Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

USING CAUSAL MODELS IN HETEROGENEOUS INFORMATION FUSION TO DETECT TERRORISTS

Paul K. Davis
David Manheim

Engineering and Applied Sciences Department
RAND and Pardee RAND Graduate School
1776 Main Street
Santa Monica, CA 90407-2138, USA

Walter L. Perry
John Hollywood

Engineering and Applied Sciences Department
RAND
1200 S. Hayes St.
Arlington, VA 22202-5050, USA

ABSTRACT

We describe basic research that uses a causal, uncertainty-sensitive computational model rooted in qualitative social science to fuse disparate pieces of threat information. It is a cognitive model going beyond rational-actor methods. Having such a model has proven useful when information is uncertain, fragmentary, indirect, soft, conflicting, and even deceptive. Inferences from fusion must then account for uncertainties about the model, the credibility of information, and the fusion methods—i.e. we must consider both structural and parametric uncertainties, including uncertainties about the uncertainties. We use a novel combination of (1) probabilistic and parametric methods, (2) alternative models and model structures, and (3) alternative fusion methods that include nonlinear algebraic combination, variants of Bayesian inference, and a new entropy-maximizing approach. Initial results are encouraging and suggest that such an analytically flexible and model-based approach to fusion can simultaneously enrich thinking, enhance threat detection, and reduce harmful false alarms.

1 INTRODUCTION

1.1 Purpose

This paper illustrates how we have used a computational version of an originally qualitative social-science model for basic research on heterogeneous information fusion bearing on detection of potential terrorists. The term “heterogeneous” highlights the diverse character of the information being fused—e.g, behavioral observations in an airport, prior-arrest records, and reports from agents of varied quality and reliability. The information is often qualitative, soft, conflicting, and even deceptive. The model assists in using such diverse and fragmentary information to piece together an estimate of the threat of terrorism posed by the individual. With respect to modeling theory, the paper illustrates the potential value of *causal* social-science models, assuming that they are used with proper respect for both structural and parametric uncertainties. The context is assisting uncertain inference about threat, rather than making point predictions or issuing firm judgments. Such fusion necessarily includes considerable subjectivity and analytic artistry, but it can be given structure and rigor, and it can include extensive and useful uncertainty analysis. Such improved fusion methods could increase the probability of detecting the rare potential terrorist, decrease false alarms, and increase the probability of exonerating individuals who might otherwise be falsely assessed. Future work will determine how much can be achieved.

978-1-4673-9743-8/15/$31.00 ©2015 IEEE

1.2 Background on Social-Science Causal Models

Earlier work reviewed scholarly social science bearing on terrorism (Davis and Cragin 2009). That began a movement toward synthesis and causal analysis by introducing easy-to-understand "factor trees" identifying the factors contributing to terrorism and public support thereof, and how those factors relate to each other. Social scientists are excellent in identifying such factors even though predicting consequences is more difficult. Even initial factor trees can elicit further expression of knowledge. Expert viewers can quickly spot omissions and ambiguities. The iterated factor trees can then be useful "thinking models"—i.e., conceptual models to structure reasoning and discussion. An earlier paper (Davis 2011) provided a primer on qualitative factor trees. More recently, case studies were conducted to "validate" a factor tree for public support of terrorism. The factor tree (Figure 1) held up well (Davis, Larson, et al., 2012) as a general qualitative theory with myriad context-specific specializations. The authors discussed what "validation" can mean here and focused on (1) tentative confirmation of factors (e.g., do the tree's factors show up in the new cases as judged from polls, news accounts, diaries, and the writings of insurgent leaders), (2) tentative confirmation regarding causality and necessity, and (3) model enhancement (if new cases reveal some additional factors or somewhat different relationships among factors, this may best be seen as "refining" rather than "falsifying" theory). Humility is important because such models cannot be validated as in the physics laboratory.

