APOLLO: An Analytical Tool for Predicting a Subject's Decision Making

Paul Sticha, Ph.D.
HumRRO
66 Canal Center Plaza
Alexandria, VA 22314
psticha@humrro.org

Dennis Buede, Ph.D.
Innovative Decisions. Inc.
P.O. Box 231660
Chantilly, VA
dbuede@innovativedecisions.com

Richard L. Rees, Ph.D.
260 Springvale Road
Great Falls, VA
RLRees@aol.com

Keywords: Multiple Competing Hypotheses, Information Sharing and Collaboration, HUMINT, OSINT, terrorism

Abstract

APOLLO is a software application—designed in conjunction with analysts—that enables the analyst to reason through a prediction of a Subject's decision making, to identify assumptions and determinant variables, and to quantify each variable's relative contribution to the prediction, producing a graphical representation of the analysis with explicit levels of uncertainty. The analyst builds Bayesian networks that integrate situational information with the Subject's personality and culture to provide a probabilistic prediction of the hypothesized actions a Subject might choose. The Bayesian network is integrated with a tool that sorts through incoming all-source reports pulled by the analysts' favorite search engines, ranks their relative salience to specific variables, and captures sourcing (the evidence) for later review by the user or others. The analyst then updates—manually or with partial automation—the predictive model, after which the APOLLO software recomputes the probabilities and alerts others when the outcome probability crosses a user-selected threshold based on new evidence. This integrated modeling, data, and knowledge management tool—linked to the analysts' existing databases and production platform—systematizes the analyses of Subject decision-making and applies accepted principles of social science to the practice of intelligence analysis.

1. Introduction

This effort, underway since 2001, addresses many of the concerns in recent critiques of the national intelligence process. The key purpose of our work is to create a method for modeling/predicting what key figures might do in specific situations. We envision this effort as being undertaken by a team of analysts with diverse backgrounds, working on a problem that will evolve over weeks or months and need periodic reporting. Our search is for a method that will neutralize various analytic biases such as recency, halo, proximity, hindsight and personalization. In addition, the method should neutralize less widely discussed but influential social biases, such as those reflected in giving undue weight to the senior expert, the "party line" or published record, the analyst with the biggest fistful of cables, or with the most dazzling personality.

This modeling method provides a natural mechanism for surfacing assumptions, logic, and new evidence for the team working the problem. In addition, APOLLO captures an auditable history of the team's thought process and supporting evidence. This software solution alerts the analyst when certain thresholds are met within the model, indicating that the evidence may warrant changing one's beliefs.

The following principles were the basis of this development effort:

- Viewpoints—the more the merrier, but make it systematic (combine many analysts with varying expertise, addressing both the situation and the subject's personality and culture, and use a rigorous analytical method for integration)
- Intel as process, not just as product—recognizing that intelligence consumers are their own analysts, engage the customer by using the model to manage the debate and questioning that often ensues when an analyst briefs a policy-maker—the software enables different assumptions or alternative hypotheses to be tested on the spot
- Continuous and real-time updating of the model—review and quantify relevant evidence and the associated probabilities for specific model variables, and explicitly inform the user when data may warrant changing judgments.

2. Modeling Method

This section introduces and justifies our use of Bayesian networks as a modeling method, and illustrates the modeling process with a sample model. The process by which the networks are developed also contributes to the effectiveness of the method. We briefly describe our experience with this process and highlight some ways that it supports the goals of the project. A number of participants have offered unsolicited kudos for the group process as much as for the tool or the output.

2.1 Processes, People, and Norms

We have developed models in two-day, facilitated meetings attended by analysts, model developers, and external subject-matter experts. The facilitator guides the participants through the steps in the development process (described in the following sections), elicits estimates of model parameters, and ensures that the requirements of the methodology are met. A second member of the modeling staff implements the model on the computer and takes notes. The model is projected onto a screen during the development process so that all participants are aware of the variables and relationships included in it. Both analysts and external experts provide the information and assessments that are incorporated into the model. The analysts usually provide critical information about the questions to be addressed by the model, while all participants provide the regional and Subject knowledge incorporated into the model.

The analysts' intelligence questions, the optional outcomes of interest—what will X do or what can we do to lead X to do Y—are debated at considerable length. The process in 2.3 below usually takes four hours at the outset of a two-day session and sets the motif—and some social norms—for the rest of the session. The mix of staff and outside panelists suggests that diversity of opinion and experience is desirable; it gives the staff permission and cover for bringing new or divergent views to the table.

