NATIONAL DEFENSE RESEARCH INSTITUTE

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis.

This electronic document was made available from www.rand.org as a public service of the RAND Corporation.

Skip all front matter:

Support RAND

Purchase this document

Browse Reports & Bookstore

Make a charitable contribution

For More Information

Visit RAND at www.rand.org

Explore the RAND National Defense
Research Institute

View document details

Limited Electronic Distribution Rights

This report is part of the RAND Corporation research report series. RAND reports present research findings and objective analysis that address the challenges facing the public and private sectors. All RAND reports undergo rigorous peer review to ensure high standards for research quality and objectivity.

Using Behavioral Indicators to Help Detect Potential Violent Acts

A Review of the Science Base

Paul K. Davis, Walter L. Perry, Ryan Andrew Brown, Douglas Yeung, Parisa Roshan, Phoenix Voorhies

NATIONAL DEFENSE RESEARCH INSTITUTE

Using Behavioral Indicators to Help Detect Potential Violent Acts

A Review of the Science Base

Paul K. Davis, Walter L. Perry, Ryan Andrew Brown, Douglas Yeung, Parisa Roshan, Phoenix Voorhies

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Support RAND—make a tax-deductible charitable contribution at www.rand.org/giving/contribute.html

RAND® is a registered trademark

Cover photo by Karl Baron via flickr

© Copyright 2013 RAND Corporation

Preface

This report reviews the scientific literature relating to observable behavioral indicators that might, along with other information, help detect potential attacks, such as those by suicide terrorists or the laying of improvised explosive devices (IEDs). The report is intended to be of interest to officials contemplating future investments amidst tightening budgets, and to researchers and analysts. It deals with individual-level indicators and does not extend to detecting society-level phenomena, such as social movements or insurgent groups.

Our research built on prior RAND Corporation efforts, notably:

- Walter L. Perry, Claude Berrebi, Ryan Andrew Brown, John Hollywood, Amber Jaycocks, Parisa Roshan, Thomas Sullivan, and Lisa Miyashiro, *Predicting Suicide Attacks: Integrating Spatial, Temporal, and Social Features of Terrorist Attack Targets*, 2013.
- Paul K. Davis and Kim Cragin, eds., *Social Science for Counterterrorism: Putting the Pieces Together*, 2009.
- Thomas Sullivan and Walter L. Perry, "Identifying Indicators of Chemical, Biological, Radiological and Nuclear (CBRN) Weapons Development Activity in Sub-National Terrorist Groups," *Journal of Operational Research and Society*, Vol. 55, No. 4, April 2004, pp. 361–374.

Comments and questions are welcome and can be addressed to Paul Davis in Santa Monica (pdavis@rand.org), Ryan Brown in Santa Monica (rbrown@rand.org), or Walter Perry in Arlington, Virginia (walt@rand.org).

For more information on the RAND International Security and Defense Policy Center, see http://www.rand.org/nsrd/ndri/centers/isdp.html or contact the director (contact information is provided on the web page).

Contents

Figures

Tables

Summary

Government organizations have put substantial effort into detecting and thwarting terrorist and insurgent attacks by observing suspicious behaviors of individuals at transportation checkpoints and elsewhere. Related technologies and methodologies abound, but their volume and diversity has sometimes been overwhelming. Also, effectiveness claims sometimes lack a clear basis in science and technology. The RAND Corporation was asked to review the literature to characterize the base in behavioral sciences relevant to threat detection, in part to help set priorities for special attention and investment.

Purpose and Approach

Our study focused on the science base for using *new or nontraditional* technologies and methods to observe behaviors and how the data gathered from doing so might—especially when used with other information—help detect potential violent attacks, such as by suicide bombers or, as a very different example, insurgents laying improvised explosive devices (IEDs). Behavioral indicators may help identify individuals meriting additional observation in an operational context such as depicted in Figure S.1. For that context, security personnel at a checkpoint are assessing (blue oval) whether an individual poses some risk in the limited sense of meriting more extensive and perhaps aggressive screening, follow-up monitoring, or intercept. They obtain information directly, query databases and future versions of information-fusion centers ("pull"), and are automatically provided alerts and other data

Figure S.1
A Contextual View of the Detection Effort
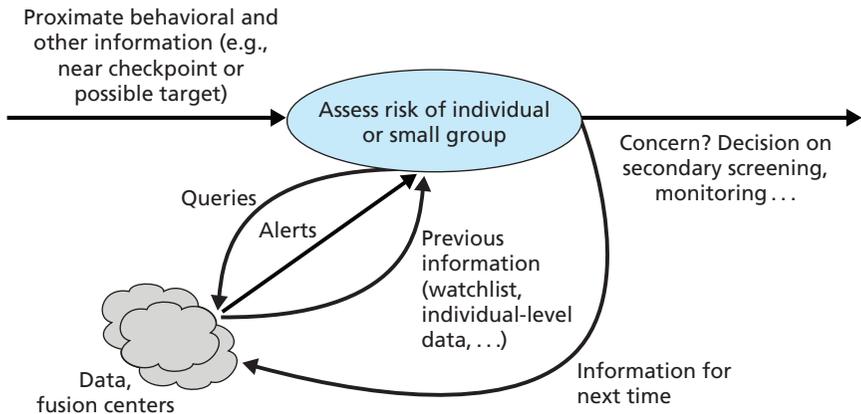
("push"). They report information that can be used subsequently. In some cases, behaviors of a number of individuals over time might suggest a potential ongoing attack, even if the individuals are only pawns performing such narrow tasks as obtaining information.

Figure S.1 refers to "other information" (top left). Although our study is concerned with detecting imminent threats rather than gathering broad information for internal security or intelligence, such information—perhaps accumulated over years—can play an important role. Where might that information be found, how might it be structured, and what indicators might be involved? We focus on what may be possible technically, without analyzing tradeoffs with privacy and civil liberties. We do, however, note troublesome issues raised by the technology and methods, point readers to an in-depth study of such matters by the National Academy of Sciences, and suggest research on ways to mitigate the problems.

Figure S.2 shows relationships among our key constructs. A base of technology and methods (left) allows detecting behavioral indicators (bottom right). Moving upward, these give signals about activities, which are grouped into activity classes called phases. Analysis can then assess whether the totality of information (including nonbehavioral

Figure S.2
Relationships Among Constructs

Other information → **Basis for concern**   **Overall assessment**, from indicators, of need for follow-up on individual or small group

Technology and methods for observation and analysis

Phases   Aggregate activities providing evidence (perhaps weak or ambiguous) that raises concern about subject

Activities   Activities providing evidence (perhaps weak or ambiguous) that raises concern about subject

Behavioral indicators of activities   Observables providing evidence of potentially threatening activity

RAND *RR215-S.2*

information) adds up to a basis for concern justifying more screening, monitoring, precautionary defensive measures, or preemptive action. Since detecting a "basis for concern" (i.e., need for further checking) will probably have a high false alarm rate, a system using this approach must be efficient and reasonable if it is to be acceptable.

Figure S.3 is a conceptual model showing phases within which lower-level activities occur. The model merely identifies where to look for information. As indicated at the bottom of the figure, there are many possible indicators of the activities and a number of technologies and methods to use in observing them. The model is merely heuristic rather than a rigorous decomposition or timeline. Activities may be performed by multiple individuals, not occur, or occur in different order. Some activities could occur in more than one phase.

Figure S.4 uses the "Developing Intent" phase of Figure S.3 to illustrate how phases, activities, and indicators relate to technologies and methods. For each of three activities, Figure S.4 shows potential indicators. The lowest box shows some relevant technologies and methods. The Developing Intent phase is unusual in that it includes

Figure S.3
Conceptual Model of Opportunities for Observation

Evidence of violent intent

Nonbehavioral data

Phases (need not all be sequential)

Developing Intent

Planning and Laying Groundwork

Immediate Pre-Execution

Execution

Aftermath

Activities

Psychological development

Psychological convergence

Recruitment or joining

Development of strategic priorities

Target identification, ISR

Materiel acquisition, testing and development

CONOPs development

Training, mission rehearsal

Other long-lead preparations

Psychological, physiological preps

Changed social patterns

Ritual practices

Deception, concealment

Logistical preparation

ISR

Deployment, positioning

Coordination, communication

Target shaping, feints

Main attacks

Post-attack reporting

Strategic communication

Adaptation

Protective measures

Indicators (nominally of specific activities, but potentially relevant to others as well, or to an overall judgment)

Technologies and methods

NOTES: ISR = intelligence, surveillance, and reconnaissance. CONOPs = concept of operations. Indicator-phase connections are nominal. Some activities can occur in more than the phase indicated. Associations of activities with phases are nominal. Some activities can occur in additional phases.

RAND *RR215-S.3*

Figure S.4
Illustration of Methodology

	Develop intent	

Activities

Psychological and emotional development	Psychological convergence	Recruitment or joining

Behavioral indicators

• Act so as to show hatred, prejudice, trauma, or shame • Exhibit strong motivation for prestige glory, status • Demonstrate intolerant black-and-white thinking • Indicate perception of "rightful place" or status, but with meager resources and attainment	• Approach life and decisions through ideological framework • Proselytize; reinforce practices and beliefs in others • Indicate reinforced commitment • Seek out materials with radical content • Change behavior at school or home	• Explore different organizations and strategies • Attend radical mosques or events • Form new connections in social network to known terrorist elements • Adhere strictly to organizational standards

Technology and methods

• All-source, real-time data fusion
• Whole-life data
• Online content "scrapers"
• Social network analysis
• Mobile device tracking
• Automated textual analysis of communications or online content

RAND *RR215-S.4*

early-in-life activities, such as might be observed by parents, neighbors, teachers, physicians, local law enforcement, and others long before an individual becomes involved in anything violent. In the main text, we discuss the phases separately, identifying generic activities and potential indicators.

Technology and Methods

We found it useful to highlight technologies and methods in three cross-cutting classes of data: (1) communication patterns, (2) "pattern-of-life" data, and (3) data relating to body movement and physiological state. Most of the methods suffer from signal-to-noise problems and false alarms; some are vulnerable to countermeasures.

Communication Patterns

Communications occurs in, e.g., face-to-face meetings, Internet chat rooms, and cell phones. Large commercial and intelligence-sector investments have yielded techniques to monitor and analyze these communications, which we also treat in three groups: online communications and analysis, text analysis and natural-language processing, and speech analysis.

Online Communications. Using data collected by online-content "scrapers" and subsequent human efforts, it is sometimes possible to infer motivations, emotions, and ideologies from online statements and actions (e.g., as with the New York Police Department following Twitter posts such as "people are gonna die like Aurora" (referring to the July 20, 2012, movie-theater shooting in Aurora, Colorado). Real-time social-media search tools can help monitor and track discussions of potential targets. Keystroke loggers and other programs can reveal past and current searches for materiel; registrations or payments to training programs; location information; and information on past searches. Changes of activity may occur when terrorists "go dark" before an attack, when logistical preparations for an attack are intense, or when calls for vengeance arise after an event such as Osama bin Laden's killing.

This and all the methods include false alarms (e.g., users may have benign reasons for their purchases or mere curiosity as they investigate troublesome websites), low signal-to-noise ratio, and vulnerability to such counters as burying information amidst innocuous communication or using anonymous or false accounts.

Text Analysis and Natural-Language Processing. Techniques to classify texts and analyze content are fairly well developed, although less so

with respect to emotion and intent. Explicit content may include bragging, ideological statements, or admiration for terrorist leaders. Textual analysis of *style* can detect word-usage patterns associated with typical attacker motivations and such emotions as anger, humiliation, and shame. Style of communication does not depend on content, i.e., on specific topic, and can suggest relationships and status within a social network. Textual analysis of style has been used for, e.g., detecting corporate fraud, terrorist interrogations, and criminal testimony. Natural-language processing can analyze massive amounts of text about which little is known in advance, classifying documents to be analyzed further by subject-matter experts. Clustering methods can identify concepts and such topics as weapons, tactics, or targets. Such mathematical techniques as latent semantic indexing can help understand concepts and have the advantage of being language-independent. Machine translation can often turn foreign-language texts into something analyzable without foreign-language expertise or language-specialized software. Speech-recognition technology can greatly increase the amount of text available for text analysis. It can also help identify individuals.

A primary shortcoming is nonspecificity; that is, detected patterns (even if apparently threatening) may be unrelated to any imminent threat, and their interpretation often depends on cultural and individual idiosyncrasies. A shortcoming of the research base itself is that much linguistic-style analysis has been done only on archival data; more testing and validation is needed with "real-life" data sets. Top researchers caution against expecting highly reliable detections or interpretations and suggest the need for very large data sets that reveal many cultural and individual differences.

Speech Analysis of Content. Several robust indicators exist for connecting vocal content and narratives with lying and deception. These include the subject (1) distancing himself from untruthful statements by, e.g., using the third person or otherwise seeming less verbally involved; (2) issuing discrepant statements; (3) providing less detail; (4) exhibiting less logical structure and less subjectively plausible stories; (5) providing less context; and (6) making less spontaneous corrections or claiming lapses of memory.

This approach's primary shortcoming in assessing deceptive or hostile intent is that interpreting lexical and vocal indicators of lying and deception depends on context, individual variability, and appreciation of nonthreatening explanations. Optimally, analysis has data on the individual's speech in a normal nondeceptive/nonhostile state. Where this is infeasible, the potential increases markedly for failed detections and intolerably many false alarms. Table S.1 summarizes

Table S.1
Considerations and Caveats: Detection and Analysis of Communication Patterns

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Online communication and activities	Extensive collection and analysis occurs today for commercial and intelligence reasons. Technologies and methods for analyzing such online activities are still relatively unproven in either academic or operational settings.	Given trends, even more and varied interactions will be available for collection.	Tools already exist. However, challenges for dealing with massive volumes of noisy data are formidable.	Methods have not been well validated in academic or operational settings. Low signal-to-noise ratio. Effects of encryption, using "code," using anonymizers, or moving offline.
Text analysis and natural-language processing	A considerable research base exists with numerous past applications. Even natural-language processing can be highly accurate in specific experimental settings.	Using operational data to train and to create baselines could improve detection of deception, hostility, or extremist patterns. Natural-language techniques, given training sets, could quickly analyze large amounts of data.	Online text is naturally occurring and publicly accessible, requiring only passive collection. Active elicitation of text or oral statements is possible in some security contexts, such as checkpoints or interrogations.	Context and cultural dependence. Inadequate testing in operational settings. Need for substantial data.

Table S.1—Continued

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Speech analysis: lexical and vocal cues	This has been validated in laboratory settings, including those specific to counterterrorism.	Advances in protocols for rapid assessment of speech patterns and content would have wide applicability for screening, checkpoint, or other situations involving conversations with security personnel.	Such analysis currently requires skilled security personnel asking questions and making judgments.	Physiological drivers, such as anxiety and changes in vocal tone, are individual-dependent. May be subject to counters, especially if criteria for judging are known.

results of our review for the assessment of communication patterns and content.

Pattern-of-Life Data

It is possible to analyze patterns of communication, travel, purchasing, and other matters using existing records and databases (many held by private industry). We discuss mobile-device tracking, using existing records, and machine learning for pattern detection. These raise profound social questions about what kind of data can and should be collected and analyzed.

Mobile-Device Tracking. Ubiquitous mobile devices provide a wealth of data on personal information, social relationships and networks, locations, patterns of movement and interactions, preferences, political opinions, the spread of information, and patterns of how opinions and preferences change. Also, mobile-device usage is related to the "Big Five" personality traits (openness, conscientiousness, extraversion, agreeableness, and neuroticism).

One shortcoming of such data is that much social networking through mobile devices is increasingly "muddy" and in many cases divorced from both intent to meet and intent to act in the offline "real

world." This complicates inference-making about patterns of communication and their connection to actual threat. And, of course, people can go offline.

Existing Records. It is sometimes possible to develop individual profiles from information about, e.g., experiences, behaviors, and relationships over time, and to provide context for assessing other incoming data. The data could come from school records, criminal records, interrogation reports, and so forth. Additionally, surveillance cameras are now common in public and business settings, allowing for the possibility of tracking individual patterns-of-life. Integrating such data requires analytic techniques, including those for all-source, real-time big-data fusion. Related analytic tools are increasingly available from such providers of cloud computing as Google and Amazon and social-media companies.

The shortcomings include, of course, the administrative, jurisdictional, legal, and database challenges of extracting and combining data across multiple sources and owners within and outside the United States. Accuracy matters for this type of analysis, whereas commercial applications often do not require high accuracy to improve the targeting of marketing efforts.

Machine Learning and Big-Data Analysis. Given the sheer magnitude of data, it is increasingly important to analyze information without the benefit of prior hypotheses or known points of comparison. "Supervised" machine-learning techniques use known data sets to train the algorithms, which can then classify data, identify relationships, and discover concepts. "Unsupervised" learning proceeds without the aid of such known prior knowledge. It seeks to find structure in the unlabeled data set. For example, researchers have used thousands of YouTube images for unsupervised detection of high-level features such as faces. Potentially, such techniques could recognize images suggesting imminent threat. Such machine-learning techniques have been applied to uncover fraud, to recognize deception in computer-mediated communication, and for predictive policing. Artificial neural-network models are promising and can be applied in real time. Video or image analysis and machine-learning techniques could be employed to find,

for example, such activities as shaving heads and prayer activities in martyrdom videos.

One shortcoming is that machine-learning techniques often require a large amount of data. At least in the public domain, sufficiently large databases of violent attacks and other events do not exist for topics such as terrorism. One innovative method for obtaining large, labeled data sets is to "crowd-source" collecting and labeling individual pieces of information.

Table S.2, which is analogous with Table S.1, is our assessment of the various approaches focused on records-based whole-life information.

Indicators from Physical Movement and Physiology

Behavioral science has identified many nonverbal behaviors that are statistically associated with emotional and psychological state and with deception or violent intent. These can be roughly categorized into (1) kinetics (including gross motor movements) and (2) observation of physiological state. As discussed in the main text, *many* findings are controversial among scientists, and between scientists and operators, but our summary assessment follows.