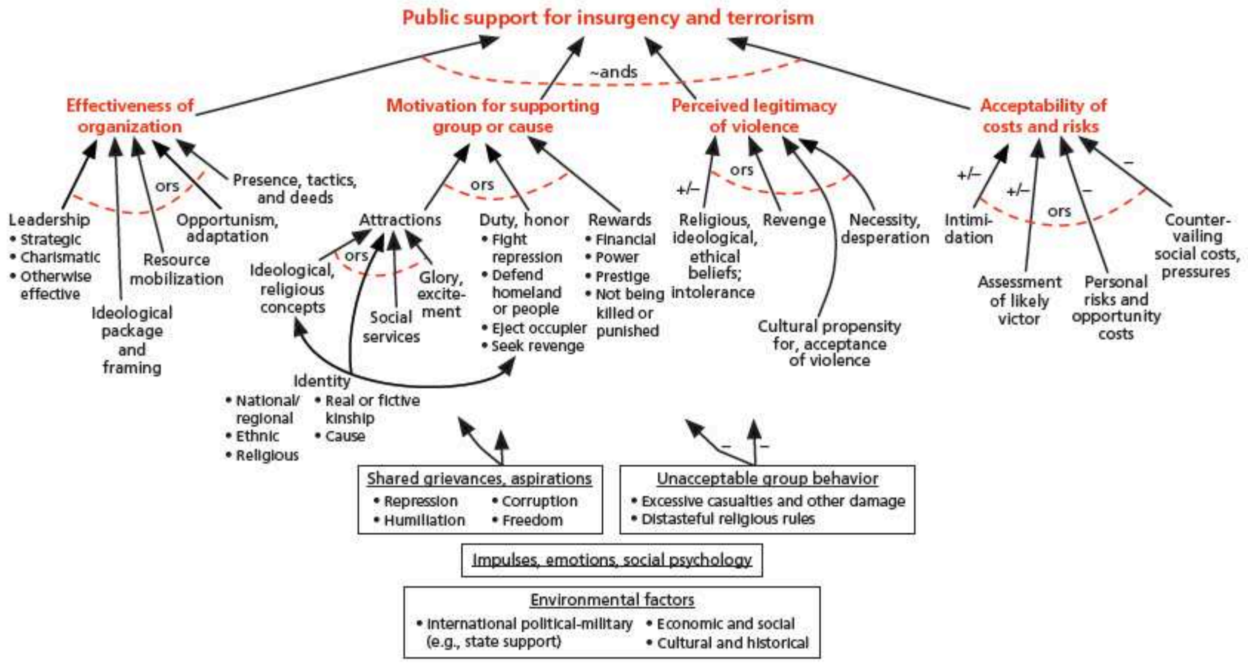
Public support for insurgency and terrorism
~ands
Effectiveness of organization
Motivation for supporting group or cause
Perceived legitimacy of violence
Acceptability of costs and risks
ors
Presence, tactics, and deeds
Leadership
• Strategic
• Charismatic
• Otherwise effective
Opportunism, adaptation
Resource mobilization
Ideological package and framing
ors
Attractions
Duty, honor
• Fight repression
• Defend homeland or people
• Eject occupier
• Seek revenge
Rewards
• Financial
• Power
• Prestige
• Not being killed or punished
ors
Ideological, religious concepts
Glory, excite-ment
Social services
Identity
• National/ regional
• Ethnic
• Religious
• Real or fictive kinship
• Cause
ors
+/–
Religious, ideological, ethical beliefs; intolerance
Revenge
Necessity, desperation
Cultural propensity for, acceptance of violence
+/–
+/–
ors
–
–
Intimi-dation
Assessment of likely victor
Personal risks and opportunity costs
Counter-vailing social costs, pressures
–
–
Shared grievances, aspirations
• Repression
• Humiliation
• Corruption
• Freedom
Unacceptable group behavior
• Excessive casualties and other damage
• Distasteful religious rules
Impulses, emotions, social psychology
Environmental factors
• International political-military (e.g., state support)
• Economic and social
• Cultural and historical

NOTES: Applies at a snapshot in time. Current factor values can affect future values of some or all other factors.

RAND *TR1220-S.1*

Figure 1: A factor tree for public support of insurgency and terrorism.

Subsequently, in a step taken with trepidation because of conceptual challenges and the uncertainties involved, a computational model was developed from the factor tree of Figure 1 (Davis and O'Mahony 2013). In doing so, the authors confronted the challenges of theory and method summarized in Table 1.

Table 1: Challenges and issues in moving from a qualitative factor tree to a computational model.

Challenge Issues	*Issues*
Define factors and factor values	How many values are sufficient (the binary case is too crude)? How can soft and fuzzy variables be defined?
Define the tree's "and" and "or" connections mathematically	How rigid should the relationship be? How can uncertainties be represented? How many alternative functional relationships are needed?
Define ambiguous and conflicting influences (+/– signs) mathematically	What does the ambiguity mean? How can it be represented?
Represent implications of line thickness in factor trees (not shown in Figure 1)	How should relative importance of factors be understood and represented in the model?
Represent uncertainty of factor values	Should this be done by giving ranges of parameter values or by using probabilistic methods?
Represent structural uncertainty of combining relationships	How can this be done?
Build model for exploratory analysis under uncertainty and assessment of confidence in estimates	How should exploratory analysis be accomplished? When should probabilistic methods be used?
Implement model in understandable high-level language	What language? How can the model be made transparent, comprehensible, and easy to re-implement for re-use?

The Davis-O'Mahony report documents the solutions found and the rationale that led to them (see pp. 73-84 for the mathematics). We describe them briefly as they apply to the work reported here. First, we define the factors (variables) of the model on an interval scale of 0 to 10, often using the discretized scale of 1, 3, 5, 7, 9 with equally spaced values corresponding to very low, low, medium, high, very high. With this type of scale (as distinct from an ordinal scale) it is legitimate to perform basic mathematical operations, albeit with caution. See, e.g., Carifio and Perla (2007) for a window into continuing debate. At the data-interpretation level it is necessary to have protocols defining how observables should map consistently into scale values. These will often use concrete examples for calibration so that a new observation can be compared subjectively to those concrete examples before classing them as, say, "very high" rather than "high" or "medium." It is an empirical matter to determine whether the results are then consistent. Such work is fuzzy but meaningful.

The actual functions describing the combined effects at each node of a factor tree may be subtle and complex. However, we concluded that much can be accomplished with a combination of a few building-block functional forms. Two, in particular, have been workhorses. They represent two more or less bounding ways to represent nonlinear effects of factors operating simultaneously. We call them Thresholded Linear Weighted Sums (TLWS) and Primary Factors (PF).