The facilitator's behavior is critical in two respects. First, the technical aspects of applying Bayesian analysis must be guided by an expert—the international relations and political science educations of most analysts don't prepare them for this methodology. The challenges of helping the group to frame questions properly—consistent with probability theory—and to keep their engagement fresh while estimating large conditional probability tables are not trivial items. In addition, the facilitator helps to keep the gate open to contrary data and judgments and healthy debate, to elicit contributions from all members, to challenge what everyone takes for granted, and to curb the natural tendencies of dominant actors to hog the stage and dictate the analysis—all while demonstrating respect for each contribution.

As the session proceeds and the facilitator leads the team through identifying the key determinative predictors and indicators of the situational variables, much debate about key variables ensues. Projecting the model on a screen as it is being developed provides a way to focus the discussion on specific issues, data, and opinions, while avoiding unproductive *ad hominem* debates. In addition, it may provide an environment that can encourage greater participation from reticent analysts. Consensus may not be feasible, but the model makes it possible to locate specific areas of agreement and disagreement and to determine the implications of this disagreement on the outcome of the model. Where there is disagreement about critical model variables, the areas of disagreement can be used to specify the requirement for additional intelligence collection.

Throughout the session, a notetaker records choices, issues and rationales for decisions to be included with the model as a history, which can help future users understand the logic underlying the model. At the end, we invite the group to review the model after a day or more—a process that can iron out wrinkles and spot deficits. Of course, many other social dynamics are managed in this process, but these are the highlights.

2.2 Why Bayesian Networks?

The general problem of predicting someone's future action is exceedingly complex. Without even considering the task of identifying the determinative variables correctly, one must deal with uncertainty, human judgment about the problem logic, relative strength of specific variables and evidence, and the dependencies of some variables on others. When we add the requirement to enable updates to the prediction as new information becomes available, we realize there is only one method that matches the problem statement—Bayesian probability (Schum 1994).

Recent advances in computer science and operations research have created Bayesian networks (Pearl 1988; Shachter 1988). Bayesian networks provide a graphical representation of the problem, using an acyclic directed graph to show the variables as nodes, the probabilistic dependence as arcs, and probabilistic independence as the lack of arcs. Conditional probabilities are captured inside the nodes. Sophisticated message passing algorithms are used to update the probabilities at all nodes based on evidence at several nodes (Buede 2001).

There is some similarity between this approach and Alternative Competing Hypotheses (ACH; Heuer 1999). In both our approach and ACH, a list of possible hypotheses is developed. Next, a set of possible indicators is brainstormed; these indicators, if true, would favor one hypothesis over the other. At this point the methods diverge. The ACH elicits qualitative statements (e.g., 1 to 3 pluses or 1 to 3 minuses) to capture the strength of the relationship between the hypothesis and each indicator. The Bayesian approach quantifies the conditional probability of the indicator given each of the hypotheses. With ACH, the result is a summary of the pluses and minuses associated with each hypothesis for identified indicators. The result of a Bayesian network is the posterior probability of the hypotheses given the identified indicators. The Bayesian approach can also incorporate causal factors that condition the probability of the hypotheses, address interactions between the indicators, and report uncertainty associated with the indicators and causal factors.

Another approach that has seen lots of applications in the last decade or so is the Situational Influence Assessment Module (SIAM; Rosen 1996). SIAM is used to compare alternate scenarios of causal factors (rather than indicators) for creating some desired or undesired outcome (hypothesis). A SIAM model looks like a Bayesian network in which there are many arrows entering a few nodes. In true Bayesian networks, this approach produces the need for an unmanageable number of conditional probability distributions. Using approximations involving labeling causal variables as promoters and inhibiters, SIAM reduces the number of conditional probability distributions to a manageable level. These independence assumptions must be considered carefully on a problem basis when judging the applicability of SIAM.

