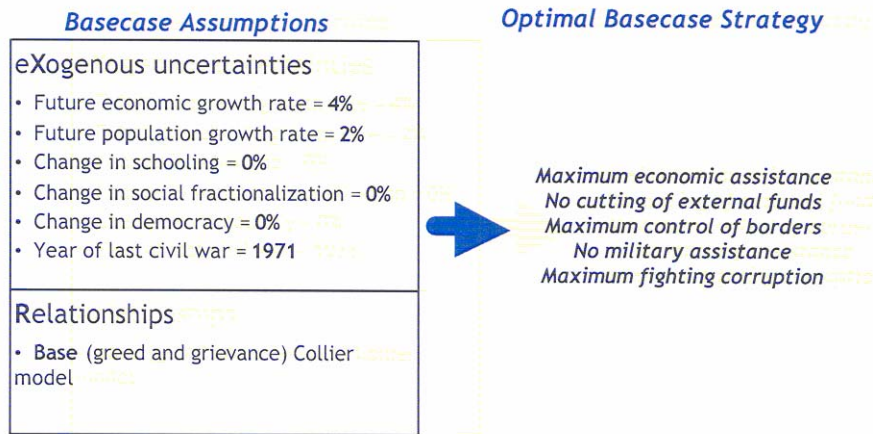
Once placed within CARs™, each model is treated as a function that maps choices about each of the uncertain Xs, and each of the controllable Ls into an outcome (M). This allows

the CARs™ interactive graphics and search capabilities to be used to assist us in exploring the behavior of these models.



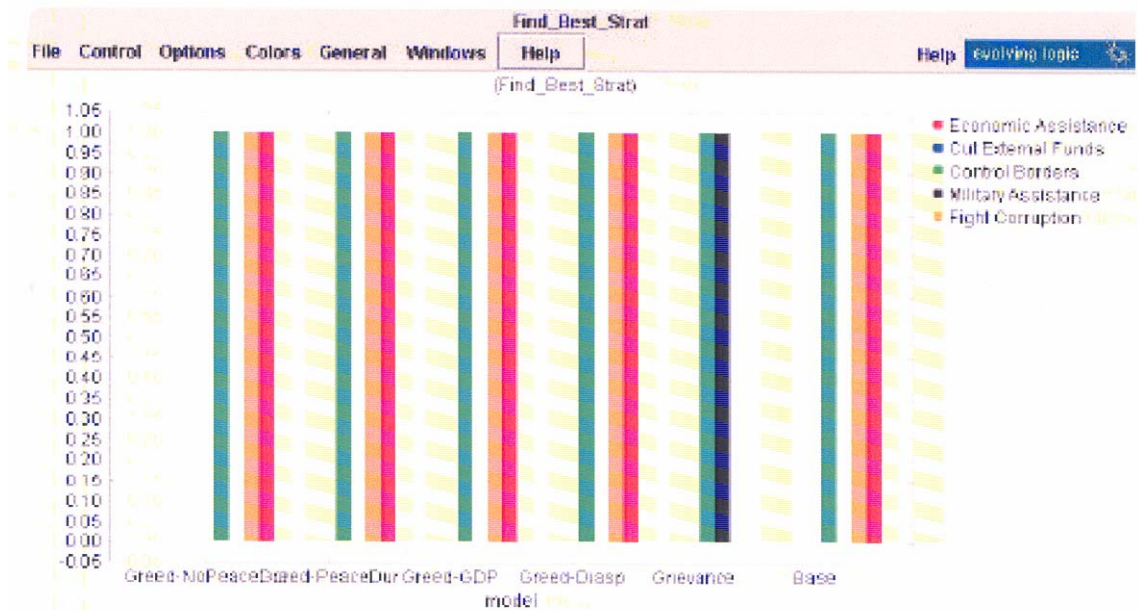
**Figure 5: RAP identifies robust strategies and key uncertainties with iterative process**

With the basis just described, we now can pursue the identification of robust strategies and key uncertainties. We first do initial interactive discovery to determine one or more candidate robust strategies. Once we have such candidates, the computer can then be tasked to search automatically for plausible future states (corresponding to assumptions about uncertainties) where it is vulnerable. Having identified such difficult scenarios, we can then renew the search for strategies, make an assessment about how much such possible stressing cases should concern us, or seek to create variants of the candidate robust strategy that shores up the detected vulnerability in the original. We can take these steps iteratively, and in doing so either discover strategies whose vulnerabilities are relatively modest, or discover classes of strategies that are robust to certain types of future, but vulnerable in others where a different strategy is robust. If no fully dominant strategy can be discovered, the choice between these classes of strategy takes the form of a wager between different classes of future scenarios.