Kinetics and Gross Movement. Existing technology can collect data on kinematic patterns (movement). Surveillance and reconnaissance platforms (e.g., tower cameras or drone systems) can monitor individuals as they maneuver before an attack. Video systems can view individuals before attacks and collect information on individuals who frequent potential attack sites, providing a baseline for identifying individuals engaged in pre-execution activities. For example, the gait of people who may be carrying weighted objects, such as IEDs, may be compared against baseline "gait signatures." Existing recordings of terrorism incidents can provide data for setting parameters on new analysis tools. The Defense Advanced Research Projects Agency (DARPA) has funded biometric technologies for identification at a distance and for early warning. Another approach seeks to automate recognition of potentially threatening body postures or poses.

Incorporating emotion into machine-learning methods may increase their future utility. To do so, "affective computing" may need

Table S.2
Considerations and Caveats for Pattern-of-Life Data

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Mobile-device tracking	Algorithms to predict individual movement patterns, preferences, etc., have been developed and validated in laboratory and experimental settings, but can benefit from more naturalistic validation.	Mobile devices will continue to add connectivity features that enable tracking (e.g., location and motion sensors, Near Field Communication chips).	Mobile device tracking may require device-owner permissions or cooperation of communications network providers.	Not traveling with or turning off device will defeat methods based on mobile-device whereabouts. Mobile-to-mobile communication is often divorced from "real-life" behaviors and intent.
Pattern-of-life data	Validating techniques to analyze large amounts of pattern-of-life data may be difficult in academic settings. Commercial data sets and analytic tools are increasingly available.	Pattern-of-life data may allow integrating disparate data types to build fuller behavioral profiles on individuals of interest. Accessing and integrating data is an issue.	Measurement does not require active or voluntary consent. However, access to various databases held by commercial or private sources may be necessary.	Pattern-of-life data may be vulnerable to "cover" activities and behaviors. Databases and algorithms for detecting threatening patterns are in early development.
Machine learning and big-data analysis	Machine-learning techniques have been extensively used and validated in experimental and some applied settings. Such techniques have been used in national security and law enforcement.	Machine-learning and big-data analysis may "discover" unknown patterns or activities hidden in large amounts of data, but massive amounts of data are needed for training.	Measurement does not require active or voluntary consent. A large amount of data or a strong hypothesis regarding relevant activity is required.	Learning techniques are probabilistic and vulnerable to noisy data. Current systems do not understand how to associate behaviors of multiple threatening individuals.

to select from various psychology and neuroscience findings and theories of emotion (e.g., "appraisal models"). Often, subsystems of monitoring and interpretation of stimuli can be computationally modeled. Improvements are possible when distinguishing between emotional states that differ in arousal, such as anger and sadness. Methods being developed to analyze gait signatures could be applied to such existing commercial technologies as cameras used for the Microsoft Kinect and Nintendo Wii game systems' motion-capture capability.

Human observers and analysts may also be employed, but results depend on such factors as training, individual talent, and observer bias. Detecting deceptive movements is easier for people experienced in employing the same deceptive movement patterns. People are best able to detect emotions associated with gait when the human walkers are expressing anger. Inference from merely a single stride can be highly accurate, suggesting that gait can be used to recognize affect. Performance varies by individual, and women may be better than men at determining actions and recognizing emotions from movements such as walking.

Analysis of kinetics and gross motor movements should apply to a wide variety of security contexts, although validation in naturalistic settings is needed and, as often occurs in looking for behavioral indicators, the indicators may arise for benign reasons, such as people being anxious at security screenings or checkpoints.

One challenge for gait analysis is that current detection systems and protocols are often built using simulated behaviors (e.g., with actors). More naturalistic (real-world) observations are needed.

Physiological State and Reactions. Observing physiological state holds promise for detecting deception and other behaviors. We touch upon polygraph testing, other measures of peripheral nervous system response, electroencephalograms (EEGs), vocal stress, and facial-expression analysis.

Polygraph testing has long been employed and found useful as *part* of an investigatory process (particularly because people often "open up" in the process), but is not by itself reliable. A great divide exists between the bulk of the academic community, who remain quite skeptical, and the community of "operators," who insist on its useful-

ness as one tool in a process. Newer approaches using some of the same physiological signals as in polygraphs (heart rate, blood pressure, etc.) are in development with respect to detection of potential deception or hostile intent.

New technologies using electroencephalograms (EEGs) allow some physiological features to be observed without "wiring up" individuals, sometimes at a distance, and sometimes covertly or surreptitiously, as with using heat-sensitive cameras to detect capillary dilation and blood flow to the face and head. There is some evidence of unique value in indicating deception or imminent action by an individual if baseline information is available for that specific individual ahead of time or if credible intelligence about a possible attack is available. Most of the technologies are in a relatively early stage of development, but there does seem to be potential. Measurement of physiological signals closer to the central nervous system (i.e., the brain) holds the most promise for detecting guilt and behavioral intent.

Evidence of vocal tension and higher vocal frequency may also be predictors of stress and deception, and a few observable aspects of speech are much more difficult for an individual to *control* than other indicators of deception, but countermeasures that obscure differences from the baseline of normalcy are definitely feasible.

Humans appear to share universal facial expressions indicative of underlying emotional and motivational states. Cultural differences seem to affect only secondary aspects of facial expressions. The seven fundamental emotions—anger, disgust, fear, happiness, sadness, surprise, and contempt—are displayed on the face with some fundamental features that are generally recognizable on all humans (barring neurological impairment). For the purposes of detecting pre-incident indicators, the most promising domain of facial expression analysis involves facial micro-expressions—involuntary expressions of emotion appearing for milliseconds despite best efforts to dampen or hide them. Whether the relevant behavior is smuggling weapons, traveling on forged documents, or hiding anger or anxiety near security officials, facial micro-expressions can be important indicators of deception or some kind of mal-intent.

At least currently, the two primary problems with using physiological indicators are (1) nonspecificity (the indicators may stem from many causes, most of them benign) and (2) individual differences (the observables that indicate attack or deception differ markedly across individuals, which requires establishing sound individual-centered baselines). Countermeasures are a problem with polygraphs, but perhaps less so with EEG methods. Even with polygraphs, empirical results have varied. Some drugs, for example, have not reduced detection rates as expected, but physical training can be effective as a countermeasure. Controlling vocal stress indicators is difficult, but countermeasures can obscure distinctions between baseline and stressed behavior. Facial expressions suffer from the same problems of nonspecificity, but they have the advantage of being more closely linked to motivational state and intent than are other physiological signals. Individual differences are also important: A psychopathic attacker, for example, might be more likely to show micro-expressions of "duper's delight" while passing through a checkpoint undetected, while a nonpsychopathic attacker might instead show micro-expressions of fear (as would a perfectly harmless nervous traveler).

While the link between micro-expressions and deception is well evidenced, utility in security-related settings is another matter. Coding emotional expressions currently involves hours of labor to analyze seconds of data, making this technique unsuitable for use in real time at checkpoints or other screening areas. However, a training system appears to increase the capacity of individuals to detect facial expressions and micro-expressions with demonstrated evidence of effectiveness in clinical populations.

Recognition of emotional expressions based on automated algorithms and computation is still in its infancy, but this is an active field of development, and improved algorithms are likely to yield greater accuracy and robustness. Furthermore, as with many pre-incident indicators of attack, emotion-recognition algorithms that fuse multiple parameters seem to perform much better than inferring emotional state simply from facial expressions alone.

Of course, checkpoints or other security environments are dynamic locations where it is difficult to capture high-resolution video

(or audio or physiological data) for individuals, but such detailed infor-
mation is often necessary for effectiveness.

Table S.3 is our assessment of how the approaches based on
detecting intent from physiological indicators stand in terms of matu-
rity, potential, measurability, and vulnerability to countermeasures.

Cross-Cutting Themes

A number of cross-cutting issues arose in our review. These suggest a
notional framework for thinking about detection systems. Although
the relevant metrics have by no means been defined as yet, much less
metrics that take into account cost-effectiveness, a goal for future
analysis might be to place something like the framework shown in
Figure S.5 on a solid scientific and analytic basis. Although it is surely
not yet "right" or well defined, this framework conveys a sense of what
is needed for sounder discussion. Further, despite its shortcomings, we

Figure S.5
A Notional Framework for Characterizing an Overall System
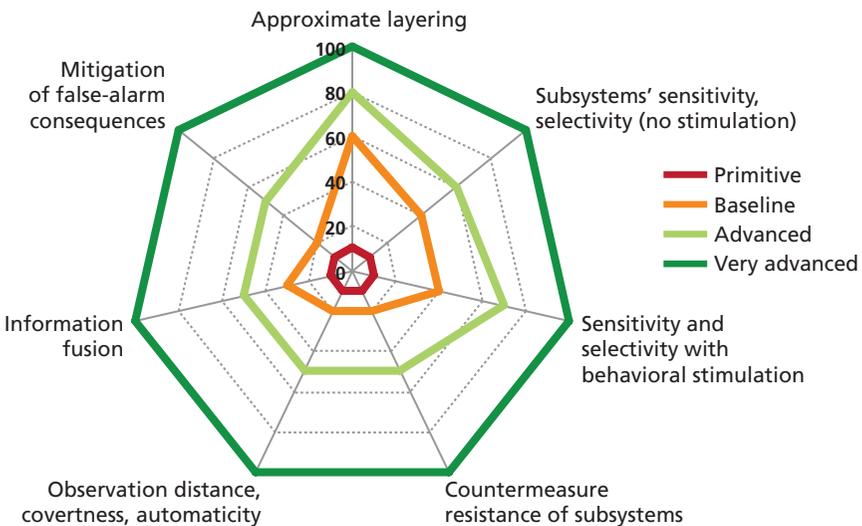
Table S.3
Detecting Hostility or Deception from Movement Physiology and Movement

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Kinetics and gross motor movement	Indicators have been validated for human observation and automated analysis in laboratory and experimental settings, including some operational settings (e.g., for gait of individuals carrying weighted objects).	Gross motor movements may reveal action, intent, or deception. On-foot motions may be unavoidable in such proximal security settings as checkpoints. Gross motor movement may be passively observed, but also actively elicited.	Some security contexts may not allow for sufficient physical movement to be interpretable (e.g., interrogation).	Masking with deceptive movements. Sensitivity to context and individual differences. Nonspecificity: triggering by diverse emotions and motivations.
Physiological state and reactions	Indicators have been validated in laboratory and experimental settings, with some experimental paradigms simulating elements of counterterrorism and some (facial) cutting across culture. In some cases (e.g., voice stress and facial indicators), automated recognition shows potential but currently has high error rates.	Internal physiological reactions are relatively automatic and difficult to control (e.g., micro tremors in speech or micro facial expressions). Probing of various sorts (even seemingly random conversations) can trigger reactions. Certain elements of facial expression are very difficult to alter voluntarily, including micro-expressions.	Currently, measurement requires direct application of sensors or the physical observation of, e.g., facial flushing, sweating. Some (e.g., facial) require lighting and proximity with currently painstaking coding feasible only for high-value interrogations. Success requires exceptional "natural" talent or training, but limited available data suggests training is effective. Measurements are most valuable when comparing against an individual's baseline, which is only feasible in voluntary monitoring or interrogation context.	Differences across contexts and individuals. Nonspecificity. Influence of drugs and training (e.g., to dampen or obscure differences between baseline and signals). Masking, in some cases (e.g., sunglasses or plastic surgery) Some differences exist (perhaps not critical) across culture. Masking (e.g., sunglasses, plastic surgery, or Botox for facial), but this may also be an indicator.

have found the framework qualitatively useful for discussing the issues that arose in our critical survey.

Figure S.5 uses a radar/spider plot to characterize a given detection system along seven dimensions, with a score of 100 corresponding to a system that has been optimized along that dimension while considering feasibility and cost-effectiveness. The score given to a lesser system is a rough and subjective characterization of how much has been accomplished relative to what would be accomplished optimally. The dimensions relate to (1) appropriate layering, (2) sensitivity and selectivity of subsystems for information when it is obtained and countermeasures are absent, (3) behavioral stimulation, (4) countermeasure resistance of those subsystems, (5) the ability to obtain the information in desirable ways that may include automated observations from a distance, perhaps without subjects being aware of the observations, (6) information fusion, and (7) mitigating the consequences of false alarms when they occur.

These dimensions relate to overall system effectiveness as shown in Figure S.6, which highlights the need for both effective detection (minimizing "false negatives") and management of the false-alarm problem (minimizing "false positives" or mitigating their negative consequences). Returning to Figure S.5, we see that it illustrates notionally

Figure S.6
Factors Affecting Overall System Effectiveness
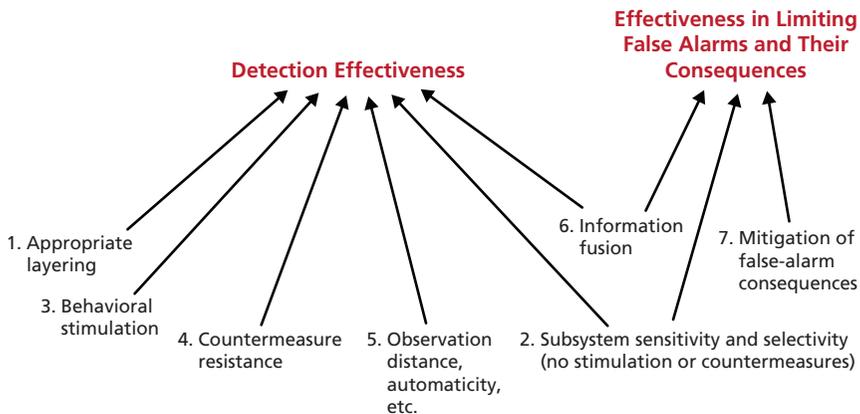
what progress might look like over time. For the example, we characterize the baseline (today's system) as more advanced in some respects than others, with layering having been taken seriously for some time, but with information fusion, for example, being relatively primitive and with too little work having been done to mitigate the consequences of false alarms when they occur. Progress would correspond to systems with contours farther and farther toward the extremity of the radar/spider plot (cautions in interpreting such plots are discussed in the main text). Subsequent paragraphs touch on each dimension.

Appropriate Layering

The value of layering in detection systems is well discussed elsewhere and merits little discussion other than to note that the fatal flaw in some assessments is assuming that the various layers are independent, which is not the case when, for example, they share lax or incompetent management. How much layering is appropriate depends on many factors.

Subsystem Sensitivity and Selectivity

Screening can be based on many types of information, such as background checks, overtly observable characteristics (including the carrying of weapons), or behavioral cues. Ongoing research is a mix of laboratory- and field-based empirical research and modeling. As an example, the Department of Homeland Security's (DHS's) Screening of Passengers by Observation Techniques (SPOT) Program was designed to help behavioral detection officers (BDOs) identify persons who may pose a potential security risk at Transportation Security Administration (TSA)–regulated airports. It focuses on behaviors and appearances that deviate from an established baseline and that may be indicative (imperfectly, to be sure) of stress, fear, or deception. BDOs may refer some of those whom they observe to additional screening. DHS conducted a large randomized trial of SPOT effectiveness and reported that the program resulted in the detection of illegal behaviors at a significantly higher rate than random selections, with a small false-alarm rate. Another empirical effort, Project Hostile Intent (later renamed Future Attribute Screening Technology [FAST]), included consider-

ation of hard-to-suppress micro-expressions as it sought remote, non-invasive, automated sensor technology capable of real-time detection of deception.

Although not behaviorally oriented, a second class of screening method should be mentioned. It is illustrated by the "Trusted Traveler" concept, which screens for those who can be *excluded* from some secondary screening. A 2011 analytical review demonstrated considerable promise, especially if measures are taken to deter attacker efforts to get into the program; aspects of the concept are now operational.

A constant issue in screening is how to trade off detection rate against false-alarm rate. Doing so should depend on context. During a period of high alert, security personnel can use less discriminate behavioral (and other) cues if they have additional temporary resources. During such a period, the public also tends to be more forgiving of inconvenience and somewhat higher false-alarm rates are tolerable.

As mentioned earlier, measuring physiological responses is most effective when measuring actual brain activity rather than downstream effects such as flushing. For example, electroencephalogram (EEG) measurements have shown high effectiveness in laboratory experiments to detect deception by subjects in mock terrorist attacks. Such methods, however, require the cooperation or coercion of individuals and expensive monitoring equipment as well as credible prior intelligence about the details of a potential attack. The approach appears to be countermeasure-resistant, but particularly aggressive individuals show less of the response being monitored, which would reduce real-world effectiveness. Nonetheless, the successes are a remarkable advance. Numerous empirical studies have also been performed on "downstream" responses, including use of polygraph methods and remote detection of peripheral physiological signals. These and other methods have shortcomings, such as nonspecificity and sensitivity to the "base rate" (the fraction of observed individuals who present an actual threat). If the base rate is very low, then false alarms are high.

Behavioral Stimulation

Probing to stimulate behavioral responses can sometimes improve detection effectiveness significantly. The basic concept has long been

familiar to law enforcement and many tangible examples exist, partly as the result of U.S. security officials learning from extensive Israeli practice. More generally, probing refers to the intentional stimulating of behavioral responses, such as by verbal questioning, anxiety-raising changes of procedure or process, subliminal stimuli, or tests with polygraph or EEG equipment. Probing may be polite, unobtrusive, or "in-your-face." Some probing can definitely improve detection-system results, but related experimentation and formal assessment has not been pursued as far as it might be. Verbal provocation and human assessment of verbal and behavioral responses can be effective in some circumstances without the use of sophisticated or expensive biological monitoring equipment. Israeli airport and other transit officials have used such techniques for many years, apparently with success (some of it deterrence-related). Subjective assessment of the plausibility of reasons given for traveling, or being at a certain location, along with the consistency of stories over time together provide the best clues about hostile or deceptive intent. Further research should address contextually distinct tradeoffs between benefits for detection effectiveness and negative consequences for civil liberties, commerce, and the perceived legitimacy of the security system.