2.3 Defining the Question

The Subject's decision, for instance, to launch an attack is not simply yes-no, but whether to make a contingent attack, one involving certain levels of force, on certain days, against certain targets, or seeking certain outcomes, or may weigh attacks versus warnings or other public acts. These alternative competing decisions represent the analyst's best estimate of the choices considered by the Subject. We have on occasion used between one and four variables, each having two to six states to define the possible prediction of a key figure's decision. These states (in each variable) need to be mutually exclusive and collectively exhaustive. Clearly, when trying to predict the future by using a discrete number (4 to 4^4) of states, we must interpret mutually exclusive and collectively exhaustive loosely. For the sample problem being used in this paper, we address a situation in which a leader must decide what to do when beset by a national strike organized by his opponents. The states defined by the intelligence analysts were:

- Leave the country;
- Make concessions to end the strike;
- Hold a voter referendum in agreement to end the strike;
- Allow a regional organization to arbitrate the strike;
- Wait out the strikers; or
- Repress the strikers using violence when necessary.

We also define the leader's strategic objectives and develop a probabilistic relationship between the objectives and the hypothesized actions (see Figure 1). The two boxes represent variables (or nodes) that may take one of several values (or states). The arc from Leader's Objective to Leader Decn in Sec Event establishes that there is probabilistic dependence between these two variables.

2.4 Modeling the Situation

When a model is being developed, quite a bit of effort is made to identify possible situational variables that might change the outcome of the leader's decision. Once the possible situational events have been discussed and prioritized, key events are picked and added one at a time. After each variable is added, we conduct several "what if" analyses (changing situational outcomes) to see if the "model" makes sense in these different situations. When these analyses reveal errors or inconsistencies in the predicted probabilities, appropriate changes are made to the model. Additional variables are added subject to the time constraints for the model development process. Figure 2 shows the national strike model with the added situational variables (one-and-a-half days of the two-day session are typically completed by this point.)

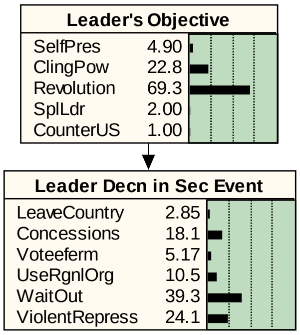

Figure 1. *Hypotheses and Objectives*

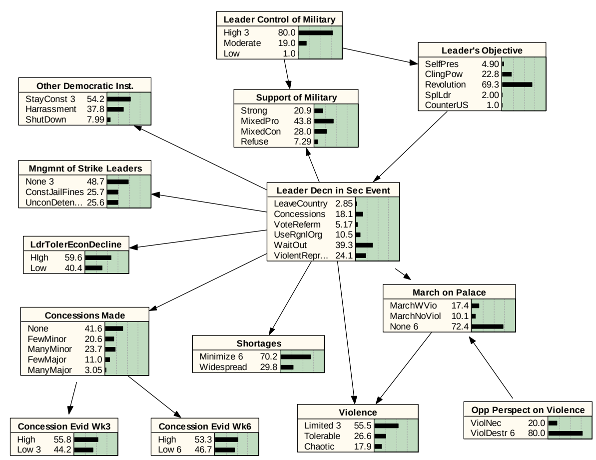

Figure 2. *Model with Situational Variables*

To develop the Bayesian network, it is necessary to represent the dependencies between variables by conditional probabilities. Table 1 shows the conditional probability table for one variable—Mgmt of Strike Leaders. There are similar tables for every node/variable in the model.

Table 1. *Conditional Probability Table – Probability of Mgt of Strike Leaders given Leader Decn*

	Mgt of Strike Leaders		
Leader Decn	None	ConstJailFines	UnconDetent
LeaveCountry	0.90	0.01	0.09
Concessions	0.60	0.30	0.10
VoteRefern	0.50	0.30	0.20
UseRgnlOrg	0.75	0.20	0.05
WaitOut	0.60	0.30	0.10
ViolentRepress	0.05	0.20	0.75

Figure 3 shows the updated probabilities after 3 weeks have elapsed, a number of intel reports have been received, and the values of some of the situational variables are known with near certainty, as shown by the shaded nodes in the figure.