**Figure 6: First Step: Identify optimal strategy for basecase future**

Following this general method, and using the best estimate model (which combines predictors of both the greed and grievance categories), we select the strategy that is optimal for the basecase when all the uncertainties take on their nominal, that is most expected or assumed values. As shown here, the optimal strategy for the case shown on the left is one that emphasizes economic assistance, controlling the borders and fighting corruption. It does not undertake military assistance or cutting the terrorists off from their funding sources.



**Figure 7: Optimal strategies for basecase Xs for different Collier models**

Here we display the results of repeating this calculation for each of the six models. The strategy space here consists of all possible combinations of yes/no choices for the five strategy components (levers), and the presence of a colored bar implies that lever was part of the strategic portfolio. Five of the models, including the “best estimate” base model, have the same optimal strategy for the nominal case. The sixth model, “Grievance”, however, has a different optimal strategy, combining military assistance with controlling the borders. Economic assistance and fighting corruption are clearly less valuable if this view of the world is correct.

Put another way, even taking no cognizance of the uncertainty surrounding the correct future values of inputs, such as Pakistan’s future economic growth rate, different strategies are recommended by different models. Unless we are certain the best estimate model is the correct model, there is a danger that the optimal strategy for that model may not be robust.

The next step in our methodology is to search for vulnerabilities of our candidate strategy. Even though the optimal strategy for the best estimate model is not optimal for the Grievance model, perhaps its performance there will be acceptable. On the other hand, perhaps by varying assumptions regarding uncertain factors the optimal strategy will fail even for the best estimate model. We conducted a series of computational experiments to illuminate these questions.

UNCERTAIN PROJECTIONS	Low	Nominal	High
Future economic growth rate	-5%	4%	10%
Future population growth rate	-10%	2%	4%
Change in schooling	-10%	0%	10%
Change in exports	-10%	0%	10%
Change in social fractionalization	-100	0	100
Change in democracy	-1	0	1
Year of last civil war	1971	1971	2020

**Figure 8: Wide ranges of Xs capture uncertainty**

All cases displayed to this point have used nominal values for all the factors that are inputs to the six models. We now will vary these inputs across the range of values we judged to be plausible for this analysis. Among the steps that can be taken is to vary these ranges to determine the maximal range of values for which a candidate strategy is robust. We do not show this step here.

By varying the inputs in a systematic way across all of the models, we generate an ensemble of possible scenarios. We regard this ensemble as a “challenge set” we will use to assess the robustness of candidate strategies. A strategy that performs satisfactorily (does not “fail”) on all members of the ensemble we will judge to be highly robust.

In order to conduct this survey, we must carefully define a criterion that will serve to assess the success or failure of a strategy. As we vary the input factors, the probability of conflict will also vary for each of the models, even holding the strategy constant. Some cases will tend to generate high probabilities of conflict for all the strategies, while others will be easy cases for all strategies. Thus, defining the criteria for success or failure in terms of the measures of interest (probability of conflict in this case) will confuse the difficulty of the case with the performance of the strategy.

Instead, we normalize the performance of any strategy against the performance of its competitors. The means for normalizing we choose here is to calculate the regret of a given

strategy in a given scenario (that is for a given choice of model and uncertain inputs). Regret is defined as the difference between the performance of a given strategy and that of the optimal strategy for that future. Thus, in every future, there will be at least one strategy with zero regret. A robust strategy is one that has small regret across all possible futures.

It should be noted that the concept of regret is quite old. In particular, Savage in 1950 argued that minimizing regret is the right way to make decisions when the probabilities of different cases are unknown.