The TLWS method is a minimally complex way to generalize from the binary-value concept that, if a binary-value node Z depends on binary-value variables X_1 and X_2 connected by "ands," then Z is 0 (false) unless *both* X_1 and X_2 are 1 (true). The generalization is that if Z is determined by a vector of variables $\boldsymbol{X}$, and we say that the elements of $\boldsymbol{X}$ are connected by "~ands," then—as an approximation—we assume that each element of $\boldsymbol{X}$ has a threshold value, implying a threshold vector $\boldsymbol{TH}$. By a threshold value, we mean that if $X_i < TH_{i,}$ then threat Z, not just X_i, is as low as possible (0 on a continuous scale or 1 on a discrete scale). We then assume relative weights for the elements, defining a vector $\boldsymbol{W}$. The TLWS algorithm, then, is

$$Z = \begin{cases} W \bullet X \text{ if } \min\left(X - TH, \text{factors}\right) \geq 0 \\ 0 \text{ otherwise} \end{cases} \tag{1}$$

Here $W \bullet X = \sum_i W_i X_i$ and $\min\left(X - TH, \text{factors}\right) \geq 0$ means that $\left(X_i - TH_i\right) \geq 0$ for all *i*.

The Primary Factors function is an alternative that applies to a first approximation if the node's value is dictated by the largest of the factors contributing to it (*P*), albeit with some possible upward adjustment reflecting the size of the second-largest factor (*S*). The continuous version of the formula is

$$Z_0 = P + \left\{\frac{S}{P}\right\}^2 \tau \text{ and } Z = \begin{cases} Z_0 & \text{if } 0 \leq Z_0 \leq 10 \\ 0 & \text{if } Z_0 < 0 \\ 10 & \text{if } Z_0 > 10 \end{cases} \tag{2}$$

where τ is a tuning parameter, which we have typically set to 2. The formula exhibits inconsequentially odd behavior for small values of P and S, but is otherwise a good heuristic.

The ambiguous influences denoted by +/- in factor trees can be due either to conflicting underlying processes (i.e., those represented by higher-resolution variables) that will be resolved in time or by stochastic effects.

Uncertainties can be represented by varying the factors parametrically, by using probability distributions, or both. In most RAND work on analysis under uncertainty the parametric approach has been favored because it retains visibility of cause-effect relations and defers assumptions about probabilities until the end of analysis (Davis 2012). An additional consideration is that the factors in such problems are often probabilistically correlated, which makes probabilistic work difficult.

To implement the ideas Davis and O'Mahony settled on the *Analytica*® modeling platform for reasons discussed in their report. The most important was that it is a visual-modeling language largely understandable to people who are not "real" programmers.. Also, the platform makes it extremely easy to conduct exploratory analyses—seeing how outputs change as numerous input variables are changed simultaneously. Figure 2 illustrates this. It is taken from their work on public support for terrorism. A single display can show the *simultaneous* effects of numerous factors, allowing the viewer to see good and bad "regions" (factor combinations). Such uncertainty analysis can assist "robust decision making" for which a considerable literature exists, much of it using a computational search called "scenario discovery" to find such regions (Lempert et al. 2006). It is far more sound to identify desirable and undesirable factor combinations than to claim to predict precisely what public support for terrorism will be given numerous uncertainties.

1.3 Background on Detecting Terrorists with Behavioral and Other Information

The other background element for our current work was a study surveying critically the science and technology base for detecting terrorists with behavioral indicators such as seen by Behavioral Detection Officers (BDOs) in an airport, officers at a military checkpoint, or intelligence officers viewing a large crowd at a public gathering (Davis, et al. 2013). One of the study's conclusions was that threat detection would continue to be very difficult and that a key element in any future success would likely be the fusion of many kinds of information. The study also concluded that the kind of information fusion needed was very different from that normally studied and that new research was called for on what the study called "heterogeneous information fusion." This meant fusing uncertain information that might be qualitative and quantitative, hard and soft, reliable and flakey, legitimate and bogus (or even maliciously deceptive), and fragmentary. Further, the fusion should be expected both to improve the probability of detecting the rare terrorist and reducing drastically the false-alarm rates that tend to plague detection efforts and cause harm to those incorrectly identified as possible threats.

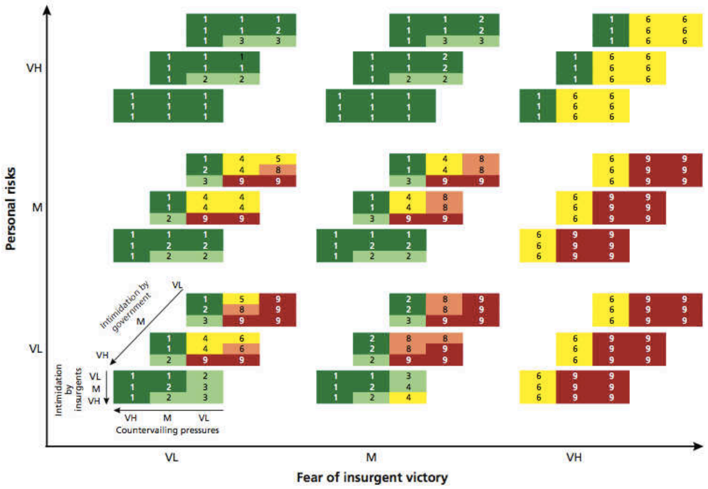

Note: Color indicates estimated strength of public support from very low to very high. The five independent variables shown are personal risks, fear of insurgent victory, countervailing pressures, intimidation by insurgents, intimidation by government. Other determinants of the results are held constant in this figure.

Figure 2: An illustrative multi-dimensional display of public support for terrorism.

In subsequent research we began studying the fundamentals of such heterogeneous fusion. This is the subject of the remaining part of the paper, but with emphasis here on the role that a model plays.