Allowing for and Dealing with Countermeasures
Much of the literature and even more of the advocacy-related discussion focuses on detecting behavioral responses in the absence of countermeasures, but countermeasures are in fact a big problem, and vulnerability to countermeasures should be a prime consideration in evaluating investment programs. That said, countermeasures often are not employed, are attempted poorly, or themselves create indicators. Thus, a balance must be struck in analysis: Worst-casing could eliminate valuable measures, but optimistic assumptions could waste resources and divert attention from more promising methods. Unfortunately, net judgments are often made informally and ad hoc. Analysis could improve this situation.

Observation Circumstances: Remoteness, Covertness, Automaticity

Many of the potentially attractive technologies and methods currently depend on such relatively benign circumstances as close-up observation by humans, sometimes with a need to minimize "noise" relevant to detection systems. Operational value, however, will be much enhanced by improved capabilities to make observations from a distance, automatically, and in some instances without the subjects being aware of the observation. Progress is being made by active technology efforts on all of these. Some of the efforts are benefiting from commercial and law-enforcement-system investments in, e.g., ubiquitous security-video recordings; supervised and unsupervised computer search of data, including "big data"; and new analysis techniques, such as those used in data mining.

The Potentially Critical Role of Information Fusion

We found nothing on the horizon that presented a "magic bullet" for threat detection, raising the potential importance of effective information fusion. We reviewed quite a number of methods for combining information, ranging from very simple to more sophisticated methods. Notably, some classes of fusion have long been used. Indeed, polygraph testing combines information from several types of physiological signal. However, we have in mind information fusion that also combines information across activities and phases. We considered a number of possibilities.

Heuristic and Simple-Model Methods include checklists and risk indexes, which are especially suitable for on-the-scene security personnel. Checklists are common already and can be of two kinds, which are sometimes referred to as positive and negative (but with different authors reversing which is which). As examples, any indicator, if met, might trigger additional screening; alternatively, if all indicators are met, secondary screening might be minimized. Index methods (scoring methods) typically characterize a risk level by summing indicator scores, or by computing a risk as the product of a likelihood and a consequence, with a score exceeding a threshold triggering additional screening. Significantly, good scoring methods often need to be nonlinear and should be empirically validated rather than ad hoc. We

also consider more complex "simple methods," such as scorecards and conditional-indicator sets.

More sophisticated integration methods are likely necessary in future information-fusion centers, which would try to incorporate behavioral indicators to overcome serious signal-to-noise and false-alarm problems. Accordingly, we reviewed mathematical information-fusion methods that might be adapted and extended (these methods are discussed in more detail in Appendix D). Bayesian updating is well understood and widely applied in other domains, but its usefulness in our context is limited by its demands for many subjective estimates of conditional probabilities for which there are and will continue to be an inadequate base, and by limitations of expressiveness. Some newer methods are based on Dempster-Shaefer belief-functions, which distinguish between having evidence *for* a proposition (such as the malign intent of someone observed) and having contrary evidence (i.e., of innocence). Evidence for both can be high, whereas if the language used were that of simple probabilities, a high probability of malign intent would imply a low probability of innocence. Dempster-Shaefer theory requires fewer subjective inputs. Ultimately, however, there are several major shortcomings in using that approach as well.

A much newer approach, called Dezert-Smarandache (DSmT) theory, has not yet been widely discussed and applied, but something along its lines has promise because it deals specifically with combining evidence from sources and sensors that produce imprecise, fuzzy, paradoxical and highly conflicting reports—precisely the type of reports encountered. For example, it allows characterizing the evidence that both A and B are true; that one or the other of A or B is true (but not both); or the evidence that A is true and the evidence that A is not true. We also reviewed, briefly, the relevance of "possibility theory," various multi-attribute theories, "mutual information" (which builds on the concept of information entropy), and Kalman filtering. The best method(s) for this problem area are not yet certain, but our review may help to generate fruitful research in this critical area.

Mitigating Costs of False Alarms

As mentioned repeatedly, a major challenge in detection systems is the tradeoff between false negatives (failure to detect) and false positives (false alarms), known as Type I and Type II errors. An understudied problem amenable to research is how the broadly construed cost of the latter can be reduced—not just by reducing the false-alarm rate, but also by mitigating such bad consequences of false alarms as wasting people's time, raising their fears, insulting their dignity, or invading their privacy. We identify three classes of initiative: (1) improve system effectiveness (a "no-brainer"); (2) reduce effects on dignity and perceived violations of civil liberties (e.g., by transparency, explanation, fairness, apology, and compensation); and (3) deter abuse by those within the security system. Progress on the latter two is highly desirable for broad societal reasons and has many precedents in law enforcement. The negative consequences of false alarms alienate people, who are then less likely to cooperate, volunteer suspicions, and support the security system.

A Core Issue in the Use of Behavioral Indicators

Many of the subjects reviewed in our study are extremely contentious. Some of the controversy is scientific, relating to whether various detection methods are scientifically sound (or, as some would have it, pseudo-science). The issue is not straightforward, because detecting attacks by subjects such as terrorists involves looking for weak signals amidst a great deal of noise in circumstances in which the "base rate" is extremely low. The consequences of detection failure are very high, but there are also profound negative consequences related to false alarms, as mentioned above.

We could not resolve the controversies in this study, but Table S.4 makes distinctions useful in discussion. It compares how various methods that use behavioral indicators can be used. All of them have deterrent or cost-imposition value (second column). Would-be attackers often fear the technology and methods and behave accordingly. All of the methods can, when properly used and in proper circumstance, be

Table S.4
Some Comparisons of Where Behavioral Methods Have Value

Method	Deterrence or Cost Imposition	Flagging for Further Routine Screening		Flagging with Prejudice for Extended Checking and Detention	Tool in Interrogation	Basis for Arrest or Conviction
		Automatic	Human			
Polygraph	Yes	No	Yes	Maybe	Yes, but	No
Voice stress analysis	Yes	Yes	Yes	No	Yes, but	No
Facial expression	Yes	Technology not well developed	Yes	No	Yes, but	No
EEG	Yes	Technology not developed	Yes	Maybe	Yes, but	No
Text or speech content	Yes	Maybe	Yes	Maybe	Yes, but	Maybe
Gait analysis	Yes	Yes	Yes	Maybe	No	No

useful in providing incremental evidence on which subjects merit closer scrutiny (third and fourth columns), although there are big variations in whether they can be used automatically, remotely, and covertly. All of the methods, if well used, can *sometimes* (fifth column) justify treating an individual with considerable concern, with subsequent assessment done "with prejudice" in the sense of being potentially extended and including detention and aggressive questioning. That "sometimes" should be understood as "occasionally," however, and the methods typically have high false-alarm rates. The sixth columns uses "Yes, but . . ." to indicate that yes, if a subject merits in-depth interrogation, most of the methods can—as part of a more complex process with skilled security officers—be useful in obtaining confessions or information, but, regrettably, they can also help generate false confessions. Abuse can occur. The last column is crucial: None of the methods, except possibly for analysis of textual or vocal content, are individually an adequate

basis for arrest or conviction. Indeed, they may not be an adequate basis for putting prejudicial information in a widely shared database (e.g., "On such-and-such an occasion, the subject manifested facial-expression behaviors correlated with posing a security risk, although other factors led to his being allowed to board the aircraft").

This illustrates one of the many unresolved dilemmas. From a purely detection perspective, and assuming a process for information fusion, it would seem desirable to collect and share all kinds of fragmentary information of varied significance and credibility. However, doing so could cause serious injustices to those affected and, in many instances, would generate suspicions when none are scientifically warranted. It is instructive that, for almost a century, the FBI has maintained "raw files" on numerous subjects of observation, with important instances of those files being misused (even though it can be argued that this occurred rarely). How much more trouble would have been created if analogous raw data had been widely shared? Such issues are matters of degree, but no common agreement exists on what is and is not reasonable. As a last example motivated by current discussions in the news (as of January 2013), consider a teenager being treated for symptoms of schizophrenia. What symptoms of violent tendencies should trigger a report to authorities that would enter a sharable database, and with what balance of positive and negative consequences? Such issues are profound. We made no attempt to resolve them except that we see a major distinction between, on the one hand, using a behavioral indicator as an increment of information in a detection system seeking to identify, without further prejudice, which individuals merit more-than-usual scrutiny, and, on the other hand, using a behavioral indicator to infer probable guilt or as the basis for arrest and conviction. It is not accidental that the U.S. justice system has major constraints on how methods such as polygraph techniques can be used.

Conclusions

We found a number of important takeaways from our survey:

- Despite exaggerations found in commercial claims and the media, there is current value and unrealized potential for using behavioral indicators as *part* of a system to detect attacks. Unfortunately, analytic quantification of that potential is poorly developed.
- "Operators" are often well ahead of the science base, which is sometimes good and sometimes bad. It is very important that programs build in and sustain objective evaluation efforts, despite budgetary pressures and the tendency to see them as mere nice-to-have items. The evaluations should be subjected to objective peer review and adequate community scrutiny, although perhaps within a classified domain. The Department of Defense and the Intelligence Community have, for example, long used the federally funded research and development centers (FFRDCs), national laboratories, National Academy of Sciences, and other special panels for credible evaluations.
- Many serious problems and errors can be avoided by up-front review of procedures by experts familiar with the subtleties of detection and screening in conditions of high false-alarm rates and low base rates. Although full validation of techniques may take years (at a time when the dangers of attack are current), existing knowledge can be used to avoide many problems that are quite significant to privacy, civil liberties, travel and commerce.
- DHS and other security organizations are experimenting with proposed methods—sometimes with laudable and ambitious scientific trials that have reported encouraging conclusions (which are difficult to judge, however, without detailed access to data and methods).
- Operators, their agencies, and the scientific community have not done enough to understand how to mitigate the bad consequences of detection systems, which invariably have false-alarm problems. Much could be done.

- Information fusion is critical if behavioral indicators are to achieve their potential. Fusion should occur not just within a given method, but with heterogeneous information across activities and phases. Methods for accomplishing this are very poorly developed. This said, it remains to be seen how much can realistically be accomplished. If the indicators being fused all have very high false-alarm rates, the fused result may be more reliable but still have a high false-alarm rate. Also, success in fusion will depend on human skill in representing fuzzy, imperfect information.
- Information generation and retrieval, integration, and sense-making will tax both automated methods (e.g., including for "big data") and perfecting human-machine interactions: Machines can process vast amounts of data, but interpretation will continue to depend on human expertise and judgment. An implications is that "optimizing" should be for man-machine cooperation, not automation.
- Very little research has been done to understand how much is enough, but, subjectively, it seems that major improvements in detection are plausible with networked real-time or near-real-time integration of information. This would include further integrating (fusing) CIA and FBI information; proximate information at checkpoints and fusion-center information; and criminal, commercial, security-related, and even whole-life information. What can be accomplished is unclear, and developing a sharper understanding of payoff potential should be a priority task for objective research and analysis.
- Such steps raise profound issues of privacy and civil liberties, but the irony is that commercial organizations (and even political parties) are already far ahead in exploiting the relevant technologies and forever changing notions of privacy.
- Investment decisions about individual technologies and methods should be informed by structured portfolio-analysis approach using something like the dimensions of Figure S.6.

Acknowledgments

Abbreviations

CCTV	closed-circuit television
CIA	Central Intelligence Agency
CONOP	concept of operation
DARPA	Defense Advanced Research Projects Agency
DHS	Department of Homeland Security
DoD	Department of Defense
DSmT	Dezert-Smarandache Theory
EEG	electroencephalogram
FAST	Future Attribute Screening Technology
GAO	Government Accountability Office
IED	improvised explosive device
R&D	research and development
SPOT	Screening of Passengers by Observation Techniques
TSA	Transportation Security Administration

Introduction

Background

Federal, state, and local government organizations have put substantial effort into detecting and thwarting terrorist and insurgent attacks by observing suspicious behaviors of individuals, whether at transportation checkpoints. Technologies and methodologies abound for contributing to such defensive activities in myriad ways. However, the volume and diversity of activities and claims has often been overwhelming. Further, claims about effectiveness sometimes lack a clear basis in science and technology. This occurs for different reasons. Sometimes operators in the field move quickly to deal with clear and present dangers without having the benefit of scientific groundwork. Other times, enthusiasts for a clever idea or new technology exaggerate its potential, perhaps by not accounting for diverse operational circumstances or for adversary adaptations.

The RAND Corporation was asked to improve the situation by conducting an analytically useful literature review of the base in behavioral sciences relevant to threat detection and by identifying tentative priorities for special attention and investment.[*]

---

[*] Another recent study relating to the prediction of violent behavior (Defense Science Board, 2012) raises and discusses policy issues that we do not discuss in this report.

Scope

Types of Attack, Attacker, and Behaviors

Deciding on the scope of our research was itself a challenge: Were we focused on countering terrorism by suicide bombers? Were we concerned with insurgent violence that is not really terrorism (i.e., attacks on noncombatants)? Were we considering crime as well? Was the scope to be defined by target, attacker's intent, or what? After considerable discussion, we concluded that the scope of our work would be as follows:

1. Detecting potential attacks by *individuals or small groups* (not large military or irregular formations), whether or not guided or supported by a larger organization (which might be a terrorist, military, irregular, or even criminal organization).
2. *Diverse contexts and targets*, such in a foreign theater of military operations or in domestic locations, such as sports stadiums or political speeches.
3. Looking primarily for *nontraditional* (that is, based on recent behavioral science rather than police or intelligence "business as usual") behavioral observations and analysis to identify individuals for increased scrutiny because they appear to be more likely than others (even if still quite unlikely) to have hostile intent or be otherwise supporting an attack. The behaviors need not, and typically would not, directly relate to hostile intent.
4. Our research, then, was on levels of analysis where individual behaviors matter. It would address technology and methods that might help detect such diverse attacks as domestic suicide bombing by a lone wolf or the laying of improvised explosive devices (IEDs) in a military theater. It would focus on going beyond traditional observation of behaviors as long practiced by sentries, guards, police, intelligence officers, and internal security personnel.

To further explain our rationale, we would *not* focus strictly on suicide bombers, because we did not want to exclude attackers who

had an escape plan. We did not want to focus exclusively on "terrorism" (attack on noncombatants), because many attacks of concern are not actually terrorism but rather "normal" violent acts committed as part of insurgency or war (insurgents may or may not use terrorism). We would not narrow down to particular motivations, because the attacks of concern, whether by terrorists or insurgents, have myriad motivations.

The third item above calls for focusing on "nontraditional" methods. What might be nontraditional? We had in mind (1) exploiting information processing and sharing, (2) looking for patterns or other indicators that have not been systematically exploited in the past, and (3) substantially improving the ability to do traditional forms of observation or enabling new types of observation (e.g., with long-distance and/or automated detection of individuals manifesting stress or attempting to avoid observation). Table 1.1 lists many traditional methods alongside the "new" methods or technologies we are interested in. Although this report certainly touches on some of the items in the middle column (traditional methods), more weight is given to those in the right column. As an example, we do not dwell on the extensive literature about polygraph tests. Instead, we give more space to newer methods for detecting deception, such as observing unintentional linguistic patterns and facial expressions characteristic of deception.

Although we did not intend to ignore any particular classes of attack by individuals or small groups, we had in mind primarily a number of relatively special attack classes, such as laying IEDs in marketplaces or against convoys; political or military assassinations; airline or other transportation attacks; a "new September 11"; post-battle killing of innocents; and genocidal raids or attacks. Some of these would have combatant targets; some would have noncombatant targets. Such examples cover a considerable range but do *not* include military attacks generally, domestic crime generally, or many kinds of violent action.

Given the time and resources available, we were unable to summarize results separately for different attack types or different operational environments, even though indicators and methods would vary across them.

Table 1.1
Examples of Traditional and Newer Technologies and Methods for Detecting Possible Violent Intent

Observations	Traditional Methods	Newer Methods
Observe physical movement	Watch for people seeking to avoid checkpoints or exploit crowds, or showing strange body language.	Increase automation and at-a-distance methods in watching for traditional indicators. Detect styles of body movement subtly correlated with suspicious behavior.
Observe personal demeanor and behavior	Detect signs of stress or deception.	Detect efforts to suppress traditional signs of stress and deception. Detect subtle and less controllable signs (e.g., facial expressions) correlated with suspicious behavior.
Observe responses to questions	Detect deception or fear of questioning.	Detect subtle efforts to deceive, hide, or suppress information.
Learn from documents and personal information	Infer information from passport, name, garb, responses to questions, and use of modest network information.	Retrieve extensive personal information, history, and relationships in real time, based on identifications from passport, biometrics, and observations.
Allocate defensive resources	Allocate guards and observers according to visual information and hunches. Surge or withhold resources based on a priori concerns and hunches about, e.g., orchestrated, sequential attacks and traps.	Supplement allocation with automated methods using sub-threshold indicators (e.g., identify multiple people for follow-up checks, rather than focusing on the first suspicious target). Supplement decisions with aids to mitigate risks and hedge.
Orchestrate, coordinate	Personally contact other defense-system individuals ad hoc, perhaps across organizations. Enter information into record. Issue community bulletins.	Increase sophistication of automated real-time push/pull methods across agencies (network-centric operations). Simplify data entry (e.g., automatically digitized verbal information). Use automated "data mining" to fuse indicators across all-source data.

Treatment of Privacy and Civil Liberty Issues

An important consideration in scoping the project was deciding on filters. In particular, would we limit discussion to methods and technologies that seem acceptable to society and are consistent with current interpretations of current laws? We concluded that such filters were not appropriate: The report's intent was more to note scientific/technical possibilities than to assess tradeoffs with civil liberties or other considerations affecting potential acceptability, now or in the future. As a result, some of the methods discussed would and should be quite controversial, even if feasible. Examples include forms of profiling; creating, exploiting, and mining databases of personal information; and sharing information across agencies in ways that would increase opportunities for misuse.

This limitation of scope was troubling, but our concerns were alleviated by the fact that excellent work has been done specifically on the many privacy issues and how to reconcile them with counterterrorism efforts. In particular, we refer readers to a National Academy of Sciences study chaired by William J. Perry and Charles M. Vest (Perry and Vest, 2008) that included panelists from law, law enforcement, information technology, computer science, and other fields. In addition, we decided to include comment from time to time throughout the study on privacy and civil liberty issues, and on where opportunities exist to mitigate related problems. Thus, while we do not analyze tradeoffs between detection and civil liberties, we make related observations where doing so appeared useful.