Ensemble = {rain, no rain}; assertion = decision X is “good”  
 Good = regret(X) is never very big

State of the World	Strategy	Max. Regret
Sunny Day	Leave umbrella home	Wet and cold
Rainy Day	Carry umbrella	Burden and potential loss
- ? -	Keep umbrella in car	Rain: damp Sun: clutter

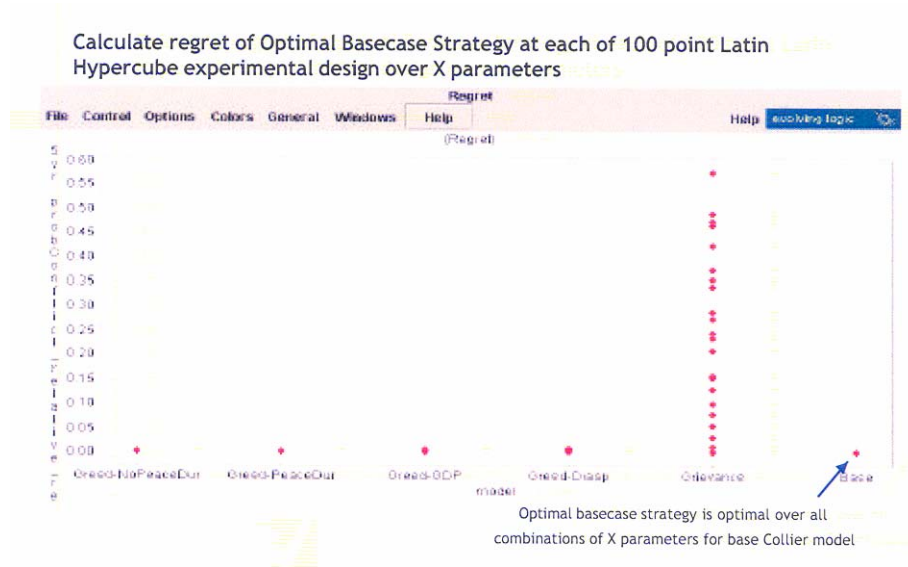
**Figure 9: Simple example of comparing regret under two scenarios and under uncertainty**

A simple example will serve to make the concept of regret and the role it can play in robust inference more obvious. Consider the decision of whether to carry an umbrella to work. The ensemble of possible futures contains just two cases, either it rains or it doesn't.

If it's going to be sunny, the optimal strategy is to leave the umbrella at home. But if we choose that strategy and guess wrong about the weather, our regret is that we are wet and cold when we wouldn't have been if only we had brought the umbrella. On the other hand, if it is going to rain the optimal strategy is to carry the umbrella. For this strategy, the regret if the weather surprises us is we have to carry an umbrella around all day long, and perhaps even lose it.

This framework gives us a basis of possibly discovering the hedging strategy of bringing the umbrella but keeping it in the car. This strategy is optimal for neither case: there is always some regret. If it rains we get slightly damp on the run to the car to fetch the umbrella. If it's sunny, we have the clutter of having the umbrella in the car. But, while optimal for neither

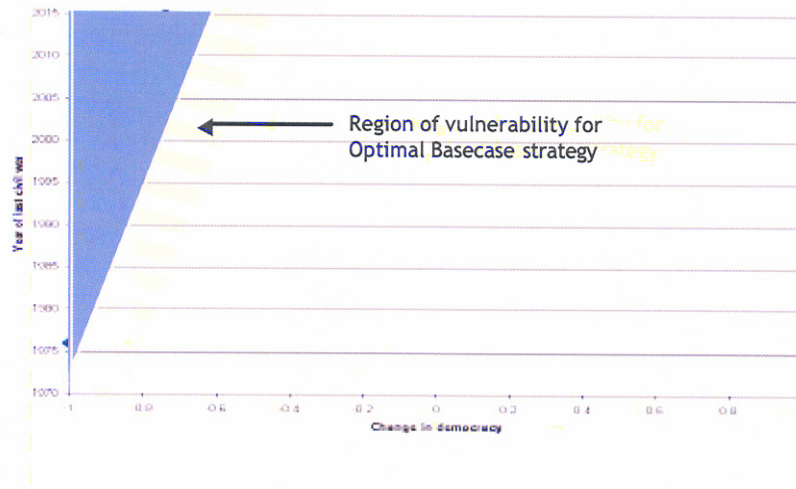
case, this strategy has a smaller maximal regret than either of its competitors. On this basis, it can be recommended as a clever hedge if the weather is uncertain.