2 HETEROGENEOUS INFORMATION FUSION FOR THREAT DETECTION

We saw our research as addressing how to deal with the disparate classes of fragmentary and uncertain information. The challenges are different from those addressed in pattern recognition, machine-learning, predictive analytics and other data-rich empirical approaches that increasingly are exploiting "big data." Our approach attempts to add theory, structure, and rigor to inference processes that not only have heterogeneous "data," often sparse, but that must also include human subjectivism and analytic art. Our information fusion is perhaps akin to that of a fictional detective who uses fragments of information to piece together notions of whether an individual had the desire, means, and opportunity to commit murder. Solid empirical data should also be used wherever possible, as in establishing "base rates" for Bayesian inferences.

In the following pages we discuss only one aspect of our research on heterogeneous fusion—the role of a *causal* cognitive/behavioral model, as distinct from an empirical-statistical model. It is fortunate that comparatively few terror plots occur in a given year in the US (typically fewer than 10) (Stom et al. 2010), far less than would be needed to build a complex empirical-statistical model. Although our current model is deliberately static for simplicity, the causal models used might in the future be dynamic, perhaps incorporating elements of system dynamics and agent-based modeling. Indeed, a factor tree can be seen as a simplified snapshot in time of a systems dynamic influence diagram. In more data-rich environments, the kinds of issues that we are addressing would also be good fodder for Bayesian-net applications with which our approach has interesting although subtle relationships.

2.1 Objectives for a Causal Model

In classic text-book accounts of information fusion the signal received includes the information needed, although often amidst a great deal of noise. In assessing the threat of terrorism posed by an individual, however, the information obtained may be far removed from what we are interested in (is the individual a threat?). We do not observe the internal workings of his mind and, only seldom, does intelligence uncover direct information about operational plans. Instead, the information may be about associations with other people, travel behavior, courses taken, or troublesome but ambiguous comments made in bars, emotion-inducing meetings, or social media. It is from such fragmentary information that judgments must often be reached. The judgments should not be precise, but rather cautious estimates of relative likelihood. A primary objective is to identify those individuals who merit closer attention, whether passive or active, cautionary or preemptive. Some of those individuals will turn out to pose no threat at all. The desire is to focus resources on those individuals that are *relatively* likely to pose a threat. A second objective is to improve estimates of that threat potential over time so that those who are not threatening are recognized as such. To put it candidly, the intent is not just to put individuals on some law-enforcement watch list, but to remove individuals from such watch lists when appropriate.

As recognized in Bayesian-net research, a reasonably good causal model is needed for turning fragmentary information into inferences about something larger. Relating this to the familiar, consider again the fictional detective who finds information over time relating to intent, means, and opportunity. The individual fragments mean little (except in the instance where the act of murder is observed directly), but combining the fragments can be very meaningful—so much so that it is a core element of our criminal justice system that the prosecution should demonstrate that all elements are present when attempting to convict an individual.

2.2 An Example: the Propensity for Terrorism Model (PFT)

2.2.1 Basic Structure and Definitions

For our study we developed a variant of the model described in Section 1.2. As the name suggests, the *Propensity for Terrorism* (PFT) model focuses on the factors influencing an individual's propensity to commit terrorism, and, thus, the threat posed by the individual. Although it has not been separately validated by social-science research, it builds heavily on the earlier work (Davis and Cragin (2009); Davis, Larson, et al., 2012; Davis and O'Mahony, 2013). Thus, it seemed to us a reasonably credible example to use in our research. The PFT factor tree is shown in Figure 3, albeit in a somewhat truncated form. In practice, we used only the top layer that asserts that the threat *T* posed by an individual is a function of that individual's motivation (*M*), perception of terrorism's legitimacy (*L*), capability-opportunity (*CO*), and acceptability of costs (*A*). The model asserts that the threat posed is driven primarily by *M*, which could be for a cause, activity, adventure, etc. Moving rightward, we encounter an important but subtle concept.

We define *L* as the degree to which the individual sees terrorist violence (attacks on noncombatants) as legitimate if motivation is at or above its threshold level. This doesn't mean that he *is* motivated. Rather, it is a definitional trick to improve the probabilistic independence of variables. With this definition, the magnitude of L reflects reasons for seeing legitimacy that hold even if there is no motivation (the reason may be sociopathy and a love of violence) and will approximate legitimacy if motivation is even higher than threshold. Errors in the approximation are irrelevant for motivation below threshold because they will not affect the estimate of *T*. There will be no error if the individual rejects the terrorist violence independent of M. The error that may exist is for the individual whose sense of legitimacy is motivation dependent. Our approximation assumes that the sense of legitimacy will not be much different if motivation is very high rather than medium (the usual threshold setting). Or, more precisely, we assume that the threat estimate T will not be sensitive to such differences. Similarly, we define *CO* is a measure of capability-opportunity for an act of terrorism, assuming threshold motivation and legitimacy. And, finally, *A* measures the degree to which the individual sees the costs and risks of his terrorist action as acceptable, given that he is reasonably motivated and has at least threshold levels of capability-opportunity and legitimacy.