Structure of Approach

Given the broad scope of the review, we constructed a conceptual model to provide structure. The model needed to be comprehensive enough to ensure that we would consider an appropriate range of possible methods and literatures.

As indicated in Figure 1.1, a nominal context for the report is of security personnel watching for individuals or small groups who may pose a threat near, in position or time, to a potential target or event.

Figure 1.1
A Contextual View of the Detection Effort
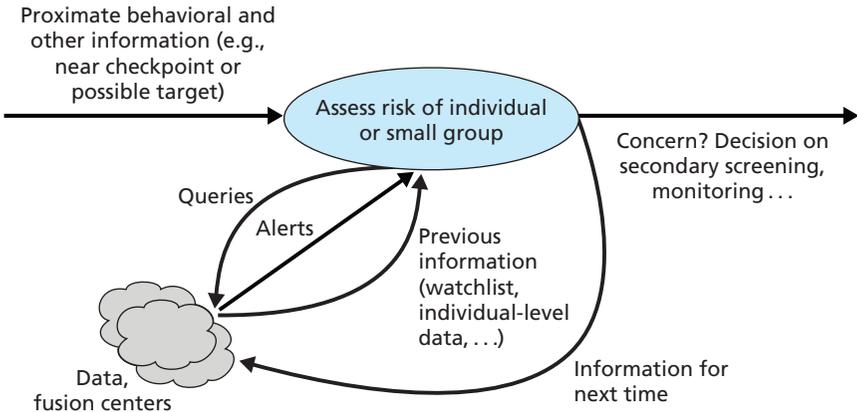
The proximate information might come from, e.g., video, audio, or security-personnel communication. The security personnel may, for example, be scanning a crowd, surveilling an area, checking identity papers at checkpoints, or reviewing relevant intelligence. Something is observed, which triggers further observations and checks. The checks include drawing on databases (or consulting with other humans) to see whether the subject in question raises concerns. For example, the subject might already be on a watch list or have troubling records of one type or another.

Figure 1.1 should be interpreted broadly. For example, the proximate information might be based on remote imaging with smart cameras. Also, "other information" may include such recently updated projections/predictions as maps of "hot spots" where attacks are most likely. It might even include information being collected at the same time as behavioral data (e.g., identification and travel history).

Figure 1.2 shows how we relate indicators, activities, and technology and methods. A base of technology and methods (left) is able to detect behavioral indicators of activities, which may then be clumped into aggregate activity classes that we call "phases" for short. Analysis

Figure 1.2
Relationships Among Constructs
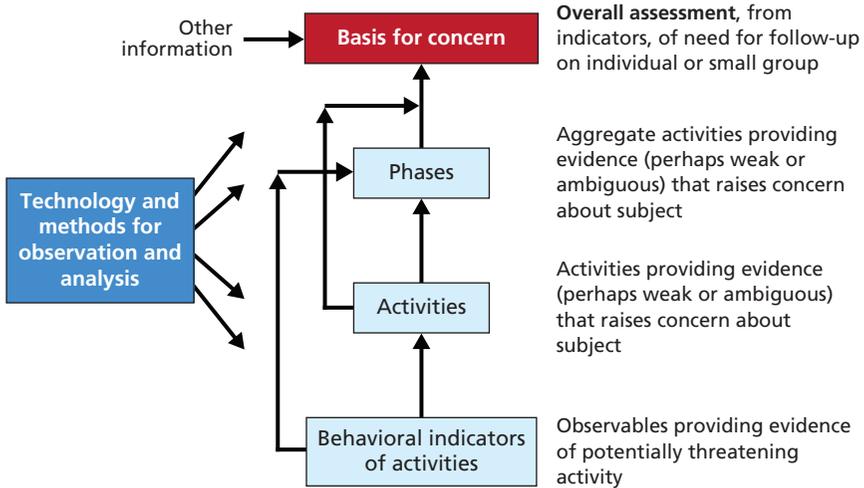
RAND *RR215-1.2*

can then assess whether the totality of what has been observed adds up to a basis for concern—i.e., justification for more in-depth observation (e.g., probing), precautionary defensive measures, continued monitoring, or intercept. A high "basis for concern" seldom implies a high absolute probability that the individual or small group is likely to attempt a violent attack. Indeed, this is a core reality: The "base rate" of target individuals is exceedingly small. A system using this approach, then, must be efficient and reasonable if it is to be acceptable.

Figure 1.3 elaborates. The meaning of the phases (top level) is fairly self-evident. The leftmost, however, is different in kind, relating to a phase in which individuals or small groups are doing things, such as studying extremist ideas or being recruited for or seeking out a terrorist organization, that may lead over time to intending hostile action. The subsequent phases in our model are conceived relative to an actual attack: the planning-and-laying groundwork phase, the immediate pre-execution phase, the execution phase, and the aftermath phase. The names of these phases may relate either to an individual (lone-wolf) attacker or an organization in which individuals are participants.

Figure 1.3
Conceptual Model of Opportunities for Observing Worrisome Behaviors
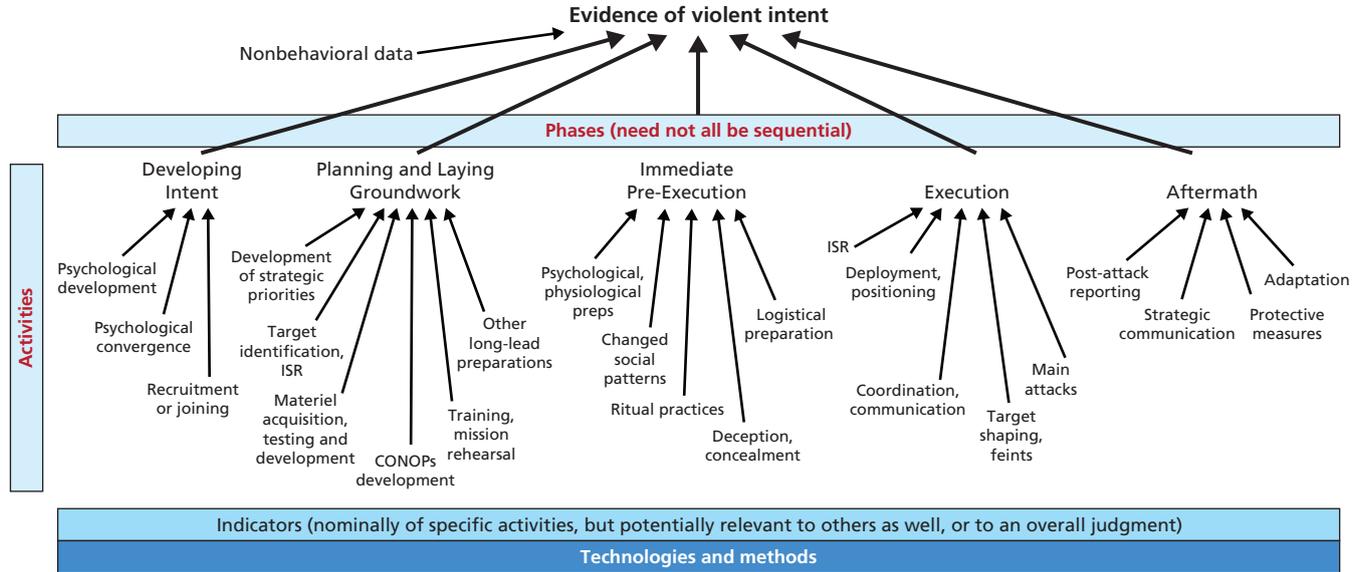
NOTES: ISR = intelligence, surveillance, and reconnaissance; CONOPs = concept of operations. Indicator-phase connections are nominal. Some activities can occur in more than the phase indicated.

RAND *RR215-1.3*

The phased model is purely for the purpose of having convenient but nominal "containers" for activities. It should not be taken too literally. In particular, a given phase may never occur; two phases may overlap in time; the order of phases may vary; and events in the execution or post-execution phase of one attack may affect motivation, planning, etc., for a next cycle of attacks. As examples, a particular individual of interest might not be involved in some phases (e.g., planning and laying groundwork, or the aftermath). Or he might have been peripherally involved in one attack and, as a result, become more interested in the organization, its activities, and its ideas. That is, the first phase-level activity might come later for some individuals, even after they have engaged in acts of violence rather than before. Again, then, the phases are simply nominal "containers" of activities without prejudice as to whether, or in what order, the individual or small group participates in the various phases. A final subtlety is that worrisome behaviors of different individuals might be observed over time, with the *accumulated* information providing evidence of an attack that the particular individuals being observed know relatively little about. Our intent, then, was to be comprehensive in thinking about "places" to look for possible indicators.*

A few other observations are worthwhile regarding Figure 1.3. This is not a decomposition diagram, as with an organization chart or system engineer's breakdown into exhaustive and independent components. Instead, it has the form of an approximate "factor tree"† showing

---

* Roughly analogous methods have been used in a variety of fields, such as with offender life cycles in criminology, where different phases are identified with respect to crime itself (one phase might be "search in a pre-criminal situation") and with respect to periods in a criminal's life, including a period of giving up crime. See, for example, the introductory chapter of Cornish and Clarke (1986). Process-model methods have been used to systematize counterproliferation research, as in identifying the numerous steps necessary to develop acquire, field, and employ a weapon of mass destruction. All such methods reflect one or another type of "system thinking," some more rigorous than others.

† "Factor-tree" conceptual models were first used in earlier RAND studies and have proven quite useful in integrating and communicating heterogeneous social-science knowledge relating to terrorism, insurgency, and stabilization and reconstruction (Davis and Cragin, 2009; Davis, 2011; Davis, Larson, et al., 2012). They can be seen as static simplifications of "causal-loop diagrams" or "influence diagrams" as used in system dynamics and policy

that what we are interested in (evidence of potential violent intent) is a *function* of a number of factors (evidence of activities and sub-activities). Further, any indicator might, in principle, shed light on any activity, or even have a direct effect on the overall assessment. That is, the concept must allow for a *network* more general than a simple hierarchy.

Table 1.2 shows much the same information, but in tabular form.

The fact that the activities in Figure 1.3 are not mutually exclusive components of a rigorous decomposition has important implications for how evidence about the various factors (activities) can be com-

Table 1.2
Phases and Activities

Phase	Activities
Developing intent	Motivational development Psychological convergence Recruitment or joining
Planning and laying groundwork	Development of strategic priorities Target identification and intelligence, surveillance, and reconnaissance (ISR) Materiel acquisition, testing, and development Conceopt of operations (CONOPs) development Training and mission rehearsal Other long-lead-time preparations
Immediate pre-execution	Psychological and physiological preparations Changed social patterns Ritual practices Deception and concealment[a]
Execution	Intelligence, surveillance, and reconnaissance (ISR) Deployment and positioning Coordination and communication[a] Target shaping and feints Main attacks
Aftermath	Post-attack reporting Strategic communication Protective measures Adaptation

[a] This activity may also be associated with other phases.

analysis. The nodes (i.e., the factors or variables) need not, and typically are not, interpreted probabilistically, as are the influence diagrams in Bayesian or influence-net research.

bined, as discussed in Chapter Seven briefly and in Appendix D in more detail.

Within each of the activities, there may be potentially observable indicators. These may be exploited by either existing or future technologies and methods. Figure 1.4 illustrates our methodology for the "Developing Intent" activity of Figure 1.3. It shows the activities for this phase (as in Figure 1.3 as well, but horizontally rather than vertically). It then lists, for each activity, the behavioral indicators that we considered and the kinds of technology and methods that are or could be applicable. For such listing and discussion, we conducted a substan-

Figure 1.4
Illustration of Methodology
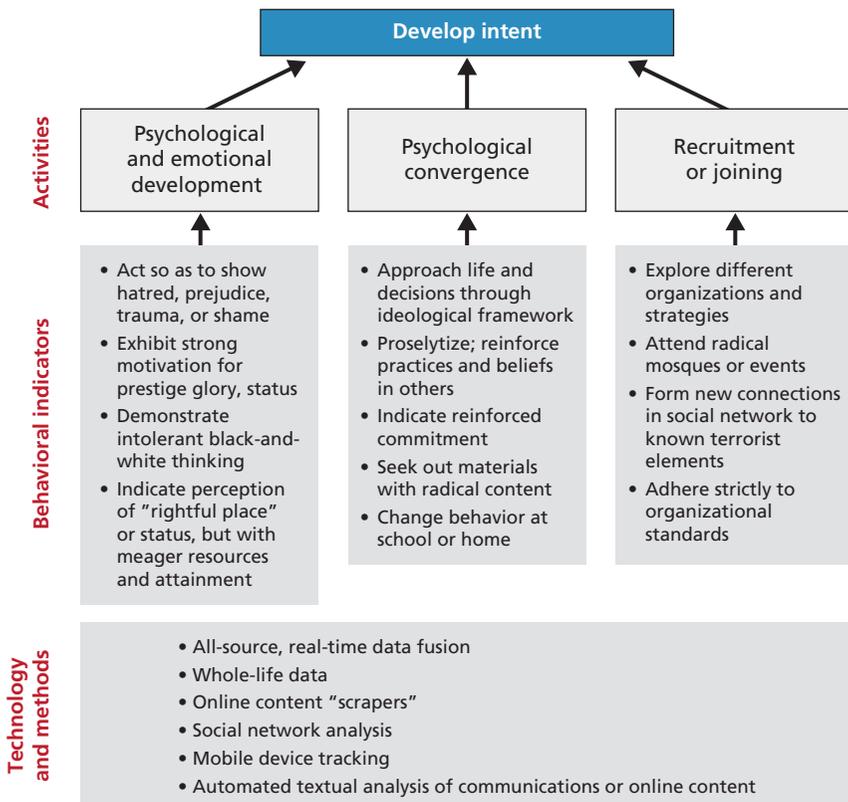
tial literature search and also interviewed prominent scientists and officials (see also Appendix A).

As noted earlier, the conceptual phase-level activities of Figure 1.3 are to some extent ambiguous and overlapping. We followed certain conventions in deciding in what phase a given activity or observable belongs. Overall, the primary issue is achieving approximate comprehensiveness, not cataloging each and every possible activity uniquely. This said, our rule-of-thumb conventions are as follows:

1. *Developing Intent.* This phase is associated more with individuals than with the organization; it is about motivation and commitment—whether to a cause, organization, or activity. An individual might be participating in organizational activities, such as meetings or even general training, that have the effect of creating motivation and commitment. If so, the activities are in the developing-intent phase. In contrast, previously motivated and committed individuals participating in the same activities might have their activities counted as part of the planning-and-laying-groundwork phase. The same observable activity might be listed in both phases.

2. *Planning and Laying Groundwork.* This phase is associated with the organization (or a lone wolf)* even though we may be observing individual behaviors. This is the phase in which the organization does its planning and prepares its people broadly for operations, perhaps with general physical training and the teaching of combat skills.

3. *Immediate Pre-Execution.* This phase, again associated with the organization's perspective, is one in which plans are finalized and resources mobilized and maneuvered so as to make subsequent execution feasible, if decided upon. It might include increased reconnaissance (or, conversely, a period of reduced visibility because adequate information has already been obtained). It might mean deploying people to the relevant country, area, or city, but putting them in holding patterns.

---

* For lone-wolf terrorists, organization and individual are the same.

4. *Execution*. This phase applies once a decision to commence the attack has been made (by the organization or, in a lone-wolf case, by the individual himself). This definition is consistent with the meaning of *execution* in military command and control. Execution may require initial activities, such as maneuvering resources to their final attack locations (perhaps moving through or around checkpoints), final reconnaissance, arming of weapons, and coordination-related communications. Almost any execution operation is contingent, in that the attack can be called off along the way. Nonetheless, until and unless the attack is called off, activities in response to an "execute" order (or decision) are regarded as in the execution phase.

5. *Aftermath.* After an attack is accomplished or an in-process attack is terminated, activities are considered to be in the aftermath phase. This might include dispersing, vacating observation posts, pulling back agents, and communications related to escape or withdrawal.

Recurring Themes

Some themes recurred throughout or research. We came to these themes based on the following conclusions: (1) Most indicators of potential interest will have low detection rates and large false-alarm rates; (2) detection rates can probably be raised only by extracting weak signals amidst a great deal of noise or by somehow causing the signals to be stronger; and (3) even where a new approach seems promising, we should anticipate adversary adaptations and countermeasures when assessing its potential. The themes, then, are as follows:*

---

* One absent theme is general data mining, such as collecting and mining behavioral data on all people in a population over time. Our focus is more on detecting attacks (e.g., at checkpoints or other defenses around targets) rather than, say, searching broadly for people with patterns of behavior that might relate somehow to terrorism, or in searching broadly for evidence that an individual is possibly subversive. Broad behavioral surveillance would raise especially profound issues of privacy, civil liberties, and the nature of pluralistic democracy (Perry and Vest, 2008).

- screening for individuals or groups meriting further monitoring and evaluation, and understanding risks and benefits of such screening, with and without probing or the stimulating of responses
- dealing with countermeasures and adaptations
- combining information (information fusion) across indicators and activities.

For each of these, we see several corollary themes that represent areas for greater focus and development:

- real-time and near-real-time networking on an extraordinary level to draw on information of disparate types and sources—both to increase detection rates and to reduce false-alarm rates
- managing the system by adjusting sensitivity of detection systems by context in recognition that in some periods maximizing detection probability is paramount, whereas more normal operations must limit the false-alarm rate because of disruptions to people and commerce, and the high costs of dealing with those false alarms
- mitigating the ill consequences of false alarms, which will assuredly occur when dealing with weak signals amidst noise.

Chapter Structure

The chapter structure for the remainder of the report is as follows. Chapters Two through Five discuss activities and possible observables for the various phases. These chapters can be skimmed if the reader is primarily interested in detection technologies and method. Chapter Six discusses such technologies and methods in more detail. Chapter Seven reviews some cross-cutting themes. Chapter Eight gives conclusions. We also include four appendixes that include our a literature review, references and cases to support historical examples and the indicator tables, and information fusion methods.