**Figure 10: Optimal basecase strategy is vulnerable in grievance model**

Figure 10 displays the regret of the candidate strategy for many randomly chosen cases on all six models. The sampling technique used was a Latin Hypercube (space filling) design. Cases were thus chosen to span the full ranges of all the uncertain variables shown in Figure 8. As can be seen on from this graph, the candidate strategy has zero or small regret on all cases sampled from five of the models, and for some cases on the sixth. However, there are also some cases for the Grievance model where the regret is substantial, as much as a 60% increase in the probability of conflict.

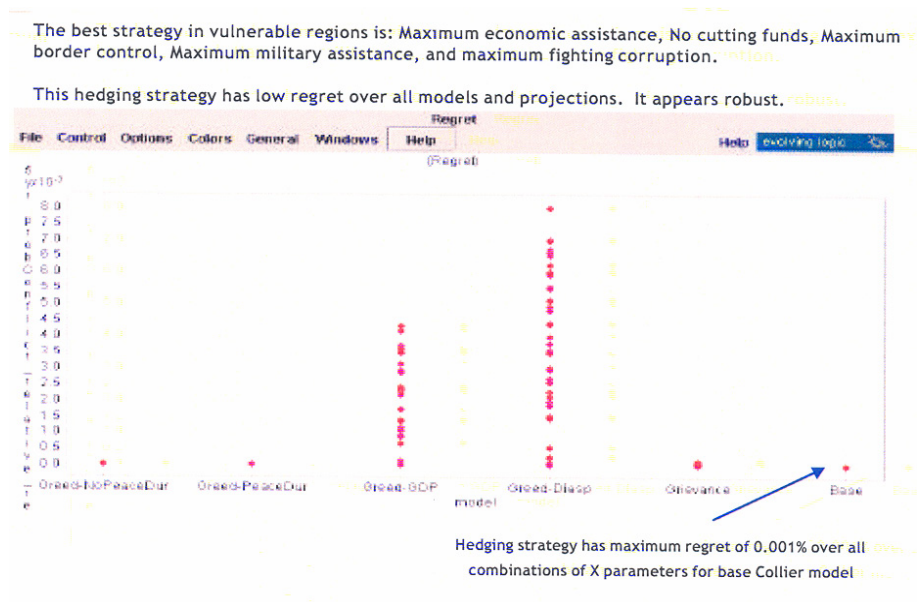
Statistical analysis suggests *Change in democracy* and *Year of last civil war* are the most important X factors in determining regret of Optimal Basecase strategy in Grievance model



**Figure 11: We can detect conditions causing optimal basecase strategy failure in grievance model**

Further analysis of the high regret cases can provide a great deal of insight. We selected only the high regret cases from Figure 10, did an analysis of variance, and display here these cases in terms of the two most influential variables in producing them. This view allows us to say that the candidate strategy performs well except if it turns out that the Grievance model is the right way to think about conflict in Pakistan, there is a significant decline in democracy in the coming years, and the terrorist infrastructure is relatively proficient. If this strategy represented our best option, perhaps the decision maker would be willing to wager against this failure mode happening. But, perhaps further analysis can provide a better candidate and avoid the necessity of placing that bet.

This brings us to the next step in our methodology, where we search for strategies that perform well against the vulnerabilities that have just been identified.



**Figure 12: Strategy exists that is robust over all models**

First a search over strategies reveals one that has the best overall performance on the high regret cases. Note that this strategy is not optimal for the nominal case for any of the models.

We then repeat the vulnerability analysis we did for the best estimate strategy, calculating the regret of this new strategy across a series of cases on each of the six models. Those results are displayed in Figure 12. This new strategy has non-zero regret for all of the models. But, its maximum regret is two orders of magnitude less than the maximum for the best estimate strategy. (Note that the scale on this graph is expanded compared to the previous graph of this form shown in Figure 10.)