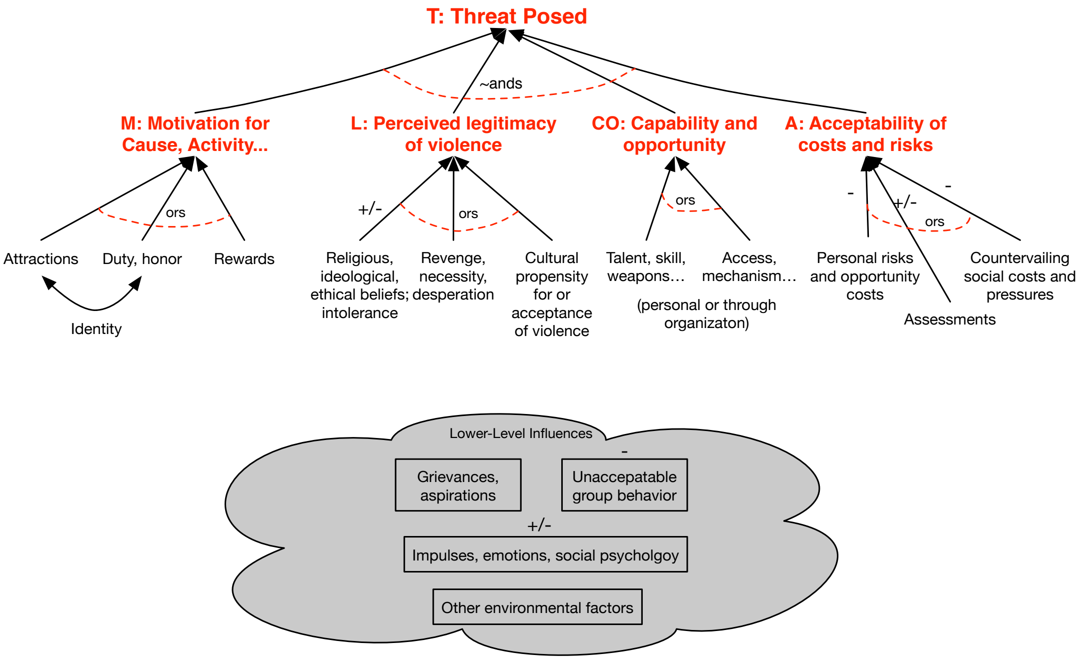
T: Threat Posed
~ands
M: Motivation for Cause, Activity...
L: Perceived legitimacy of violence
CO: Capability and opportunity
A: Acceptability of costs and risks
ors
Attractions
Duty, honor
Rewards
Identity
+/-
ors
Religious, ideological, ethical beliefs; intolerance
Revenge, necessity, desperation
Cultural propensity for or acceptance of violence
ors
Talent, skill, weapons…
Access, mechanism…
(personal or through organizaton)
-
+/-
-
ors
Personal risks and opportunity costs
Countervailing social costs and pressures
Assessments
Lower-Level Influences
Grievances, aspirations
Unacceptable group behavior
-
+/-
Impulses, emotions, social psycholgoy
Other environmental factors

Notes
1. “ands and ors” apply strictly only to binary case
2. (+),-, +/-: influence is positive, negative, or a mix
3. Lower-level influences may affect multiple nodes

Figure 3: A truncated factor tree of propensity for terrorism.

The innovation of this unusual set of definitions simplified the problem: the definitions mean that the factors of the model are approximately mathematically and probabilistically independent. Thus, if $\Pr(M,L,CO,A)$ is the joint probability for the four variables defined in this way,

$$\Pr(M,L,CO,A) \approx \Pr(M)\Pr(L)\Pr(CO)\Pr(A) \qquad (3)$$

This is because, as defined, L, CO, and A do not depend on each other (although they depend on the thresholds). In practice, the quality of the approximation depends primarily on whether the values of *M, L, CO,* and *A* are estimated within the sprit of the factors’ subtle definitions. A given source of information on *M* and *L* might, instead, estimate *L* as being high because that source believes that motivation *implies* legitimacy. If so, the data will be correlated even though the theory depends on it not being. Because of such possibilities, our method includes mechanisms for over-riding Eq. (3) to insert specific correlations where they are recognized. Also, our method includes adding parameterized correlation functions to see how strongly threat assessments depend on the assumption of independence. At a meta level, we are working on methods accounting for correlations across reports, as when they stem from sources with shared biases.

2.2.2 Beyond the Rational Actor

Another innovation was that the model structure of Figure 3 is deliberately not that of the rational-actor model. A rational actor would merely compare costs and benefits of different actions, focusing on their subjectively assessed “expected value” as implied by a utility function and notions about the consequences of options. That approach does not adequately capture important aspects of human decision making. In particular, individuals are often driven by emotions to do things that in retrospect they regard as unwise. Individuals also make numerous judgments based on wired-in heuristics that incorporate cognitive biases reviewed in Daniel Kahneman’s book for a broad audience (Kahneman 2011) and his earlier Nobel Prize speech (Kahneman 2002). In contrast, other important sources discuss the benefits of intuitive or naturalistic decision making despite human "biases" (Gigerenzer and Selten 2002; Klein et al. 2006a; Klein et al. 2006b). After years of sometimes acrimonious debate, scientific closure is occurring between those who emphasize watching out for cognitive biases and those who emphasize the virtue of intuitive decision making. The answer is that "both are right," but that the balance depends on context (e.g., deliberate decision making in peacetime versus the heat of battle). Some of this is reflected in Kahneman's 2011 book, but an earlier review written while the debate was still hot draws implications for decision-aiding that are still valid (Davis et al. 2005).

Another fundamental problem with the rational-actor model is that it depends on the individual having a *stable* utility function, with the actor not changing what he or she is trying to optimize over time. In reality, de facto utility functions often *emerge* from the course of events. As a result, they may be path dependent and temporally unstable as new events occur and context changes.

It is sometimes noted that more nearly realistic results can be obtained with the rational-actor model if the utility function captures such individual values as altruism and religion. For example, some terrorists truly believe that martyrdom will mean a bliss-filled eternity, will serve a deity, and will advance the cause of which the individual is part. That, however, does not address the instability problem. Suppose that an individual acts in the moment believing that his action is appropriate, but—a month later—realizes that it was foolish, wrong, or even evil. Yes, one could say merely that his utility function has changed, but is that useful? Is it not better to acknowledge that such individuals do not have stable utility functions? That said, the PFT model incorporates rational-analytic decision making as a special case.