Developing Intent

"Developing Intent" includes developing a motivation, disposition, or inclination that may lead to a violent act in the context of terrorism or insurgency.* We divide this phase into three lower-level behavioral activities: (1) motivational and emotional development, (2) psychological convergence, and (3) recruitment or joining.

Motivational and Emotional Development

Cognitive and Emotional Underpinnings

Some cognitive and emotional characteristics developed relatively early, perhaps under harsh conditions and even oppression, could support later involvement in a terrorist or insurgent attack. Behavioral indicators of such developments typically provide only very weak and ambiguous signals. For example, social disaffection can lead to involvement with criminal groups, to mental illness, or to addiction—all unconnected with insurgency or terrorism. Hatred of a regime, or an idealistic drive for change, can lead to peaceful revolution. Even weak signals, however, may be useful in recognizing individuals who are more likely than others to be part of violent attacks. Examples might be thrill-seeking or antisocial behaviors, as with the shooters of the Columbine High School massacre (Kass, 2009), or an extensive history of having been bullied in school, as with Timothy McVeigh (Smith,

---

* This phase may unfold over months or years, may occur shortly before an attack, or may not even occur separately.

Damphousse, and Roberts, 2006). Even if criminal activity occurs early, however, its significance is modest. Empirically, the majority of criminals seem to be motivated by thrill-seeking or peer influences that are limited to adolescence and early adulthood (Moffitt, 2006; Moffitt, Caspi, Harrington, and Milne, 2002; Kallioniatis and Macleod, 2010). Such large longitudinal studies have shown that only a small minority of early criminals develop into hardened, pathological criminals in adulthood. Many, and perhaps most, people with such background indicators grow up to become law-abiding and sometimes exemplary adults.

Also, interpretation of early-life indicators is tricky and depends in part on distinctions. For example, "acting out" in class is apparently a much weaker indicator of future violence than is sadism to animals. No single indicator has proven effective, but aggregate indicators have shown significant correlations with future violence (Office of the Surgeon General et al., 2001):

> The bulk of the research that has been done on risk factors identifies and measures their predictive value separately, without taking into account the influence of other risk factors. More important than any individual factor, however, is the accumulation of risk factors. Risk factors usually exist in clusters, not in isolation. Children who are abused or neglected, for example, tend to be in poor families with single parents living in disadvantaged neighborhoods beset with violence, drug use, and crime. Studies of multiple risk factors have found that they have independent, additive effects—that is, the more risk factors a child is exposed to, the greater the likelihood that he or she will become violent. One study, for example, has found that a 10-year-old exposed to 6 or more risk factors is 10 times as likely to be violent by age 18 as a 10-year-old exposed to only one factor. . . .

Demographic Indicators

Empirically, such demographic variables as education level, poverty, and unemployment are, in isolation, poor predictors of future involvement in terrorist activities (Kreuger and Maleckova, 2003; Berrebi, 2009), as are many other indicators related to so-called "root causes"

(Noricks, 2009). Demographic indicators also yield many false positives. For example, most attackers are fighting-age males, but knowing that is not especially helpful. And, of course, some attackers are different from the statistically based norm, as was the female suicide bomber who killed 43 people at an aid distribution center in northwestern Pakistan on Christmas of 2010 (Associated Press, 2010). Because such indicators have not worked well, and because of deep concerns about civil liberties, there exists a sometimes-fierce public debate about using demographic information for profiling and screening in national security (Harris and Schneier, 2012). We discuss these issues further in Chapter Six.

Early Behavioral Indicators

It is only logical to imagine that, in principle, useful warning-sign indicators of personal behavior could be found in the course of a person's development by educators, police, physicians, or computers studying web-using patterns of individuals and inferring other matters.

The most studied issue is probably whether there are warning signs before mass killings by violently crazed individuals. In retrospect, it is often possible to find such warning signs. Perpetrators often have records, whether with schools, law enforcement, or physicians, that include what could have been seen (and sometimes were seen) as indicators of possible future trouble. So also, interviews with family and friends not uncommonly show that *some* of the people saw odd behaviors. One study focused on school killings found that

> Most attackers engaged in some activity, prior to the incident, that caused others concern or indicated a need for help. (Vossekuil et al., 2002, p. 34)

The study also found that most attackers had difficulty dealing with significant losses or personal failures. Many had felt bullied, persecuted, or injured. It remains unclear how valuable the information

would have been even if it had been shared.* When viewing the evidence for the context of this study, however, we see two important distinctions:

- First, an indicator that would by no means justify an arrest or enforced medical treatment might, in connection with other information, have value at a checkpoint or an intelligence center pondering a tip about terrorist action. The result might be to increase caution, double check a discrepancy, look for more information, put the person in question under surveillance, or conduct an interrogation.
- A person-specific behavioral indicator is very different from one based on attributes such as nationality, race, and gender, although some of the same issues arise (e.g., for a young male to be risk-taking or aggressive is merely to be in a very large category within which only a small fraction would be of concern).

If such early indicators were to be useful in a later assessment of whether an individual deserved greater-than-normal scrutiny in some security context, they would have to be known at the time. In imagining this as a possibility, a number of issues arise:

1. If certain indicators are observed, should they be reported and stored (e.g., in police records) or remain confidential within a school system or clinical setting?
2. If indicators are stored, should they be shared?
3. If they are shared, would the information be pushed or pulled, and, if pulled, with what authorization (e.g., a request from the

---

* It is often concluded that none of the earlier information about an individual involved responsible for a mass killing had been actionable because, after all, most angry or depressed people do not kill other people, and behavioral changes may be due to any number of reasons, possibly temporary. There is understandably great reluctance to report individuals, much less involve government authorities, unless absolutely necessary. Even reporting an incident to a teacher or school principal is not undertaken lightly by most people. Such reluctance can be especially high in authoritarian countries, where the consequences of reporting can be particularly unpredictable, severe, and irreversible.

counterterrorism center [CTC] is different from a request by a random police officer)?

4. Could any such shared data be "raw," or would it need to be filtered and packaged so as to minimize misinterpretation and misuse?

The same kinds of issues arise in debates about gun control and sexual crimes, so they are familiar. That does not make them easier to deal with.

Finding Additional Indicators

With respect to where more useful indicators might be found with additional research, we found three topics of potential interest.

Psychological Autopsies. Psychological autopsies have been conducted for years after domestic and military incidents of suicide, whether or not related to terrorism. They can turn up remarkable information about personality; interpersonal relationships; prior indications, such as recent changes of behavior; and probable intentions. It is possible that more such work in connection with terrorist events would yield new insights.[*]

Unfortunately, such work has been largely unsystematic with large variations in the background and qualifications of those conducting the autopsies, no common doctrine for doing so, and few if any controls (Pouilot and De Leo, 2006). Thus, generalizing conclusions is difficult. In our reading of this literature, an even bigger problem—for our context rather than, say, reducing the suicide rate of military personnel or teenagers—is that the focus has largely been on finding evidence of mental illness. Substantial evidence indicates that terrorists, including terrorist suicide bombers, are not unusually afflicted with such illnesses. Rather, they often have deep beliefs and values but operate in a cultural environment in which suicide attacks are seen as worthy and heroic. Parts of Palestinian society, for example, idolize suicide bombers (Merari, 2010). The subculture of al Qaeda does so as

---

[*] We were pointed to this subject in discussions with Dana La Fon, who has much relevant experience. A number of references describe psychological autopsies (La Fon, 2008; Department of the Army, 1988; Cavanagh et al., 2003).

well, as illustrated by the attackers of 9/11 and others subsequently. We conclude the following:

- More intensive work in psychological profiling might pay off if broadened to include more attention to the cultural and environment factors affecting the attackers, as illustrated by ongoing work on how terrorist groups manipulate cultural narratives to recruit suicide bombers (Hafez, 2007).

Additionally, profiling should be undertaken with the recognition that relatively weak and unreliable signals can sometimes be useful, if judiciously used in a context of information fusion, as discussed in Chapter Seven and Appendix D.

More Discriminating Profiles. We believe, as do a number of scientists and practitioners, that the potential for "profiling" potential terrorists has been underestimated because prior efforts have been relatively crude (e.g., those focused on demographics and mental illness) and because of a desire to find strong correlations rather than informative but weak evidence.*

Speculatively, we would not be surprised if, for example, a mindset of intolerance and purely black-and-white thinking had some warning-sign value. Such thinking, of course, would also correlate with many other classes of people, including visionaries and idealists—not just the kind of people referred to derisively as intolerant "true believers." Nonetheless, information about such a personal style might have value (as indeed it does in our everyday lives as we judge whom to trust and with whom we are able to work). Some published research, although difficult to interpret, is suggestive about the need to "look harder" for profile-relevant information. As a puzzling example, sociologists Diego Gambetta and Steffen Hertog found that a much larger-than-expected number of known terrorists had backgrounds in engineering (Gambetta and Hertog, 2009). The authors speculate about reasons

---

* The term "profile" has multiple meanings, which causes a good deal of trouble. Also, the *use* of various profiles varies from benign and responsible to injurious and illegal. Some of the issues will be discussed more fully later in the report as we discuss challenges in detection theory and information fusion.

and attempt to test their hypotheses. They argue that the most plausible explanation is the result of the confluence of mindset and professional frustration. Near the end of their article, they observe (p. 227):

> The only other case in which we find a trace of engineer preeminence outside of Islamic violent groups is, consistent with the mindset hypothesis, among the most extreme right-wing movements, especially in the United States and Germany.

Additional empirical evidence comes from survey researcher Tom Rieger for the Gallup Organization, which used the Political Radicals model (POLRAD) to identify two different classes of what the report called "radicals," one of which was seen to be intolerant, elitist, ideological, distrustful of government, and able to thrive in areas where safety concerns are especially strong (Rieger, 2008).[*]

However these issues turn out after further study, these strands of research provide evidence that cognitive style and personal behavior may fare better than crude demographics in a search for indicators based on empirical correlations.

Much personal information, of course, is usually off limits for networked crosschecks, but a huge spectrum exists between having records *never* available (until after a tragedy) and having them routinely available to authorities. Investigations by police or the FBI frequently draw on personal information. Indeed, it is common for television shows and movies to assert remarkable capability today, as in the popular *NCIS* (Naval Criminal Investigative Service). And, of course, some intelligence agencies today depend heavily on networked databases for diverse purposes.

Behavioral Targeting. A domain with unquestionable potential (but many pitfalls) is the mining of web data for personal information, patterns, and plausible inferences. The terms used are "behavioral targeting" and "predictive behavioral targeting." A small public literature exists on the subject, although such matters were largely outside the scope of our study. A well-researched and thoughtful review, from an

---

[*]   We benefited also from personal discussions with Thomas Rieger about survey-research conclusions from Iraq and Afghanistan, including discussions of intolerance.

academic press although intended for a public audience, is *The Daily You* (Turow, 2011). Whether or not desirable (see Perry and Vest, 2008, for in-depth discussion), much pattern development is already being accomplished. Related technologies are discussed further in Chapter Six. The primary conclusion here is that:

- *Technically*, an enormous amount of personal-level information could be tapped in near real time with massive networking. Some of this is happening already for commercial or other purposes.

As a point of contrast, current-day background checking at security checkpoints such as airports is based on a remarkably small class of data, such as whether the individual's passport number has been flagged, whether the person is on a watch list, and what countries he or she has visited. It may be that significantly more could be done before reaching the point of serious conflicts with civil liberties. For example, a "risk score" might be provided to a checkpoint based on data held elsewhere (in a future type of fusion center?) regarding prior criminal convictions, formal security-related investigations, or associations with known terrorists. To be sure, however, collection and use of such data could be to start down a very slippery slope.

Intellectual, Ideological, Religious, and Other Motivations

The "Developing Intent" phase often involves ideological, religious, or political motivations, but may sometimes relate instead to the desire to be part of something larger, to have comrades, and to experience action and excitement. Other times, it is associated with the reality or perception of oppression or direct challenge to one's religion, culture, or people. That is, the range of motivations is broad, and it is a mistake to imagine that terrorism is uniquely related to Islam, to religion more generally, or to political ideology: Sometimes it is, and sometimes it is definitely not. This subject of motivations has been reviewed elsewhere,[*] but it is relevant here because motivations play such a strong

---

[*]   See Helmus (2009) and Paul (2009) in a review of social science for counterterrorism (Davis and Cragin, 2009). See also subsequent work discussing social-movement considerations in Davis, Larson, et al. (2012). An Army-sponsored report reviews diverse social and psychological theories of radicalization (Crossett and Spitaletta, 2010). A Department of

role and because people are often relatively overt as they develop the motivations, even if they later become more secretive. They typically work with and need a larger organization, which increases the potential opportunities to find indicators of motivation-building activity (e.g., radical prayer-group meetings in the Middle East or expressions of hate and braggadocio by domestic right- and left-wing groups).[*]

Looking at Cases for Insights and Indicators

Case histories can provide insights and possible indicators. A few examples are possible from the public record, drawing on direct testimony of individuals involved. These include the case of American-born jihadi Omar Hammami, who was president of his high-school class and well liked in a small Alabama town (Elliott, 2010).[†] The Elliott article discusses his intellectual development in high school and college, his highly observable search for meaning, and his progression toward violent extremism, the outcome of which does not at all seem to have been inevitable. Another case is documented in a BBC film *My Brother, the Jihadi* (Leech, 2011). Again, the jihadi's development was highly observable.[‡]

Defense white paper collects short papers on the "perception-to-intent dynamic" and another compilation discusses social, neurobiological, and complexity-related issues relating to motivations (Astorino-Courtois et al., 2012). Dipak Gupta (Gupta, 2008) provides an excellent life-cycle-of-terrorism book. Some of the primary authors on these matters include Bruce Hoffman (Hoffman, 2006) and Marc Sageman (Sageman, 2004; 2008).

[*]   A well-studied case of right-wing radicalism in the United States is based on experiences of Thomas Martinez, who—after years of involvement and ties to leadership in the white supremacist organization The Order—broke away and became an FBI informant (Martinez and Guinther, 1988). Martinez grew up with anger, resentment, and disillusionment; he fell into crime and associated with strong people who offered him a path, comradeship, and excitement (private communication with Martinez). Left-wing radicalism has also yielded some insightful discussions, such as a book on true believers and charismatic cults by an author with personal experience in the 1960s (Lalich, 2004).

[†]   In mid-2012, when he was barely surviving in Somalia, Hammadi gave a remarkable news interview expressing his attitudes about Islam, Sharia, and the United States (Putzel, 2012).

[‡]   The filmmaker's half-brother, Richard Dart (who renamed himself Salahuddin), was arrested by British police shortly before the 2012 Olympics on charges of preparing for acts of terrorism (Guardian, 2012).

A third and more problematic example is an "autobiography" by Omar Nasiri (a pseudonym), who was a jihadi in the 1990s but came also to spy for Western intelligence services and, eventually, to be dropped by them (Nasiri, 2006). The author describes a lengthy and fascinating history of alienation, street crimes, drug dealing, early minor-league involvement with radical Islamists in prayer meetings and other motivational settings, arms smuggling, associating with numerous terrorists and expediters, going to training camps, and—along the way, but in a decidedly non-intellectual way—"picking up" belief in the jihadi cause based on a sense of brotherhood with other Muslims. He shifted back and forth over time in his attitudes about large-scale terror attacks.[*] Assuming reasonable credibility, Nasiri's book suggests that someone like himself could leave an informative digital trail for years, a trail of the type that could be valuable when making sense of fragmentary information.

Another class of ongoing research has provided indicators based on information directly from terrorists. John Horgan and colleagues have been studying and interviewing those who have left terrorism for varied reasons (Bjorgo and Horgan, 2009; Horgan, 2009). This seminal work began with the Irish Republican Army and their families but has been extended to Middle Eastern cases.

Less Specific Motivations

One underlying cause for young people who seek meaning, ideology, and perhaps a movement to join involves some combination of broad disorientation, cultural confusion, and alienation. It can be said that if a person suffers from these and is angry, he need not have a specific reason for violence.[†] Since these feelings are common in some societies

---

[*]   Nasiri's book has been disputed, with suggestions that it was significantly embellished by the publisher and (by implication) British intelligence (Moniquet, 2006). Nonetheless, the author was interviewed by CBC and *New York Times* reporters, who were able to verify some aspects of his story. Michael Scheur, previously head of the bin Laden desk for the CIA, said in an introduction to the book that the book rung very true with him.

[†]   We draw on workshop observations by and discussion with UK sociologist Frank Furedi in research about extremism in Europe (Arana, Baker, and Canna, 2010). During the workshop, Furedi also argued that extremism and terrorism in Europe should be recognized as

for many and complex reasons, it is unlikely that they will be the basis of useful behavioral indicators. We mention them in part to counterbalance the tendency to look for more pointed causes and indicators.

Psychological Convergence

Although overlapping in part with activity of "psychological convergence," the previous section largely focused on the relatively early development of motivational substrates (a mixture of social, emotional, and cognitive) that can, in some cases, lead to participation in terrorist or insurgent attacks. Psychological convergence focuses on indicators that an individual or group is *committed* to involvement in a cause, belief system, or group with violent intent. As such, this activity generally represents a more advanced developmental state (or at least a psychological state more proximal to an actual attack) in which ideological radicalization and/or the commitment to involvement in an attack is increasingly solidified. Whereas an individual may previously have been searching and experimenting with ideas, causes, or groups, this activity may include such indicators of radicalization, as (1) statements of ideological adherence or attempts to proselytize to others, (2) public statements of ideological beliefs or explicit declarations of malign intent against a state or civilian target population or group, and (3) attendance at events that might support ideological commitment to a cause with terrorist underpinnings (e.g., attendance at radical mosques or political meetings). At some point, of course, a radicalized individual intent on violent action may become much more covert about such activities, or may assume a cloak of more restrained motivations.*

---

a *lifestyle*. He also noted the huge gaps between cultures as to how matters are viewed and what is "extreme." He mentioned anti-consumption, anti-materialist, anti-modernist, and conspiratorialist themes within the European communities of concern. Consistent with a caution throughout our report, he warned against seeing a linear, deterministic, and consistent radicalization process.