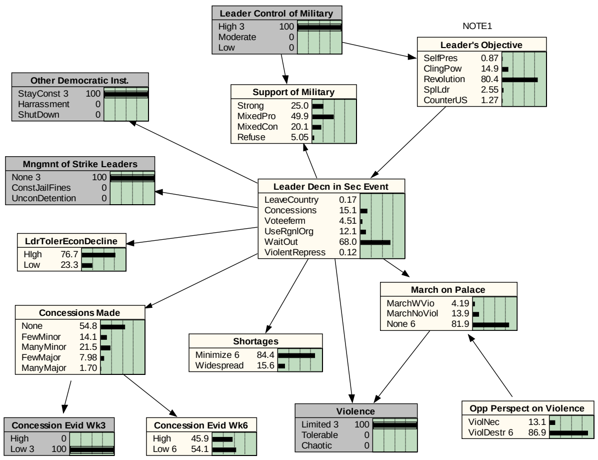

Figure 3. *Updated Model after 3 Weeks*

2.5 Adding Personality

There are many approaches to modeling personality; two of the more relevant and academically tested come from political psychology and personality psychology. After a detailed review of the personality literature and a consensus session with some of the leading researchers, we identified the following variables from the political psychology literature: positive image of others, internal locus of control, need for power, conceptual complexity, general distrust and suspicion, and acceptance of risk. (Sticha *et al.* 2000) From the personality researchers within psychology, the emphasis is on the five-factor model (neuroticism, extraversion, openness, agreeableness, and conscientiousness) (Costa and McCrae, 1985). Our early attempts at model building demonstrated that the leadership analysts were more familiar and comfortable with the concepts from political science than the five-factor model. However, the five-factor model has substantial research backing it up. In an effort to synthesize user acceptance and empirical foundations, we decided to integrate the two sets of personality factors. Following standard practice in psychological measurement, we adjusted the variables for intercorrelation, for reliability, and for correlation with behavioral outcomes (since personality does not perfectly predict behavior.).

Psychologists at HumRRO related the 30 facets from the five-factor model of personality to the six personality characteristics from political science/political psychology. The facets are the second tier elements of the five-factor model; each of the five factors has six facets. Table 2 shows the relationships established between the two models.

The second major element of the personality model is the incorporation of data and associated error. There are several ways to report and assess data. The NEO (a commercially available personality test with a form for knowledgeable informants) is a well-known, validated measure of the facets. Profiler+ (Young, 2001) is a content-analysis approach that analyzes first-person verbalizations according to Hermann's (1984) personality theory of leadership. Finally, HumRRO psychometricians developed a short, third-party evaluation form based on our variables. The estimated error of each kind of assessment is considered in the model.

Table 2. *Linkage Between Two Personality Models*

Political Psychology	**Facets from 5-Factor Model**
Positive Image of Others	Positive Emotion (Extraversion) Trust (Agreeableness)
Internal Locus of Control	Vulnerability (Neuroticism) Depression (Neuroticism) Assertiveness (Extraversion) Competence (Consc.) Self-Discipline (Consc.)
Need for Power	Compliance (Agreeableness) Achievement Striving (Consc.) Assertiveness (Extraversion)
Conceptual Complexity	Openness to Ideas (Openness) Openness to Values (Openness) Openness to Actions (Openness)
General Distrust & Suspicion	Trust (Agreeableness) Angry Hostility (Neuroticism) Warmth (Extraversion) Compliance (Agreeableness)
Acceptance of Risk	Openness to Actions (Openness) Anxiety (Neuroticism) Deliberation (Consc.) Excitement Seeking (Extraversion) Vulnerability (Neuroticism)

The third major element of the personality model specifies how it should be connected to the situational model including the hypothesis (decision) node and leader objectives. Initially, we connected the hypothesis node directly to the six political psychological characteristics. While this approach "worked" by producing interesting and believable results, it was too time consuming and required analysts to make very difficult judgments involving both the problem and psychology. After several intermediate attempts, we created a set of intervening variables from the political psychology literature that help express the relationship of traits to actions. The following six intervening variables link actions to personality, in that they can be considered to be both action characteristics and behavioral proclivities.

- Conflict versus cooperation (regarding opponents);
- Follow through required versus not required;
- Consistent with position versus not consistent;

- Unilateral versus collaborative (regarding colleagues);
- Substantive versus protocol; and
- Challenges constraints versus no challenges.

HumRRO, along with Intelligence Community psychologists estimated the quantitative relationship between the personality traits and the six behavioral proclivities; these correlations are embedded in the model. When we model a particular decision, we draw dependencies between hypothesis (decision) node and the action characteristics, thus specifying which proclivities are relevant to the decision. The process recognizes that not every proclivity relates to a particular hypothesis node (See Figure 4). To complete the integration of the situational and personality models, the analysts complete probability tables that specify the variability of decision options with respect to the action characteristics. We have found this procedure to be much easier than the initial approach was.