Such issues are discussed in a recent National Academy report (National Research Council 2014, 35ff) to which one of us (Davis) contributed on this subject. It points to a significant literature (page 36), including a thoughtful book on deterrence theory (Morgan 2003) that discusses how political leaders often do not even know their values and utility function until *after* engagement, debate, negotiation, and iteration. This is closely related to the phenomenon described in the literature on "wicked problems," which notes that solutions often *emerge* rather than being the predictable solution to the original problem conception (Rosenhead and Mingers 2002).

2.3 Alternative Fusion Methods

Because our application is so different from that in more usual data-driven work and because of the heterogeneity of information, we had to develop a number of alternative fusion methods. These drew, of course, on the classic literature of Bayesian analysis (Gelman and Shalizi 2010) and Dempster-Shafer theory (Shafer 1976), and also the more recent Dezert Smarandache theory (Smarandache and Dezert 2009a; Smarandache and Dezert 2009b), and several others as surveyed briefly in our earlier work (Davis, Perry, et al., 2013).

We considered five types of fusion method: (a) purely subjective, (b) nonlinear algebraic, (c) Bayesian, (d) quasi-Bayesian, and (e) a new entropy maximizing method (MEMP), which actually optimizes a weighted sum of entropy-maximizing and penalty functions. We also paid considerable attention to Bayesian-net methods, although we did not employ them because most related research focuses on data-rich circumstances rather than those of our study. Nonetheless, the literature is highly relevant and some Bayesian-net approaches also emphasize the core importance of underlying causal models (Pearl 2009).

We had to derive novel features for methods (a)-(e) because of our context. Our nonlinear algebraic methods make use of the TLWS and PF methods discussed earlier. Our Bayesian method required us to concoct alternative "generic" likelihood functions and to demand that analysis experiment with the range of such functions because the "real" likelihood function, to the extent that it exists, is often unknowable. Our quasi-Bayesian approach is a variant in which the fusion analyst is urged to construct a subjective context-specific likelihood function. A "sticky" variant of the Bayesian methods allows the analyst to hedge by sticking—to some extent—with the prior assessment rather than replacing it with a Bayesian update. This is a special case of how our fusion methods have to account, yet again subjectively, for the credibility, salience, and ultimate reliability of the various reports on an individual. Finally, we introduced a new approach that pivots from the perspective of information-theory entropy. Technically, the approach (developed by Hollywood) uses nonlinear programming for fusion. It maximizes an objective function that includes a weighted sum of entropy-maximization terms and terms minimizing contradictions with reports, such as a claim that a person's motivation is in the medium-to-high range, but no lower or higher. The method yields estimates of threat level that are as conservative (i.e., uncertain, in an information-theoretic sense) as possible given what has been reported, but with recognition that the reports' assertions may not be correct (i.e., they imply "soft" constraints. The method has no difficulty fusing directly conflicting assertions and does not depend on the order in which reports are processed. Further, its complexity grows rather slowly with the number of assertions rather than exponentially, as with other fusion methods.

Another significant aspect of our approach was to design for analytic flexibility. For example, we recognized that, depending on the information available, it may be better to fuse first at the factor level (improving estimates of factors across reports), to estimate threat in each report and then fuse those threat estimates across reports, or some mixture. Further, we sought an approach encouraging competitive streams of analysis with different causal models, assumptions, and analyst judgments. Fusion across streams might be needed occur early, along the way, or only at the end. Thus, the approach and the software platform for experimentation had to allow for and facilitate analytic artistry.

3 ILLUSTRATIVE RESULTS SUGGESTING A VISION FOR HETEROGENEOUS FUSION

To pursue our basic research on heterogeneous fusion we constructed synthetic data reflecting the kinds of fuzzy, ambiguous, conflicting, and sometimes misleading information that might be available in a real application. This served a function analogous to scenario spinning in other fields. Figure 4 illustrates results for one cases. Figure 4a shows the probability densities obtained after fusion across reports in Stream A of analysis of the threat posed by an individual named Harry. Results are shown as a function of fusion method used to fuse across reports. Figure 4b shows an aggregate summary: the mean probabilities that Harry is a threat (T between 6 and 10), in the gray area (T between 4 and 6), or a non-threat (T between 0 and 4). The "slicer bars" at the top indicate major contributors to the result beyond the choice of fusion method. For the example, the TLWS method was used to estimate threat by combining factor values, factor values were fused first before estimating threat, nominal threshold values were used in the TLWS calculation, the reports were processed in the order received, Bayesian calculations were accomplished with Quasi Bayesian likelihood functions, report weights were entered factor-by-factor for each report, and the primary-factors calculation ignored reports with quality factors below 0.5. Such parameter values can be changed interactively by clicking through their menus (note arrows). This is *exploratory analysis*, i.e., viewing results as a function of numerous variables as they are changed simultaneously. Other inputs are also uncertain, but are suppressed for the example. In some of our synthetic cases, the threat assessments are even more starkly different as a function of fusion details. It sometimes even matters in what order reports are processed because of heuristics and simplified likelihood functions. Such details are irrelevant here but the examples illustrate how in our work uncertainties are highlighted, rather than suppressed. For the example, there are substantial differences in result across method (and the values of the other variables). When that occurs, the analyst needs to go back into details and judge which of the methods and assumptions are likely to be most and least reliable, or most apt, for the specific context.