* The terms "radical," "violent," and "extreme" have diverse and contradictory meanings in both normal language and within the official and scholarly communities. To be "radical" or even "extreme" is viewed positively in some communities, with no implication of violence,

Such indicators led the father of Umar Farouk Abdulmutallab to report his son as a danger to two CIA officers at the U.S. Embassy in Nigeria. Abdulmutallab began attending radical religious meetings and spoke of leaving "for the course of Islam" before attempting to detonate an explosive on a plane from Yemen to Detroit (Hosenball, Isikoff, and Thomas, 2010).

As discussed in Chapter Six, indicators of activities in the psychological convergence category can sometimes be mined from social media and unstructured text.

Recruitment or Joining

The final activity in this phase relates to participation in terrorist or insurgent organizations. Examples would be the recruitment and radicalization of the September 11 hijackers (National Commission on Terrorist Attacks, 2004) and those involved in the 2005 London-subway bombing (Sciolino, 2005), who were carefully selected, vetted, and trained.

Recruitment and joining obviously does not apply to lone wolves. Also, in areas of heavy ongoing conflict, individuals may skip straight to recruitment or joining without lengthy motivational development or psychological convergence, as when some suicide bombers are moved to violence by a recent death of a family member or house demolition (*Los Angeles Times* Staff, 2012). This "direct-to-recruitment" developmental pathway may be especially likely where involvement in resistance movements is socially normative and a majority of the population is directly involved or complicit in its support—as in certain areas of Northern Ireland, Sunni neighborhoods in Iraq, and Kandahar in Afghanistan. In such environments, involvement in violent terrorist or insurgent groups might be as normative and prevalent as involvement

---

much less terrorism. In other cases, the terms "extreme" (and sometimes "radical") are used specifically to distinguish the violent or potentially violent from others who may also have strong views. One consequence is a tendency to use multiple adjectives, as with "violent extremist organization."

in police or military organizations in the United States (or, as another example, involvement in gangs in certain neighborhoods).

Recruitment or joining may be relatively formal or informal and may involve clandestine or completely open procedures, including use of training videos.* Membership in terrorist or insurgent groups might also be part-time, as in Afghanistan, with major activity during the fighting season and "piecework" actions, such as burying a jug of homemade explosives, for quick cash, with individuals otherwise make a living as subsistence farmers. Recruitment into terrorist organizations may involve deception or coercion. For example, in 2004 a Palestinian boy suffering from severe Down's syndrome accepted an offer to wear a bomb vest to an Israeli checkpoint with the promise of later compensation (Daraghmeh, 2004). In others, volunteers may travel long distances or even pay money out-of-pocket for a chance to officially join terrorist operations and be a "part of the fight." Of course, these aspects of variability in the recruitment and joining process affect the presence as well as the relative ease of observation for behavioral indicators associated with recruitment and joining terrorist and insurgent organizations.†

The Internet has many video segments showing recruitment and handling of suicide bombers. Some are anonymous YouTube uploads; others come by way of reputable news media. A review essay on suicide terrorism provides a good overview (Crenshaw, 2007).

Summary Activities and Indicators

Table 2.1 pulls together the topics discussed in this chapter. The first column lists generic indicators. Subsequent columns illustrate how they may relate to the several activities (they might sometimes occur

---

*   A number of recruitment-related videos are readily available on the web. See, for example, Zubaydah (no date). Much more warlike and action-filled videos can be found readily with an Internet search.

†   Some official documents (e.g., U.S. Army Training and Doctrine Command, 2006) draw on past cases to characterize suicide bombers, their recruitment, and related matters.

under activities). Appendix C gives more specifics, as well as references to the literature.

Table 2.1
Behavioral Indicators for Developing Intent and Nominal Association with Activities

Generic Indicator(s)	Motivational and Emotional Development	Psychological Convergence	Recruitment or Joining
Reveal hatred, prejudice, trauma, or shame	●	●	
Exhibit motivation for prestige, glory, status	●		
Approach life and decisions through ideological framework		●	
Explore different organizations and strategies		●	●
Proselytize and adhere strictly to organizational standards		●	●
Impose or reinforce practices and beliefs in others			●
Show signs of reinforced commitment		●	
Form new connections in social network to known terrorist elements			●
Seek out, read, or post radical content		●	●
Attend radical mosques or events		●	
Change behavior at school or home		●	

Planning and Laying Groundwork

Except for spontaneous attacks that use immediately available weaponry or explosives, a planning phase usually exists in which individuals or organizations select a target, acquire or develop and test the necessary explosives or weaponry, develop a plan of attack, and train or rehearse. One example was the Oklahoma City bombing. Timothy McVeigh carefully selected his specific target: knowing the attack would receive much press; he picked a building with a large window façade to create a more dramatic post-explosion image. Also, he used a variety of discrete ways to acquire bomb-making material, preformed numerous tests to ensure correctness, developed a plan of placement, and rehearsed the drop (Smith, Damphousse, and Roberts, 2006). As with the previous phase (Developing Intent), not every attack involves all of the activities described below, and the entire phase may be skipped or play a fairly minor role—especially if targets, opportunities, and weapons are readily available to attackers. In this chapter, we pay particular attention to aspects of planning and laying the groundwork that yield an "out of the ordinary" behavioral signature potentially detectable by current or emerging observational technologies. We identified six contributing activities: (1) development of strategic priorities; (2) materiel acquisition, testing, and development; (3) target identification, surveillance, and reconnaissance (ISR); (4) development of concepts of operations (CONOPs); (5) training and mission rehearsal; and (6) long lead-time preparations.[*]

---

[*] As a reminder, these need not all occur and need not be sequential. Further, an individual or group may be conducting activities associated with more than one phase at a time, and may move fluidly back and forth among such phase-level activities.

Development of Strategic Priorities

Terrorist or insurgent organizations, and even individuals, often think carefully about the *types* of targets to attack, based on their strategic or tactical objectives. It was no accident that the World Trade Center was targeted on September 11th, and had been targeted previously in the 1993 bombing. As known from bin Laden's post-9/11 communications, the towers were exactly the type of targets that al Qaeda prefers: those that are high-profile, have potential for many deaths, and are symbolic of free enterprise trade (National Commission on Terrorist Attacks, 2004). Organizations may issue corresponding directives to its operatives, such as al Qaeda's call for "homegrown jihad" on soft targets in U.S. and European cities. Even individuals consider what types of targets most fit purposes, as did Timothy McVeigh when he specifically targeted a U.S. government building (Smith, Damphousse, and Roberts, 2006). To the extent that strategic priorities are expressed in official decrees, individual communications, or other media (e.g., chatroom communications), such information can help mark not only the possible locations of impending attacks, but also the stage of planning.

Target Identification, Intelligence, Surveillance, and Reconnaissance

Beyond deciding target-type priorities, attackers must identify *specific* targets. This usually requires physical reconnaissance and surveillance, which may involve clandestine movement or camouflage to avoid detection if the target is heavily guarded or monitored. In some cases, would-be attackers will attempt to gather intelligence about targets through unwitting or complicit insider sources. They may even test the security perimeter around particular targets as part of their target selection and target reconnaissance and surveillance. This was the case in 2009 when Hosam Smadi attempted to detonate a bomb in a Dallas skyscraper. Before the attempted attack, which involved a dud weapon obtained from the FBI in a sting operation, Smadi did considerable reconnaissance to assess where best to place the bomb and where secu-

rity levels would allow him to do so (Trahan, Gillman, and Goldstein, 2009).[*] This type of activity involves deception and concealment and, sometimes, behavior designed to provoke and test existing security. Surveillance and reconnaissance can also be completed in a nonphysical manner, as when the 2008 Mumbai attackers used Google Earth to locate weaknesses in the security perimeter and locations to hide from security forces (Moreau and Mazumdar, 2008). The ease of remote surveillance and reconnaissance is growing as more advanced information becomes readily available to the public, especially via the Internet.

Materiel Acquisition, Testing, and Development

Unless the necessary weapons or explosives are readily available near the desired target (e.g., as with insider attacks on security-force trainers with weapons), the attackers must acquire these materials. We did not consider acquisition or movement of large payloads of material by insurgent or terrorist organizations (i.e., large-scale weapons trafficking), but rather we were interested in signs that individuals and small groups who might conduct an attack are in the process of acquiring, developing, or testing weapons. In some cases, this involves active experimentation, which can lead to detectable chemical or explosive hazards.[†] In other cases, this might be indicated by purchase records for explosive precursors or retail sales of firearms. Reports surfaced after the 2009 shooting at Fort Hood that, just weeks before, Nidal Malik Hasan entered a local gun store and abruptly asked for "the most technologically advanced weapon on the market and the one with the highest magazine capacity" (McKinley and Dao, 2009). In other cases, materiel is obtained through small-scale trafficking networks similar to those used by criminals.

---

[*]   The FBI complaint to obtain an arrest warrant provides considerable official detail (U.S. District Court, 2009), including lengthy discussion of Smadi's preparations.

[†]   Discussions with John Horgan, March 28, 2012.

Development of CONOPs

To prepare an actual attack, attackers often must develop a plan to train and move personnel to the target site, implant explosives, create decoys or other diversions, and execute the attack itself. This can either be a very simple plan (open fire on nearest cluster of soldiers, such as in the Fort Hood attack) or a much more complicated plan (e.g., deploy multiple groups of attackers simultaneously, as with the September 13, 2011, Kabul attack or the 2008 Mumbai attack) (Sengupta, 2009). The behavioral indicators of CONOP development itself are likely to be scant or difficult to detect, unless there is some leakage of communication among the team members or the attack involves repeated visits to the target site, a known safe house, or planning location to confirm or develop details of the attack CONOP.

Training and Mission Rehearsal

Once a broad plan of attack has been settled upon, it is sometimes necessary to acquire or improve marksmanship or other skills. Nidal Malik Hassan, who carried out the shooting at Fort Hood, visited an outdoor shooting range several times just prior to the attack, where he allegedly became adept at hitting silhouette targets at distances of up to 100 yards (Brown and Graczyk, 2010). Also, early reports indicate that James Holmes, the Aurora-theater shooter, had sought membership to a gun range (Associated Press, 2012b). Enrollment records for training programs and related travel records might be observable clues, but, of course, very few individuals who acquire such training conduct attacks. Another difficulty with this class of indicator is that training activities may occur shortly or long before an actual attack. As with almost every behavioral indicator, information on training is useful only in combination with other indicators as part of a more holistic analysis. For example, before the Columbine massacre, Eric Harris and Dylan Klebold bragged about newly acquired skills verbally and through online

communications and videos, thereby leaving a communication trail. They also had web presence with many troubling features.[*]

In some cases, attackers conduct a mission rehearsal, although seldom at the exact target site because of the risks of doing so. Although rehearsals soon before the intended attack will likely attempt to maximize operational security, they may still yield indicators of special value in suggesting details about intended targets, preferred mode of attack, and perhaps even the personnel involved. In some cases, groups or an individual engage in such "partial" rehearsals as challenging security personnel or alarm systems to test their sensitivity, but then quickly depart the area. An example of this is, as noted before, the "dry run" by the attacker on the 2009 attempted bombing of a Dallas skyscraper (Trahan, Gillman, and Goldstein, 2009).

Various federal, state, and local agencies study continue to develop doctrine on how to construct and make use of "suspicious activity reports, while respecting civil liberties" (U.S. Department of Justice, 2008).

Long Lead-Time Preparations

In rare cases, attackers cultivate inside sources, as did double-agent Humam Khalil Abu-Mulal al-Balawi, a Jordanian doctor who gained acceptance with the CIA before a suicide attack at the CIA's Camp Chapman that killed seven CIA officers. Another example was the Taliban's ruse about a desire for negotiations, which allowed the "turban bomber" to receive an audience with former Afghan President Rabbani, at which time he detonated a bomb in his turban (Rubin, 2011). Another fairly long lead-time example was the Irish Republican Army

---

[*]   Much of the online information about the Columbine massacre is erroneous, some of it even hoax material. Journalist Dave Cullen's book (Cullen, 2009) draws heavily on massive documentary evidence about videos, notebooks, and school assignments, as well as extensive interviews. The book illustrates well the massive number of prior indicators, but also the difficulty in observing them or making use of them before a major crime, especially when the duo was extremely deceptive and sometimes persuasive. See also Erickson, 2001, for an official report.

planting a bomb in the Brighton Hotel weeks before a planned visit by Prime Minister Thatcher. In terms of behavioral signatures, such preparations are not necessarily different than shorter-term preparations, except for the need for attackers to keep preparations and sources secret for a much longer period of time, which in turn yields more potential for detection. In retrospect, of course, there were many such observables prior to the 9/11 attack because complex preparations were indeed made for months. As with attack rehearsals, attackers with at least moderate capabilities and experience are likely to maximize secrecy and operational security in such long lead-time preparations. Indeed, skilled or experienced organizations are the most likely to conduct these types of complicated and highly involved attacks, which usually involve high-value targets or the intent to inflict massive casualties. Because they occur so far in advance of an actual attack, these actions may only be obviously connected with the attack after it has occurred (unless corroborating or supplementary suspicious behaviors are observed).

Illustrative Planning Behavioral Indicators

Table 3.1 records generic indicators (left column) and nominal associations of indicators with each of the activities discussed above. More details are in Appendix C.

Table 3.1
Illustrative Behavioral Indicators of Planning and Nominal Association with Activities

Generic Indicator(s)	Develop Strategic Priorities	Target Identification and ISR	Materiel Acquisition	CONOP Development	Rehearsal	Long-Lead-Time Preparations
Seek information on construction of weaponry and explosives			●			
Visit training camps and seek aviation or marksmanship training						●
Acquire dual-use electronics, explosives, ignition devices			●			
Conduct surveillance of target		●			●	
Use dry runs to simulate and practice attack					●	
Try to provoke or test security responses near target		●			●	
Release information or discuss how to harm or influence target population	●					
Experiment with chemical or explosive weapons			●			
Purchase explosive precursors or firearms			●			
Maneuver clandestinely or with camouflage near potential targets		●			●	

NOTE: For references and instances observed, see Appendix C.

CHAPTER FOUR
Immediate Pre-Execution

The "Immediate Pre-Execution" phase refers to the behavior of attackers and such support personnel as drivers and handlers in the period immediately preceding the attack (usually 24 hours or less) in what could also be termed "final preparations" for the attack. Due to the temporal proximity of these behaviors to attack execution, this phase could be the most useful and relevant for security and intelligence services attempting to detect and disrupt an attack before (or during) its occurrence. This temporal proximity also has relevance for a range of psychological and physiological changes and preparations, as well as alterations in social behaviors and actual physical preparations made by attackers in advance of attack execution. Because attackers are preparing to either risk or in some cases intentionally end their own lives, a number of significant social and psychological rituals and other potentially detectable processes tend to occur in the period immediately preceding the attack.

Some of these behaviors are unintentional physiological responses to the stress and cognitive burden of evading detection during clandestine movement, interactions with security personnel, or such other requirements of the attack as detonating an explosive device or firing on civilians. These unintentional "tells" include outward signs of nervousness, facial-expression "leaks" indicating deceptive communication or hostility, or even patterns of gross motor movement indicating deception or hostile intent. Such behaviors have received considerable attention by the Department of Homeland Security (DHS) in the past few years. Specifically, DHS and other agencies have explored the abil-

ity of human observers, as well as of technological tools, to detect these subtle indicators during routine observation or screening. Below, we pay particular attention to such behavioral indicators.

We identified five contributing activities: (1) psychological and physiological preparation for operation, (2) deception and concealment, (3) ritual practices, (4) changing patterns of social interaction,[*] and (5) logistical preparation for operation. As with activities in the other chapters, these do not always appear, may appear simultaneously, or in any order.

Psychological and Physiological Preparation for Operation

Changes in psychological behavior or physiological body function can be relevant indicators, although most of the following examples occur much more frequently in people who have nothing at all to do with planned violence. In any case, the current state of the art in research on the underlying physiological and neurological mechanisms for aggression and violence support the existence of alternate biological and behavioral pathways to violence. A review of recent research supports the existence of both individuals who tend to show a more "reactive" pattern of aggression and those who show more of a "proactive" or instrumental aggressive response (Scarpa, Haden, and Tanaka, 2010). The reactive pattern is typified by more pronounced physiological reactions to stressful stimuli (including heart rate and skin conductance response), while the proactive/instrumental pattern is typified by reduced physiological reactions to stressful provocation, and in some cases pronounced calmness (Patrick, 2008). These two behavioral and physiological extremes represent alternate responses to similar situations and stimuli. Which occurs in a given individual may depend partially on his or her personality and psychological profiles.

---

[*]   The changes referred to here, such as withdrawing from normal contacts, even with family, are different from those mentioned earlier, such as joining extremist groups.

Suicide bombers in particular seem to present one of two profiles (interview with Naval Postgraduate School professor Nadav Morag, 2012): an unusually relaxed and disconnected state (sometimes with indicators of euphoria), or an extremely nervous and tense state.

Studies of suicide bombers have noted that some attackers may go into a deep dissociative state, exhibiting little visible emotional response, feeling subjectively disconnected from ongoing events, and in some cases disconnected from physical sensation (Speckhard, 2008). Although discussed later as a separate activity, both personal and social rituals can help to induce this dissociative state, which has also been described among various groups of nonconventional combatants across the globe, especially after traumatic stress (Schauer and Elbert, 2010). A review discusses the associations among trauma, disassociation, and violence, noting that a cycle sometime occurs (Moskowitz, 2004, p. 38), which we indicate schematically in Figure 4.1.

Outward signs have been described as showing "no obvious emotion," having a "blank stare," or appearing to be in a "trance-like state" (Mullaney and Costigan, 2010). Significantly, this phenomenon only sometimes occurs, and eyewitnesses of other terrorist incidents have described perpetrators chatting or otherwise carrying on normally.[†]

While the proactive pattern of aggression is rare and tends to be associated with sociopathy and psychopathy (Patrick, Bradley, and Lang, 1993), it applies to at least some of those associated with terrorism and mass killings. Other attackers show different outward signs that might be described as a hyper-arousal and hyper-vigilance. This suggests strong short-term activation of both the sympathetic nervous

Figure 4.1
A Cycle Involving Trauma and Violence

Exposure to trauma → Dissociation → Violent behavior

RAND *RR215-4.1*

[†]   Personal communication with John Horgan.

system and the hypothalamic-pituitary-adrenal (HPA) axis (Cacioppo, Klein, Berntson, and Hatfield, 1993; Kagan, 1997; Meaney et al., 1996). It may include excessive sweating, shaking, nervous glancing, and other physiological and micro-behavioral signs of hyper-arousal (Mullaney and Costigan, 2010, p. D39).