Figure 4 summarizes the entire modeling process. The situational model (on the left) is connected to the personality model (on the right), as described above. The analysts then enter as much personality information as is available on the subject in question. The results of this effort modify the probabilities (based on situational factors) that the subject will make each of the decisions specified in the states of the hypothesis node.

2.6 Performing "What If" Analyses

"What if" analyses are defined as tests of a model made by setting model parameters to see (a) how the changes in antecedent variable affects the outcome, and (b) if the results at that setting make sense. In addition, comparisons are made across multiple "what if" tests to see how changes in antecedents affect the relative results of the model from different settings and how they may make relative sense.

For example, we would move the probability setting on one of the situation variables from one extreme to another to observe the impact on the probabilities of the hypothesis. We would also compare the results of this analysis to those for other variables, first singly to observe the relative magnitudes of effect one variable makes with respect to the others. Then we compare the effects introduced by interactions among the situation variables. When we find results that do not make sense, we check for changes in the probability tables that would produce the results the analysts feel make sense. Sometimes, the analysts desire to make the changes to the probability tables; other times, they prefer to leave the probability tables as they were since they make more sense than the desired results. The software is very flexible in handling these calibrations and instantaneous in revealing results of multiple "what if" propositions.

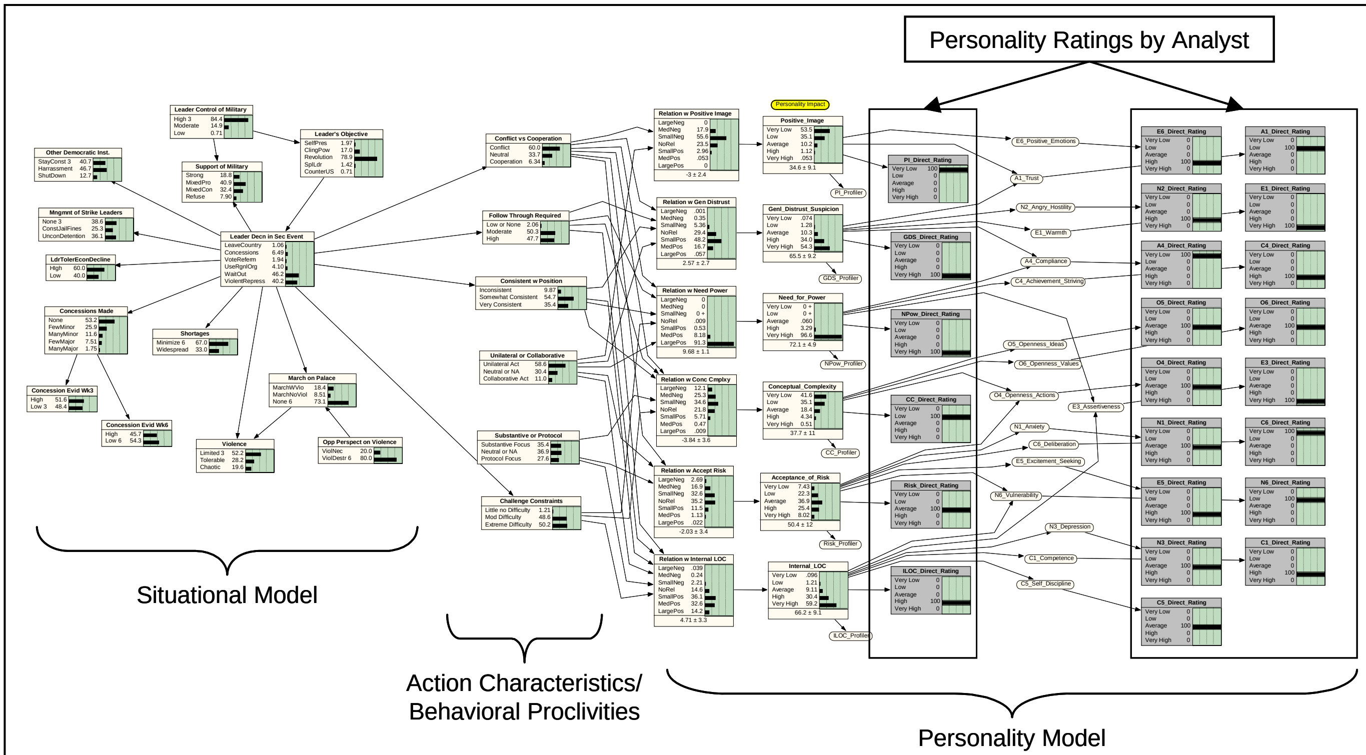

Figure 4. *Integration of Situational and Personality Models*

2.7 Assessing the Sensitivity of Variables

In most Bayesian network software implementations, the user can designate a node and calculate the mutual information between the selected node and other nodes, one at a time. This calculation identifies (based on the mutual information metric) the relative impact that changes in the probabilities of other notes will have on the probabilities of the designated node.

2.8 Automating the Integration of Reports—Future Capability

The Bayesian network will be tied into AIPSA (Automated Intel Processing for Situational Awareness) being developed by the Pacific Northwest National Laboratory (PNNL), which is based on the INSPIRE engine. AIPSA will sort through incoming all-source intelligence reports and Web documents using the analysts' favorite search engine, rank their relative salience to each situational variable, and captures the source information for later user and management review. The analyst then updates—manually or with partial automation (available in 2006)—the predictive model. As a result of the evidence changing the model, APOLLO recomputes the probabilities and alerts the analyst or others when the outcome probability crosses a user-selected threshold based on new evidence.