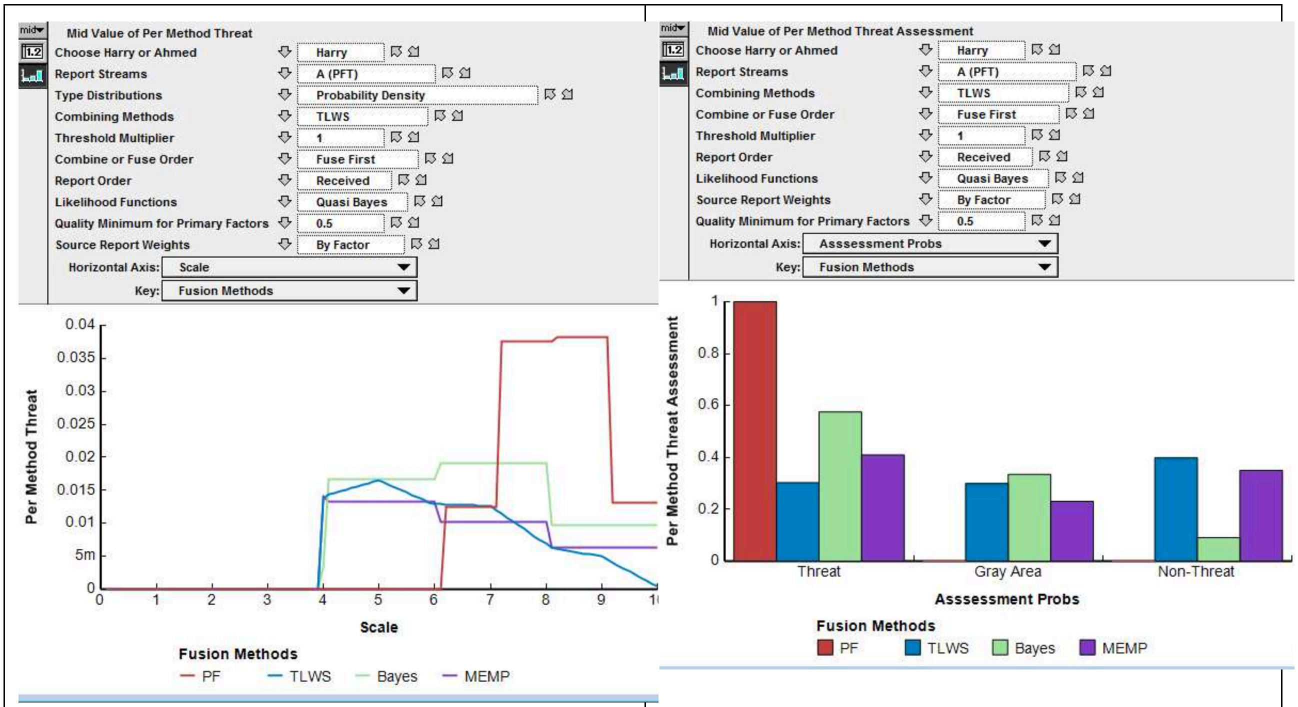

Mid Value of Per Method Threat
Choose Harry or Ahmed: Harry
Report Streams: A (PFT)
Type Distributions: Probability Density
Combining Methods: TLWS
Threshold Multiplier: 1
Combine or Fuse Order: Fuse First
Report Order: Received
Likelihood Functions: Quasi Bayes
Quality Minimum for Primary Factors: 0.5
Source Report Weights: By Factor
Horizontal Axis: Scale
Key: Fusion Methods

Mid Value of Per Method Threat Assessment
Choose Harry or Ahmed: Harry
Report Streams: A (PFT)
Combining Methods: TLWS
Combine or Fuse Order: Fuse First
Threshold Multiplier: 1
Report Order: Received
Likelihood Functions: Quasi Bayes
Source Report Weights: By Factor
Quality Minimum for Primary Factors: 0.5
Horizontal Axis: Asssessment Probs
Key: Fusion Methods

a. Probability Densities

b. Probabilities of Being Threat, Non-Threat, or in Gray Are

Figure 4: Illustrative results from prototype experiments.

4 CONCLUSIONS

Our initial experiments were gratifying. The methods were falling into place, the prototype analytical platform was operating, and we saw significant and useful consequences of going about heterogeneous fusion in different ways. This was “good,” not “bad,” because in this domain it is necessary to experiment with different ways to process information and different assumptions within the process of doing so. In our domain, analysis would be different in a context of trying desperately to identify individuals most plausibly posing a threat, so that resources could be immediately allocated to look into them, and in a more usual context of attempting to objectively assess the threat posed (or not posed) by an individual. In the former case, one would be “looking for trouble” and interested in distributional tails; in the latter case, one would concerned about characterizing knowledge as fairly and soberly as possible. In still other cases, the analysis might be about identifying what fragment(s) of information have been pivotal in identifying someone as a threat. With knowledge of that, it might be easier to spot inappropriate leverage of very dubious information or to identify what new information might lead to exoneration.

Although we are still speculating about what may be possible, our initial theoretical work and experimentation has been encouraging. Perhaps needless to say, much remains to be accomplished. Most obviously, we need to experiment using “real” data and interacting with “operators.” We are merely at the beginning of what should be years of research into heterogeneous fusion.

ACKNOWLEDGMENTS

The research described here was sponsored by the Office of Naval Research and the Office of the Secretary of Defense.