Pre-execution activities may be indicated by Kinetic patters (i.e., body movement or gait), such as a demeanor indicating hostile intent, attempted clandestine movement, or carrying a weapon or bomb. For instance, as suicide bombers approach populated areas or other targets on foot to place an IED, they need to physically conceal it, as by attaching it to the target, covering it with road debris, or burying it in the road. Attackers using light weapons such as guns or knives may, of course, have to move very close to their targets. Research has identified different footstep types and rhythms by individuals attempting to avoid detection near a target location (Rowe et al., 2012).

Gait also provides clues about emotional state (Karg, Kühnlenz, and Buss, 2010). The relevant features include both movement and posture (Roether, Omlor, Christensen, and Giese, 2009). Hostile intentions may be exhibited by strong visible emotions, such as anger. Indeed, anger appears to be more easily detected than other emotions, and to aid in identifying human motion amidst random noise (Chouchourelou, Matsuka, Harber, and Shiffrar, 2006). However, emotions may sometimes be produced by multiple patterns of body movement (Dael, Mortillaro, and Scherer, 2012), implying a wide range of emotional gaits about which knowledge is incomplete. Anger, these authors note, was "very well discriminated from any other emotion" and "encoded with a very specific response profile, characterized by high rates of communicative and emphasizing gestures combined with forward body inclination." That is, angry individuals tend to gesture visibly and generally lean forward in posture. Research has also identified both hesitance and purposefulness of gait as behavioral indicators of imminent attack close to target locations (Kull et al., 2009, p. D35).

Changing Patterns of Social Interaction

While some attackers are lone wolves, eschewing all but the most casual or utilitarian forms of interaction (Pantucci, 2011), many others are not. For attackers who have been recruited from otherwise socially connected lives, the ramp-up toward an attack often involves progressive isolation from nonterrorist or non-insurgent elements of their social network, including friends and family members. Summing up evidence gleaned from interviews with Islamic militants as well as the analysis of several documented attacks and other sources, Guss, Tuason, and Teixeira (2007) describe the most common trajectory in patterns of social interactions in the time leading up to an attack:

> To dismiss further doubts and possible confrontation with contradicting views, the volunteers are often isolated from their families and friends. The volunteer's need for affiliation is met by belonging to a group of people who think similarly, and who are most likely to have had experienced the same oppression, outrage, and helplessness. (p. 426)

Isolation from elements of the attacker's social network who are not involved in planning, supporting, or coordinating the attack serves dual roles. First, this isolation helps ensure psychological commitment and restricts potential regret or other moral emotions that might interfere with having the conviction to conduct the attack; that is, it ensures that the master narrative of humiliation, justification, revenge, and redemption that drives much of terrorism and insurgency (Hafez, 2007; Post, 2005) is not questioned or interrupted. Furthermore, this social isolation helps to ensure secrecy and increase operational security by limiting contact with those outside the terrorist organization who might leak information or interrupt the plot at hand.

Ritual Practices

Preparations leading up to an attack may include not just practical actions and involuntary physical reactions, but also physical and social

rituals. These prepare the attacker for the psychological intensity of their upcoming actions (up to and including intentional death through suicide bombing). The rituals usually serve to cement resolve and may make it more difficult for them to renege on their commitment to attack. For example, the organizations that recruit, train, and deploy Palestinian and Iraqi suicide bombers create "martyrdom videos" in which attackers claim responsibility for their impending attacks (Kimmage and Ridolfo, 2007). These are released after the attack has been completed (Guss, Tuason, and Teixeira, 2007; Hafez, 2007). Making these videos and recording one's last words for public consumption make it more difficult for individuals to withdraw from the process. This process is part of a suite of social mechanisms and sanctions intended to ensure that suicide bombers follow their operational plan or else face embarrassment, shame, or even more severe social sanction (Ferrero, 2006; Bloom, 2005).

Rituals can also involve more specific activities tailored to the self or body in final preparations on the evening before or morning of an attack. While ritual preparation for battle is not unique to terrorist or insurgent organizations and has many historical counterparts, preparation for suicide or high-risk attacks particularly emphasizes the transition to the afterlife. This frequently involves shaving and washing the body in ritual preparation for entering heaven after death. Instructions for final preparations in incidents of Islamic terrorism are replete with religious references and familiar rituals or prayer and ablution. As such, this weaves steps of the attack with religious understandings and rituals designed to calm attackers and steel their will. Bruce Lincoln's analysis of al Qaeda's official textual instructions to the 9/11 attackers for their final 24 hours illustrates the role of rituals in preparing the attackers for an attack (Lincoln, 2006):

> "Shave excess hair from the body and wear cologne." In the last paragraph of the same section . . . cleansing one's body is described as ablution: a ritual act of self-purification that helps secure salvation.

> "Pray the Morning Prayer in a group and ponder the great rewards of that prayer. Make supplications afterwards, and do not leave

your apartment unless you have performed ablution before leaving, because the angels will ask for your forgiveness as long as you are in a state of ablution, and will pray for you." (p. 9)

Ritual preparation practices can also include instructing attackers to demonize their civilian targets (Guss, Tuason, and Teixeira, 2007) and to visualize themselves during the attack, seeing themselves as part of a great, mythical battle with religious implications (Lincoln, 2006).

Deception and Concealment

As mentioned earlier, under "Physiological and Psychological Preparation for Operation," one of the most important components of a successful terrorist or insurgent attack is deception and concealment, whether to pass through checkpoints, acquire materials, implant an explosive, manipulate a source close to the target, or avoid detection and interruption in the final approach to the target. A large and sophisticated body of research exists on deception and its detection, much of it sponsored by the Department of Defense (DoD) and various U.S. law enforcement agencies. Most instances of deception and concealment, of course, have nothing to do with planning attacks. Nonetheless, detecting deception and concealment are considered to be major elements of attack-related security. Here, we focus on elements most strongly connected with terrorist and insurgent attacks. We do not exhaustively review or evaluate the use of polygraph tests, but we do touch upon measurements used in these tests (galvanic skin response, heart rate, etc.) and whether the potential exists for such measurements from standoff positions or during screening. Such matters are discussed in Chapter Six.

The upshot of recent research focused on national security settings is twofold. First, people intending to commit a future criminal act and lying about it produce less plausible stories about their intent, as determined from subjective ratings of plausibility. Second, people lying about past criminal acts produce stories that are less plausible, less detailed, less internally consistent, and less consistent across inter-

views (Vrij, Granhag, Mann, and Leal, 2011a; Vrij, Leal, and Mann, 2011b). These results are based on subjective ratings conducted by lay (untrained) personnel, which suggests their usefulness for detection in basic screening procedures, rather than only in lengthy interrogation sessions.

Logistical Preparation for Operation

On a more practical level, terrorists and insurgents must attend to such logistical details as ensuring that explosives and other weapons are functional and distributed to those who will use them. Vehicles must be fueled and prepared, and individual attackers given their final briefings. In some cases, attackers travel to a safe house or other location near the target. Additionally, final surveillance and reconnaissance of the target is common (Hafez, 2007).

Logistical details might seem mundane, but are often given symbolic or religious meaning by terrorist organizations or the handlers assigned to suicide bombers. For example, an al Qaeda propaganda and training video illustrates how Abu Osama al-Maghribi (one of the Iraq UN Headquarters bombers) was overjoyed to hear that his wife had given birth to a new son on the day of his attack and gladly ran toward his IED-carrying vehicle to conduct jihad, which he accepted as his "new wife" (Hafez, 2007, p. 105). Similarly, Lincoln describes how the mundane detail of making sure the box cutters carried by the 9/11 attackers were sharp was given religious significance and imbued with meaning by construing the flight attendants who were to be killed as an animal sacrifice. Al Qaeda's textual instructions read,

> Check your weapon before you leave and long before you leave. You must make your knife sharp and must not discomfort your animal during the slaughter.

Table 4.1 summarizes illustrative indicators for generic indicators and relates them to the various activities of the pre-execution phase (see also Appendix C).

Table 4.1
Behavioral Indicators and Nominal Associations with Pre-Execution Activities

Generic Indicator(s)	Psychological and Physiological Preparation	Deception and Concealment	Ritual Practices	Changing Patterns of Social Interaction	Logistical Preparation for Attack
Separate from nonterrorist elements of social network				●	
Increase communication with terrorist elements				●	●
Take specialized actions to motivate self and co-attackers			●		
Give inconsistent responses to questioning		●			●
Give nonverbal signals of deception and lying		●			
Hesitate near target					●
Show accelerated heart rate, sweaty palms, thermal indicators	●	●			
Show micro-expressions of fear, hostility, deception, detachment	●	●			
Show indicators of instrumental aggression	●				
Exhibit body-movement patterns indicating hostile intent, clandestine movement, or weapon carrying	●	●			●

NOTE: For references and instances observed, see Appendix C.

Execution and Aftermath

This chapter combines the phases of Execution and Aftermath. We included these phases in a deliberate approach to look for plausible observables everywhere we could think to do so. Indicators from one attack's execution and aftermath could be valuable in detecting a future attack. We drew on existing historical cases, logical thinking, and knowledge from what occurs with criminal behavior to generate a long list of behavioral indicators. Probably none are truly new or unusual *except* if looking for indicators of these types by exploiting massive networked computer searches in near real time, by drawing on prior knowledge about individuals, and by fusion seeking to detect potential signal amidst a great deal of noise. For example, it might be that a search of travel data around the time of the attack (from a day or so before until the attack itself) could be focused on all of the individuals (pooled across databases) having even weak "tags" about possible associations to the insurgent/terrorist organization. It is possible that the hit rate would be small enough to have value, even if only to add marginally to the tagging of individuals arising in the search. A current example of this technology in use can be found in the Total Domain Awareness System employed by the New York Police Department in coordination with Microsoft (New York City, 2012).

Some terrorist and insurgent attacks are almost instantaneous—for example, those that begin and end with a single suicide bomb detonation. However, terrorist and insurgent attacks are sometimes "complex attacks," meaning they combine tactics, such as using multiple suicide bombs or firearms of various sorts and/or using multiple per-

sonnel attacking in stages. In such cases, attackers are emitting behavioral cues *during the attack itself* that might hold information about how subsequent attack stages will unfold and what people may be of concern.

Furthermore, after an attack has been completed, behavioral cues left in the forensic evidence of the attack can help provide a timeline of how attackers behaved over the course of the attack—for example, where and when they were positioned, how many rounds were expended and via what kinds of firearms, approximately when suicide bombers detonated, and the force of explosion/type of explosive material and detonators, etc. Such information can be combined with direct observational and other evidence (e.g., closed-circuit television [CCTV]) to provide a more complete picture of how the attack was carried out, providing information on insurgent and terrorist tactics, techniques, and procedures (TTPs) that can be leveraged to anticipate operational behaviors in future attacks or hints about identities of organizational members still alive. Forensic data are discussed in more detail in Chapter Six.

Finally, attackers continue to display behaviors after an attack. For example, surviving attackers may flee the scene (or general area), coercively silence (or kill) people close to the attack, or even brag about the success of the attack in the ensuring hours or days. Similarly, operational planners and groups will often engage in public "spin"—declaring the attack a success, inserting details or footage of the attack into propaganda or recruitment videos, or even arguing with opposing forces (for example, the battles between the Taliban and International Security Assistance Force [ISAF] on Twitter) about the success and moral righteousness of the attack. These post-attack behaviors can again provide clues about the targets, TTPs, and other features of potential future attacks, as well as surviving members of the organization.

This chapter briefly steps through the various activities associated with terrorist or insurgent attacks at the execution and aftermath stage, frequently relying on information about the September 13, 2011, attack on ISAF headquarters and the U.S. Embassy as an example but drawing on other attack cases as well. The contributing activities are (1) intelligence, surveillance, and reconnaissance, (2) deployment and

positioning, (3) coordination and communication, (5) target shaping and feints, (6) main attack(s), (6) post-attack reporting and strategic communication, and (7) protective measures and adaptation.

Intelligence, Surveillance, and Reconnaissance

Attackers perform final surveillance and reconnaissance on their target individuals or locations in the very moments leading up to an attack. For example, Palestinian suicide bombers are often given some leeway in the precise timing and targeting of their attacks. This may lead to hesitation near the moment of detonation as suicide bombers decide whether there is a large enough concentration to create significant casualties (interview with Naval Postgraduate School professor Nadav Morag, 2012). Meanwhile, bus bombers may look around and attempt to time detonation to maximize Israeli casualties or target specific groups, such as Israeli soldiers (Butterworth, Dolev, and Jenkins, 2012). Although asking for directions is hardly an indicator of something troublesome, it is interesting that attackers who have traveled a long distance may even directly speak with individuals to determine whether they have found the correct target.

Deployment and Positioning

As an attack commences, the initial positioning and subsequent movement of attackers can yield clues about the next stages of the attacks. For example, attackers during the September 13, 2011 attacks on the U.S. Embassy and ISAF in Kabul, Afghanistan, used suicide bombers at multiple checkpoints, perhaps hoping to distract from the bulk of the attack. Attackers who position themselves within line of site for particular targets could (if they were otherwise suspicious) be giving away clues that they intend to fire on these targets. For example, the

September 13 attackers in Kabul took up positions in an abandoned construction site with a clear view of ISAF and the U.S. Embassy).*

Coordination and Communication

During complex or multistage attacks, attackers must communicate and coordinate with each other to carry out a multistage operational plan (e.g., the Mumbai attacks) (Moreau and Mazumdar, 2008). This may involve direct vocal chatter over radios, although in certain environments this is avoided because adversary forces (i.e., the state or occupying powers) are known to have superior listening capabilities. In some cases, attacks may employ different personnel as "eyes on the ground" to fire on their targets and indicate visually (e.g., signaling with a mirror) when to detonate a roadside bomb. If video feeds are sufficient or attackers can be directly observed, it may be possible to pick up on communication patterns signaling clues for impending operational activities and subsequent attack stages. In the hours before the 2005 suicide bombings of the London subway system, the attackers communicated via mobile phone while getting into their positions to be able to detonate in quick succession (Sciolino and van Natta, 2005).

Target Selection, Shaping, and Feints

Attackers not only collect information and perform final surveillance on their targets before attack; in some cases they select, shape or attempt to shape the target itself. Last-minute target selection was illustrated by the "Passover Massacre" in 2002, in which the attackers drove around for some time before settling on the Park Hotel in Netanya, Israel (Intelligence and Terrorism Information Center at Center for Special Studies [Israel], 2004). Feint attacks were arguably illustrated in the September 13 attack on Kabul, in which attackers attempted to draw

---

*   Information on the September 13 attacks comes from several sources such as Rubin, Rivera, and Healy (2011); Aikins (2012); and live videos available on the web. Some of the interpretations are speculative.

protective measures away from the main attack area. Meanwhile, multiple-detonation suicide-bomb attacks, such as occurred frequently in Israel during the Second Intifada and which occur currently in Iraq, intentionally target rescue personnel or onlookers who gather at the scene of the attack (Butterworth, Dolev, and Jenkins, 2012). In some cases, attackers may try to clear the target area of individuals they consider to be friendly, such as occurred in moments just before Eric Harris and Dylan Klebold opened fire on students and faculty at Columbine High School. Just before the massacre, Brooks Brown, a classmate who had recently patched up a longstanding series of disagreements with Eric Harris, was warned by Eric, "Brooks, I like you now. Get out of here. Go home" (Merritt and Brown, 2002). As another example from the Columbine Massacre, shortly before arriving at the school, Harris and Klebold placed a small bomb in a field located approximately one mile away from Columbine High School. The bomb's explosion was set as a diversion for emergency personnel (Cullen, 2009).

Main Attack(s)

Behavioral indicators of a main attack are many and varied, but some include overt communication and coordination, moving rapidly (perhaps even in crowded automobiles) to attack points, running checkpoints, taking firing positions, etc.

Hostage-taking sometimes signals plans for a long standoff. The Beslan school siege in September of 2004 involved the capture of over 1,100 people as hostages (including 777 children). The attacks ended only after Russian security forces stormed the building with tanks, days after the situation began (Osborn, 2004). And just years before, during the Nord-Ost siege, forces decided they had to end the long hostage standoff by pumping chemicals in the ventilation system, which caused the death of 128 hostages. Indeed, hostage-taking as part of the attack strongly indicates that the attackers are prepared for a lengthy confrontation (Osborn, 2004).

Post-Attack Reporting and Strategic Communication

If part of a group, attackers who survive and flee may communicate back to headquarters about the success of their attack. Post-attack reporting may contain clues about the perceived success of the attack, as well as the overall mood of the attackers, planners, and leaders (exuberant, demoralized, etc.) (Jenkins, 2011).[*]

Meanwhile, public "spin"—whether directed to the media at large, other group members, or potential recruits—contains clues about ways that organizations mobilize and motivate potential recruits, including the "grand narratives" that speak to recruits' personal and moral sensibilities. For example, jihadist recruitment videos feature a narrative of humiliation and Muslim leaders' impotence, followed by the redemptive heroic action of suicide bombers and insurgents (Hafez, 2007). These materials are designed to bolster the commitment of current group members and drive future recruiting; thus, the "Developing Intent" phase begins again.