**Figure 13: Placing both results on same scale**

Here we have placed the regret data for both strategies in the same graph. The worst case for the hedging strategy has a relative regret of less than 1%. While further searching might reveal an even more robust strategy, this one is clearly much more robust than the original candidate, which was optimal for the best estimate model.

Had we used only the best model as the basis for our analysis, we would have recommended a strategy that was more vulnerable to surprise than the hedging strategy. The discovery of the hedging strategy was made possible by our use of the suite of alternative models. These different views, all derived from the same data, provide us with more information than any single model would have provided us. If we are interested in robust strategies, the use of multiple alternative models provides a clear advantage.

## 5. RESULTS AND DISCUSSION

Having demonstrated the virtues of suites of alternative models, we turn next to multiple models of different types. We retain the six Collier models, and include two additional models of different type. CAST digests daily news reports, and combines them using an expert designed weighting scheme to produce 12 indicators of potential conflict. For this exercise, we reimplemented the CAST algorithm in java.

We also created a model based on structured interviews with military officers and government officials. This model represented the effects differing strategies might have on various aspects of the future situation in Pakistan. Using the XLRM methodology (from Figure 3), we captured not only their consensus belief about the effects that various policies might have, but also our uncertainty about the size of these effects.

	Uses Real Time Data & Expert Judgment	Uses Econometric Data & Statistical Estimation	Supports Exploring Decision Options
CAST	Yes	No	No
Collier Models	No	Yes	No
Elicitation Model	No	No	Yes
Fused Models	Yes	Yes	Yes

**Figure 14: Fusion of complementary models into a suite allows better utilization of the data**

In total, we then have a suite of 8 models of three different types. They are representative of classes of model that might be used in a future PCAS system. CAST responds to rapidly changing data to produce indicators of trends and is effectively an expert system. The Collier models are examples of the use of statistical regression to capture the pattern in historical data, and utilize information that is slower to change in order to extrapolate events farther into the future. The expert elicitation model is less well grounded in social science than the other two, but reflects the pragmatics of potential operations, based on the experience of the operators.

Each of these classes of model brings different information into the analysis. If we restricted ourselves to only one type, the other sources of information would not be available. If we can use the models jointly however, in principle the insights of all types can be utilized.

The inputs and outputs of the various models are not unrelated. In fact, the lever inputs to the Collier models and to the Effects model (Expert Elicitation model) are identical by design. We can further relate the 12 indicators that are outputs of CAST to various X inputs to the Collier models. And outputs of the Collier models, particularly the probability of conflict by year serve as an input to the Effects model.

Our technology allows us to avoid expressing this relationship between the models as “hard wiring” created by revising the software. Such hard wiring often results in difficulties when models of aspects of a deeply uncertain problem are combined in this way. The scheme for joining two models in this way is itself an act of modeling, and as with other models of deeply uncertain phenomena, the assumptions embodied in the wiring itself represents an uncertainty whose implications should be explored. Failure to contend with this difficulty has at time led to nonsensical results in past efforts at model fusion.

Instead of hard wiring, we continue to treat each model as a separate platform for computational experiments, but allow the strategy for case generation for experiments on various platforms to be coupled or correlated in arbitrary ways. This allows us to treat the constellation of models as a single unified model if we wish to, but also allows the confederation of models to be exercised in other ways, to answer other questions than “what would happen if...?”

As we caution that wiring or gluing models together may not be the most effective means of inference using them jointly, it may be that “model fusion” is poor terminology to describe what we believe can be effective joint use. We have continued to use this term in this briefing as there is no generally recognized alternative yet available.