To support the production process and to add efficiency, the evidence (source documents) supporting each version of the Bayesian network, the entire model, the probabilities associated with the variables, and any annotations the analyst makes are stamped with date and user and stored for future editing or reference.

3. Summary of Models Built to Date

In the last two years we have built the following models while working with teams of intelligence analysts and expert consultants:

- Invasion
- National strike
- Domestic threat *
- Missile testing
- Support for the Global War on Terrorism
- Dispute over contested territory
- Peace/cease-fire negotiation *
- Use of WMD *
- Monetary devaluation *
- Establishment of a new caliphate*
- Operational planning in a terror cell*

Those models marked with an asterisk (*) are forward-looking models for which the answer was not known when the model was built. We plan to compare the predictions of these models to actual events to estimate the validity of the models. The models without asterisks were post-dictions of historical events. Although these models cannot be validated in the same sense as the forward-looking models, we will investigate the extent to which a model of a historical event can be applied to a similar situation with a different Subject.

4. Summary and Conclusions

The APOLLO program, underway for three years now, is currently delivering a software-based tool to intelligence analysts that supports the development of Bayesian network models to address a wide range of situations in which a leader is making a decision, the effects of which will evolve over several weeks/months. A library of models has been under development during this time period as a proof of concept and as a resource for analysts to use as part of bootstrapping their efforts. The models span many different topic areas (invasions, national strikes, missile testing, WMD, and economics).

References

Buede, D.M. 2001. *Application of Bayesian Networks to Intelligence Analysis*. Innovative Decisions, Inc. Technical Report.

Costa, P.T. & McCrae, R.R. 1985. *The NEO Personality Inventory manual*. Odessa, FL: Psychological Assessment Resources.

Hermann, M.G. 1984. Personality and Foreign Policy Decision-making: A Study of 53 Heads of Government. In D. Sylvan & S. Chan (eds.), *Foreign Policy Decision-making: Perception, Cognition and Artificial Intelligence*. (pp. 53-80). New York: Praeger

Heuer, R. J. 1999. *Psychology of Intelligence Analysis*. Washington, D.C.: Center for the Study of Intelligence.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kauffman.

Rosen, J.A. and Smith, W.L. (1996). "Influence Net Modeling with Causal Strengths: An Evolutionary Approach", in 1996 Command and Control Research and Technology Symposium.

Schum, D.A. 1994. *The Evidential Foundations of Probabilistic Reasoning*, New York: Wiley.

Shachter, R. 1988. Probabilistic Inference and Influence Diagrams. *Operations Research*, 36(July-August), 589-605.

Sticha, P.J., Handy, K., and Kowal, D.M. 2000. *Predicting Policy Decisions: Literature Review*. HumRRO Technical Report 2000*N09200*000.

Young, M.D. 2001. Building worldview(s) with Profiler+. In M.D. West (ed.), *Applications of computer content analysis*. Westport, CT: Ablex.