REFERENCES

Carifio, J. and R. Perla 2007. "Ten Common Misunderstandings, Misconceptions, Persistent Myths and Urban Legends about Likert Scales and Likert Response Formats and their Antidotes." *Journal of Social Sciences* 3(3): 106-17.

Davis, P. K. 2011. "Primer for Building Factor Trees to Represent Social-Science Knowledge." *Proceedings of the 2011 Winter Simulation Conference*

Davis, P. K. 2012. *Some Lessons From RAND's Work on Planning Under Uncertainty for National Security*. Santa Monica Calif.: RAND Corp.

Davis, P. K., and Kim Cragin, eds. 2009. *Social Science for Counterterrorism: Putting the Pieces Together*. Santa Moncia, Calif.: RAND Corp.

Davis, P. K., Jonathan Kulick, and Michael Egner. 2005. *Implications of Modern Decision Science for Military Decision Support Systems*. Santa Monica, Calif.: RAND Corp.

Davis, P. K. Eric Larson, et al. 2012. *Understanding and Influencing Public Support for Insurgency and Terrorism.* Santa Monica, Calif.: RAND Corp.

Davis, P. K., and Angela O'Mahony. 2013. *A Computational Model of Public Support for Insurgency and Terrorism: A Prototype for More General Social-Science Modeling*." Santa Monica, Calif.: RAND Corp.

Davis, P. K., Walter S. Perry, Ryan Andrew Brown, Douglas Yeung, Parisa Roshan, and Phoenix Voorhies. 2013. *Using Behavioral Indicators to Help Detect Potential Violent Acts*. Santa Monica, Calif.: RAND Corp.

Gelman, A., and C. R. Shalizi. 2010. "Philosophy and the Practice of Bayesian Statistics."

Gigerenzer, G., and R. Selten. 2002. *Bounded Rationality: The Adaptive Toolbox*. Cambridge, Mass.: MIT Press.

Kahneman, D. 2002. *Maps of Bounded Rationality: A Perspective on Intuitive Judgment and Choice (Nobel Prize Lecture)*.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Klein, G., B Moon, and R. R. Hoffman. 2006a. "Making Sense of Sensemaking 1: Alternative Perspectives." *IEEE Intelligent Systems* 2 (4): 70–73.

Klein, G, B Moon, and R R. Hoffman. 2006b. "Making Sense of Sensemaking 2: A Macrocognitive Model." *IEEE Intelligent Systems* 21 (5): 88–92.

Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes. 2006. "A General Analytic Method for Generating Robust Strategies and Narrative Scenarios." *Management Science* 4, April 514–28.

Morgan, P. M. 2003. *Deterrence Now*. Cambridge: Cambridge University Press.

National Research Council. 2014. *U.S. Air Force Strategic Deterrence Analytic Capabilities: An Assessment of Methods, Tools, and Approaches for the 21st Century Security Environment*. Washington, D.C.: National Academies Press.

Pearl, J. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge, Mass.: Cambridge University Press.

Rosenhead, J, and J. Mingers. 2002. "A New Paradigm of Analysis." In *Rational Analysis or a Problematic World Revisited: Problem Structuring Methods for Complexity, Uncertainty and Conflict*, edited by Jonathan Rosenhead, and John Mingers, 1–19. Chichester, UK: John Wiley & Sons, Inc.

Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton, New Jersey: Princeton University Press.

Smarandache, F. and J. Dezert, eds. 2009a. *Advances and Applications of DsMT for Information Fusion*. Rehoboth: American Research Press.

Smarandache, F. and J. Dezert. 2009b. "An Introduction to DsMT." In *Advances and Applications of DsMT for Information Fusion*, edited by Florentin Smarandache, and Jean Dezert, Rehoboth: American Research Press.

Stom, K., J. Hollywood, D. Snyder, K. McKay, and J. Boon. 2010. "Building on Clues: Examining Successes and Failures in Detecting Terrorist Plots, 1999-2009." RTI International. Research Triangle Park, North Carolina.

AUTHOR BIOGRAPHIES

PAUL K. DAVIS is a senior principal researcher at RAND and a professor of policy analysis in the Pardee RAND Graduate School. He is a graduate of the University of Michigan (B.S.) and Massachusetts Institute for Technology (Ph.D. in Chemical Physics). His research has included such diverse subjects as strategic planning; deterrence theory; counterterrorism theory; modeling, including cognitive modeling of adversaries and multiresolution modeling more generally; and complex information fusion to assist threat detection. His email address is paul_k_davis@me.com.

WALTER L. PERRY is a senior information scientist at RAND. He received his Ph.D. at George Mason University after retiring from the U.S. Army's Signal Corps. He has taught electrical engineering, computer science, statistics, and mathematics. His research has included leading official reviews for the Army of operations in Kosovo, Afghanistan, and Iraq. More technically, his research has included methods for data fusion and information-processing, and for complex information fusion to assist threat detection. His email address is walt@rand.org.

JOHN S. HOLLYWOOD is a full operations researcher at RAND and a professor of policy analysis at the Pardee RAND Graduate School, where he applies qualitative and quantitative analytics to security policy, including criminal justice, homeland security, counterinsurgency, and defense systems. He holds an SB in Applied Mathematics and a Ph.D. in Operations Research from the Massachusetts Institute for Technology. His email address is johnsh@rand.org.

DAVID MANHEIM is a doctoral fellow in the Pardee RAND Graduate School and an Assistant Policy Analyst at RAND. His current focus is disaster risk, recovery, and mitigation, including both natural hazards and terrorism. He holds a B.S. in Mathematics from Lander College, where he concentrated on abstract mathematics and financial modeling. His email address is dmanheim@rand.org.