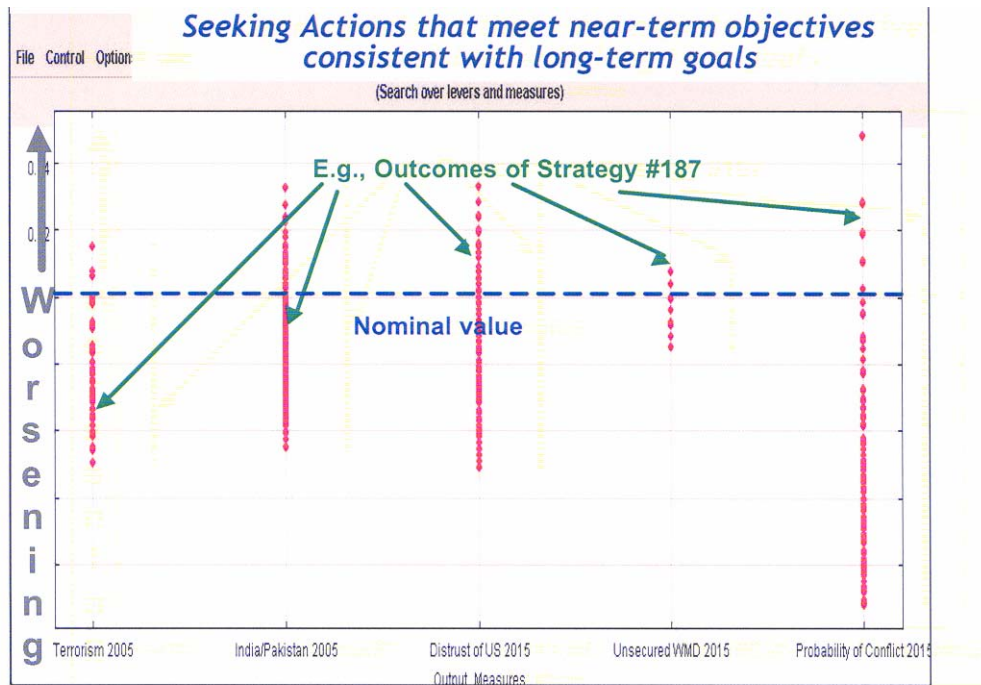
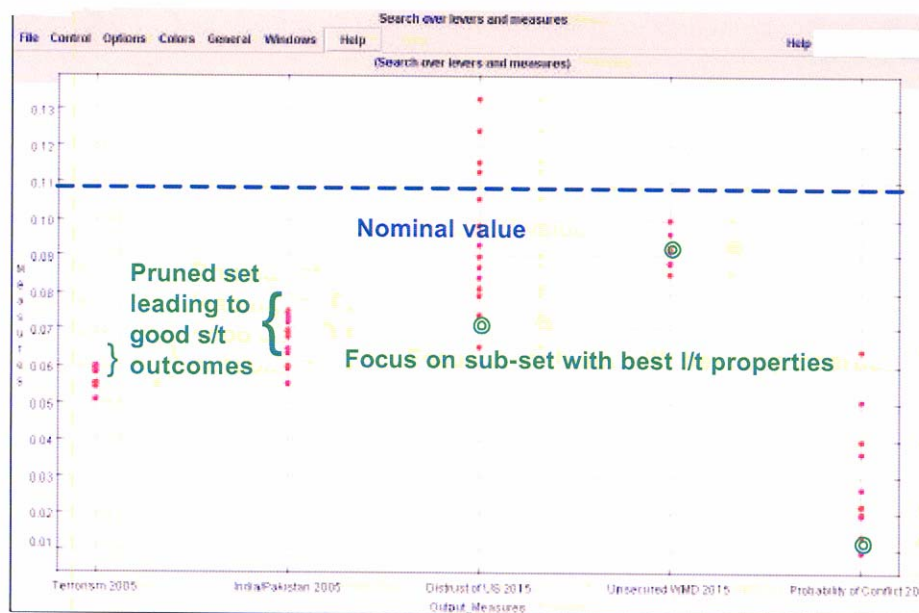


Figure 15: Seeking actions that meet near-term objectives consistent with long-term goals

The constellation of models allows us to ask questions regarding ranges of plausible outcome and robustness of candidate strategies, similar to our demonstration with the six alternative models above.

Our focus at present however is on the use of this suite of models to support devising robust shaping strategies. By this we mean strategies that are robust not only to the uncertainties (Xs) but also to the choice of measure, corresponding to the array of near term and long term goals we would like to achieve.

Depicted here is the outcome of a broad sampling of alternative strategies assessed on five different measures representing two points in time. For each measure, there are strategies that lead to better outcomes, and strategies that lead to worse outcomes. The trick in devising shaping strategies is to find an option that performs satisfactorily in both the short and long term, and also across a variety of different goals and constraints. In this case, we desire a strategy that confronts terrorism in 2005 without increasing tensions between India and Pakistan, but which also promotes desirable outcomes in 2015, in avoiding or mitigating public distrust of the US, unsecured WMD, and the probability of internal conflict.



**Figure 16: Understanding tradeoffs among best near-term strategies**

Various computerized search techniques might be employed to help discover good options for robust shaping strategies. Here we demonstrate the potential utility of this approach by the simple expedient of filtering the large set of strategies on the previous slide, to show only those with good properties in the short term. As among this remaining set of strategies are a few that also have desirable long term properties (marked in green), we can

see that by examining many options, ones with specially desirable properties may be discovered.

In an eventual decision support context, we would aspire to discover as many candidates of this kind as possible, so that the commander or other user will have the ability to select among them based on additional criteria not represented in the computer. Rather than the machine suggesting a course of action that users must either accept or reject, presenting a series of options allows the human to merge his other insights with those produced by the machine.

## 6. CONCLUDING REMARKS

While we have had some success in using constellations of models in concert, various questions remain regarding the issue of model fusion.

First, it is worth noting that model fusion involves issues of some depth and subtlety. The fusing of complex models such as simulations has been attempted for many years and has often proven much more difficult than originally anticipated. Especially where models have not been designed with fusion in mind, many problems can arise. The software engineering problem of plugging one model's output into another model's input has been solved with the success of HLA (High Level Architecture) as a clear example. But, the software engineering problem is much less difficult than is bridging the differences between different assumptions that have been used in differing models. These assumptions get reflected in approximations, techniques for addressing uncertainty, and forms for representing important phenomena.

Thus, it must be emphasized that using models together is itself an act of modeling, and must be driven by a strategy for solving a particular problem. It is a much deeper issue than finding a way to connect the models.

As all models are approximations, an approach to fusion that does not account for differing assumptions in the pieces can result in nonsense. (Recall that from A and not-A, you can deduce anything.)

We believe that our technology and our approach to the problem of model fusion provides a much better foundation for using suites of models than is possible through "wiring the models together". Our approach allows the "wiring" between the models to be set at "run time" and hence varied in response to changing hypotheses and goals. This allows us to use strategies for specifying groups of computational experiments to account for uncertainties inherent in the model fusion itself.

Our approach allows the problem of model fusion to be transformed into a version of the machine learning problem. Improvements in our ability to use large numbers of models in concert aggressively will come from improved algorithms for search or adaptively sampling from large spaces of possible computational experiments, -- and in deriving conclusions from the results of those experiments.

In recent years, the machine learning community has been moving to ensemble-based methods. Techniques such as bagging, boosting, and stacking have proven to provide value

across a wide range of applications and constituent model types. This is related to the approach to model fusion we are investigating. We can demonstrate that ensembles of models often contain more information than any single model does. A growing portfolio of means for exploiting ensembles of models is becoming available.

We are interested in exploiting techniques adapted from the Machine Learning literature to decision support problems, and believe that continuing innovation in this area will provide for discontinuous improvements in our ability to confront the really difficult challenges, such as those posed by PCAS.

We have addressed key issues in model fusion and reasoning strategies for exploiting suites of models and selected a vignette for demonstration purposes. We identified and assembled a suite of models and data and built software for effective model fusion. By doing this, we demonstrated we can discover shaping strategies robust across multiple objectives and models. The use of multiple models can improve decisions by incorporating insights from:

- Different fields (economics, military, political, sociological, cultural)
- Different data input types (statistical time series, area expertise, media feeds)
- Different perspectives on cause and effect (agency perspectives on policy)

Support tooling can assist in developing shaping strategies robust against multiple scenarios of future developments and shaping strategies from multiple models may be superior to those from single models.

We have demonstrated the feasibility of our approach. Further work is needed to take these techniques and make them ready to operational use. We have confidence that moving these techniques to operational use is feasible, and that doing so could bring great benefits.

## 7. LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

PCAS	Pre-Conflict Anticipation and Shaping
CARs™	Computer Assisted Reasoning® system
GNP	Gross National Product
RAP™	Robust Adaptive Planning™
XLRM	eXogeneous, policy Levels, Relationships, Measures
WMD	Weapons of Mass Destruction
HLA	High Level Architecture