For the same purposes, attackers extend strategic communication post-attack. Oftentimes, because the attackers themselves are captive or dead, post-attack communication originates from the attackers' associates. After Major Nidal Malik Hasan opened fire on soldiers at Fort Hood, his "religious advisor," Anwar al-Awlaki, posted an online message, allegedly endorsed by Hassan before the attack, asking Muslim U.S. soldiers to follow in Hassan's footsteps (Hess, 2009). Attackers still living after their planned attack have also engaged in strategic communication. American teenager turned Islamic radical Zachary Chesser, who threatened the lives of American writers, wrote public letters to senators from prison, urging the creation of different foreign policies. Others use their highly publicized courtroom appearances as

---

[*]   This section deals with post-attack communications by the attackers. In some instances, attacker communications before or during an attack are intercepted but not analyzed until afterward. For example, according to ABC News, the National Security Agency intercepted and recorded two conversations in Arabic on September 10, 2001. One said, "Tomorrow is zero day," and the other said, "The match begins tomorrow." The messages were not translated until after the attacks (Ross and ABC News Investigative Team, 2011). Nonetheless, from the viewpoint of our analytic structure, these indicators belong to pre-attack or attack-phase indictors.

a stage for communication. A recent example is Anders Breivik, the right-wing extremist behind a bomb-and-shooting massacre that killed 77 people in Norway. Breivik has entered and exited every courtroom appearance with the Nazi salute and has crafted his defense statements in support of his views. In his plea, Breivik admitted to the acts but not to criminal guilt, saying the attacks were necessary "to protect Norway from being taken over by Muslims" (Associated Press, 2012a).

As for the attackers whose lives are lost in their attack, they still can, and have, personally engaged in post-attack communication. As noted earlier, it is common for Islamic attackers planning on dying in a suicide attack to make "martyrdom videos," which are released after their death (Kimmage and Ridolfo, 2007). The videos typically include a statement of purpose by the attacker preparing to be a martyr (Kimmage and Ridolfo, 2007). While the videos are usually released immediately after the attack, some, like the martyrdom video of the CIA double agent, surface later. Almost a year after he detonated a suicide bomb, killing seven CIA employees, al-Balawi appeared in a video calling the American team a "gift from God" (Oppel, Mazzetti, and Mekhennet, 2010): "Look, this is for you," he said to the camera, strategically appealing to his base and potential recruits, "It's not a watch. It's a detonator to kill as many as I can, God willing."

Protective Measures and Adaptation

Vacated offices or safe houses may contain clues about the attack, including specific tactics. The subsequent search of the LAX would-be-bomber's "safe house" in 1999 revealed significant planning details and valuable information about al Qaeda's organization, recruitment, and training (U.S. Court of Appeals for the Ninth Circuit, 2010). This was a fortunate find, since terrorist groups often attempt to cover their tracks, either through destroying evidence of planning or silencing collaborators who might leak vital information. For example, in response to infiltration, Palestinian militant groups took punitive actions against suspected "collaborators" (with Israel) in the Occupied Territories (Jackson et al., 2007).

If a terrorist or insurgent group concludes that an attack's impact was deterred or lessened by prior knowledge (or if post-attack arrests were made), this may tip them off about collection and detection methods. As a result, such groups adapt, continuously evading detection and maximizing the effectiveness of their attacks. For example, Palestinian militant groups have dressed suicide bombers as religious Israeli Jews or Israeli soldiers to enable them to escape detection by CCTV and get closer to their targets, while Jemaah Islamiyah has used camouflage to avoid aerial surveillance (Jackson et al., 2007). Adaptation and countermeasures are discussed further in Chapter Seven. Aftermath activities provide opportunities to connect the dots, understand target selection and purpose, and otherwise understand something about potential future attacks.

Behavioral Indicators of Execution and Aftermath

Based on the preceding discussion, Tables 5.1 and 5.2 show generic indicators and their relationship to activities of the execution and aftermath phases. See also Appendix C.

Table 5.1
Behavioral Indicators of Execution

Generic Indicator(s)	Deploying and Positioning	ISR	Coordination and Communication	Target Shaping and Feints	Main Attack(s)
Take action portending next stage			●		●
Indicate intent by nature of initial targets or give clues about future attacks					●
Drive car packed with fighting-aged males	●			●	●
Run checkpoints or security barriers				●	●
Split into groups (signaling multiple points of attack)	●			●	●
Run into buildings with cover or line-of-sight, indicating intent to engage targets	●				●
Shape population composition of target area				●	
Collect intelligence		●			
Interact with security personnel		●		●	
Prepare attack or feints				●	
Conduct main attack					
Communicate post-attack "spin" to media					

NOTE: For references and instances observed, see Appendix C.

Table 5.2
Behavioral Activities in the Aftermath

Generic Indicator(s)	Post-Attack Reporting and Strategic Communication	Protective Measures and Adaptation
Take responsibility or lay blame	●	●
Call to action	●	●
Idolize attackers	●	●
Communicate with headquarters about operational success or failure		●
Silence (kill, threaten, etc.) those with protected information about the attack		●
Clean evidence from safe houses and planning areas		●
Report troubles—including interdiction or interruption of attack	●	●
Develop new tools and CONOPs as workarounds		●

NOTE: For references and instances observed, see Appendix C.

Technologies and Methods

Previous chapters describe behaviors displayed by terrorists and insurgents prior to, during, or after attacks. This chapter addresses technology and methods for detecting such behaviors, grouping them in three cross-cutting categories of information: (1) communication patterns, (2) "pattern-of-life" data, and (3) indicators relating to body movement or physiological changes. The items in each category can be useful in observing behaviors in the activity classes used throughout this report: developing intent, planning and laying groundwork, immediate pre-execution, execution, and aftermath. Because our review is to inform research and development (R&D) management and investment, it is selective rather than exhaustive, and the topics covered vary greatly in development status and robustness. Thus, we include paragraphs on cautions and tables assessing development status, upside potential, measurement requirements, and shortcomings. Some of the issues are quite controversial, both scientifically and with respect to privacy and civil liberties. We touch on the primary points of controversy.

Detection and Analysis of Communication Patterns

Individuals communicate in many ways involving, e.g., face-to-face meetings, Internet chat rooms, and cell phones. Techniques exist to monitor these communications and to analyze their content. Indeed, both commercial and intelligence sectors invest heavily in such techniques. What follows draws only from work in the public domain. The first three subsections discuss online communication, text analysis,

and speech analysis. Drawing in part of the first three, the fourth sub-section addresses more explicitly threatening communications, which have been well studied in their own right.

Online Communication and Activities

It is possible to learn about some activities in the "Planning and Laying Groundwork" and "Immediate Pre-Execution" phases by tracking online communication and activities. Online statements and actions may reveal or suggest thoughts, emotions, or even intent. Thus, related tools and methods for analyzing online content and communications may be particularly helpful. Data collection itself can be performed manually, but is more efficiently done using online-content "scrapers."* These can pull in content constantly from particular sites or individual authors and can "flag" specific types of content.

Monitoring social media discussion for threatening communications is often the responsibility of human analysts, such as the New York Police Department's social media unit (Rock, 2011). The NYPD investigated threats posted on Twitter (e.g., "people are gonna die like Aurora") following the July 20, 2012, movie-theater shooting in Aurora, Colorado (Ruderman, 2012). Real-time social-media search tools can facilitate monitoring for discussions relating to potential violence. They may also track general discussion around such potential targets as landmarks, military bases, or upcoming events. Such social media search tools as Kurrently and Social Mention illustrate the tools available. Some large social media services, such as Twitter (Ruderman, 2012) and Skype (Timberg and Nakashima, 2012), have made content and user information available to law enforcement.

Knowing identities is crucial for tracking communications and interactions. Many users, however, seek anonymity—for any of many reasons—by creating accounts with no or false information. This tactic is also useful to those who plan, acquire radical ideologies, or discuss violent attacks. The magnitude of the fake-account problem was con-

---

* Related terms include "web harvesting" and "web data extraction." The technology is closely related to that for "indexing," which is central to the work of familiar Internet search engines. Numerous scraper tools are readily available for download.

firmed in our interview of a Facebook engineer who deploys detection and review tools for questionable content or inappropriate usage. There may be close to 83 million fake Facebook accounts (Wasserman, 2012).

Understanding online activities and their meaning requires knowing what individuals post or share and also what they read and consume. Scrapers or human analysts can track the former, while keystroke loggers or even downloadable computer malware can capture online consumption. Perusing an individual's documents or browsing history may show online searches for materiel (e.g., dealers, instruction manuals) or online registrations or payments to relevant training programs. Subtler methods leverage the increasing personalization of online services. For instance, if an individual is logged into such Google services as Gmail, Google attempts to autocomplete search terms, revealing previous searches. Google searches may also provide information on immediate location or other preferences. Marketing research exploits such methods heavily.

Spikes or trends in online activity may reveal shifting patterns of social interaction activity in the "Immediate Pre-Execution" phase, as when people stop posting or browsing, or cut off contact with family or other social contacts. This may relate to "going dark" before an attack. Conversely, an uptick in the communications of suspected terrorist elements or radical groups can relate to logistical preparations for an imminent attack or—as in the aftermath of the raid that killed Osama bin Laden—communications among those eager for vengeance. A number of attacks have been linked to such calls for vengeance, the first being a suicide bombing killing 80 people in Pakistan (Brulliard and Hussain, 2011).

A conference of DoD and commercial stakeholders sponsored by the Office of Naval Research articulated such key challenges as establishing the predictiveness of social-media data causality between social media data and future events, validating social media data amidst misinformation and deception, and the accuracy of sentiment analysis (Lyon and Afergan, 2012). Even tools that have not overcome these challenges, however, can flag potentially significant events for subsequent in-depth analysis. Given the tools' present status, such coupling of manual and automated analyses is an effective way of monitoring

social media data (Elson et al., 2012). We doubt that this need for man-machine cooperation will change soon.

A good area for future research is the relationship between online activity and content (what is posted and read) and physical actions (attending meetings, training, etc.).

Cautions and Shortcomings. Online communications and activities may reveal interactions between individuals or groups planning violent attacks, but false alarms arise when, for example, people search for information or purchase "dual-use" material that could be used either for attacks or—much more commonly—for entirely benign purposes. Also, indicators of actual hostile intent may be hidden within seemingly benign communications. Finally, high-quality encryption is increasing as companies such as Apple increase security options available to developers and users of smartphones (Garfinkel, 2012).

Text Analysis and Natural Language Processing
Content

The content of online communications can be studied with automated textual-analysis techniques, a number of tools for which are available for such purposes as academic study of text for signs of clinical depression or different styles of cognition. Text analysis can also help detect violent intent.[*] For instance, explicit content may include bragging or making ideological statements, as well as engaging or expressing admiration for known violent extremists such as the late jihadi leaders Osama bin Laden and Anwar al-Awlaki.[†]

---

[*]   Once again, we remind the reader that such indicators are seldom specific and may not be actionable.

[†]   As one example, Major Nidal Hasan (the Fort Hood shooter) wrote a series of emails to Anwar al-Awlaki, subsequently released by the FBI. One referred to Hasan Akbar, an American Army soldier who killed two fellow soldiers and wounded 14 others in Kuwait in 2003. The email, reported by CCN (Shaughnessy, 2012), said (with grammar errors retained),

> There are . . . many Muslims who join the armed forces for a myriad of different reasons. Some appear to have internal conflicts and have even killed or tried to kill other us soldiers in the name of Islam i.e. Hasan Akbar. . . .

Linguistic Style

Text analysis of *how* people write and talk can also shed light on thoughts and feelings. One prominent example is Linguistic Inquiry and Word Count 2007 (LIWC) (Pennebaker, Booth, and Francis, 2007; Pennebaker, Chung, Gonzales, and Booth, 2007), simple word-counting software identifying word-usage patterns statistically associated with motivations, attitudes, emotions, and other psychological states. These patterns can be analyzed to extract topics of discussion (Pennebaker and Chung, 2008).

Tools such as LIWC may uncover such psychological states as yearning for prestige and status, anger, humiliation, or shame—all relevant to the Developing Intent, Planning and Laying Groundwork, or Immediate Pre-Execution phases. Analyzing linguistic style has the distinct advantage that people cannot easily manipulate word-usage patterns in everyday conversation (Chung and Pennebaker, 2011). Further, linguistic style information is available regardless of the content or specific topic being discussed. Thus, while content analysis provides insight about topics, linguistic style analysis provides insights about the writer or speaker (Elson et al., 2012)

Linguistic style may also help flag extremist thinking. For example, use of third-person plural pronouns (e.g., "they," "them") may suggest "that the group is defining itself to a large degree by the existence of an oppositional group" (Pennebaker and Chung, 2008). Linguistic style, when combined with techniques identifying topics of discussion, may help identify potential targets of attack.

Analysis of linguistic style has already been applied to such topics as corporate fraud (Keila and Skillicorn, 2005), terrorist interrogations, and criminal testimony (Skillicorn and Little, 2010). Keila and Skillicorn (2005), for example, observe:

Would you consider someone like Hasan Akbar or other soldiers that have committed such acts with the goal of helping Muslims/Islam (Lets just assume this for now) fighting Jihad and if they did die would you consider them shaheeds (martyrs)?

As often happens, the e-mail could have been read at the time as not yet threatening.

Deception theory suggests that deceptive writing is characterized by reduced frequency of first-person pronouns and exclusive words, and elevated frequency of negative emotion words and action verbs. We apply this model of deception to the Enron email dataset, and then apply singular value decomposition to elicit the correlation structure between emails. Those emails that have high scores using this approach include deceptive emails; other emails that score highly using these frequency counts also indicate organizational dysfunctions such as improper communication of information. Hence this approach can be used as a tool for both external investigation of an organization, and internal management and regulatory compliance.

Individual Differences. Linguistic style may also reveal such attributes as age (Pennebaker and Stone, 2003), gender (Newman, Groom, and Handelman, 2008), health status (Pennebaker and Mayne, 1997), or emotional state (Alpers, Winzelberg, and Classen, 2005). Pennebaker and Stone found that, as they age, people use fewer first-person singular words and more first-person plural words. Newman and colleagues examined over 14,000 text files from 70 studies to analyze gender differences in language use, finding that men more often discussed "external events, objects, and processes," while women more often discussed "people and what they were doing." Greater use of first-person singular pronouns (e.g., "I," "me," "my") suggests self-focus and is sometimes statistically associated with depression (Rude, Gortner, and Pennebaker, 2004) and, perhaps, of tendencies toward suicide (Stirman and Pennebaker, 2001).* Use of second-person plural pronouns (e.g., "you," "yourselves") suggests attention paid to or focus on others and is associated with better health (Cohn, Mehl, and Pennebaker, 2004; Gortner and Pennebaker, 2003; Stone and Pennebaker, 2002). Using more positive than negative emotion words, such as when discussing a particularly traumatic experience, is linked with better physical health (Pennebaker and Mayne, 1997). Some work suggests that personality

---

* We observe, however, that such statistical associations are often weak and context-dependent. The heavy use of "I" and "we," for example, is normal (and part of good communications) in many contexts, but a bit "off" in others.

traits may be unrelated to linguistic indicators as measured by LIWC (Kahn, Tobin, Massey, and Anderson, 2007), but other researchers show more positive results with automated inference about personality based comparing a number of algorithms' performance experimentally (Mairesse, Walker, Mehl, and Moore, 2007).

Social Dynamics. Examining people's communications can indicate the nature of their relationships with others. This may help identify social networks and individual positions within them. For example, Tausczik and Pennebaker (2010) review research showing that people's language use varies according to their status relative to the listener. Higher-status individuals speak more often and more often use words referring to others (e.g., first-person plural pronouns, such as "we"). Lower-status individuals use more self-focused words (e.g., first-person singular pronouns, such as "I"). Knowledge about such social hierarchies may indicate, for example, ideological leaders or operational planners in terrorist cells.

Natural-Language Processing. Natural-language processing techniques should be useful for analyzing large amounts of text about which little is known in advance. However, these techniques attempt to solve extremely difficult problems in computer science. Accordingly, many competing algorithms and applications have been proposed.

These techniques can help detect behavioral indicators of violent attacks. Document classification can sort large amounts of text for subsequent analysis by appropriate technical or subject-matter experts. They might examine blueprints, how-tos, instruction manuals, statements of policy or intentions, news articles, or religious sermons. Topic mapping, such as by using clustering methods to identify concepts and topics discussed in text, could be useful for indicators in the Planning and Laying Groundwork phase by indicating particular searches for information or advice on weapons, tactics, or acquiring materiel. Mathematical techniques such as latent semantic indexing (or latent semantic analysis) can be used to understand concepts within the text and relationships among them. They have the advantage of being language-independent, with word order not playing a role. Machine translation can render foreign language texts into a form analyzable without foreign-language expertise or software specialized to the target

language. Speech recognition—determining textual content of spoken language—can greatly increase the amount of text available for text analysis. Similarly, voice recognition can help identify individuals from speech samples.

Text analysis and natural-language processing are complementary ways of learning about the content and implicit meanings of text. Linguistic style, however, is nonspecific. Changes in these indicators may suggest either individual or emotional differences, but interpretation is sometimes culture- or individual-dependent. This is one reason that natural-language processing can be used to categorize text or topics, but manual analysis of the content will likely still be required.

Cautions and Shortcomings. Despite the considerable past research, further work will be required before—if ever—linguistic style analysis can be *reliably* used to detect deception. Noting that much prior work has used archival emails, a Deloitte report (Mosher, 2010) argues that LIWC-based deception research needs further testing and validation on "real-life data sets." Similarly, in a national-academy review volume (Chauvin, 2011), authors Chung and Pennebaker (2011), pioneers in such work, point out the need to adequately understand the perceiver/listener of potential deception and the individual differences or situational factors that influence his/her judgment. They also say,

> Given the current state of knowledge, it is inconceivable that any language assessment method—whether by human judges or the best computers in the world—could reliably detect real-world deception or other psychological quality at rates greater than 80 percent, even in highly controlled datasets.

It should be emphasized that the validity of inferences based on linguistic cues depends heavily on such context. Cultural orientation for example, helps shape language use. To illustrate, a study of American and Japanese texts found that American authors used far more first-person plural pronouns, in a distant, royal-we manner, as compared with Japanese authors (Fieldler, 2007). Significantly, cultural orientation does not necessarily signify race/ethnicity, but rather the culture with which one identifies. For instance, a Caucasian-American living abroad could—but would not necessarily—have integrated

deeply into Japanese culture. It is also possible for cultural orientation to involve adopting beliefs and behaviors from religion or groups defined in ways other than race/ethnicity, such as a large corporation or military. Understanding cultural differences in language use may thus suggest ways to identify a person's affinity for particular groups or cultures from the way they speak.

Content Analysis of Speech
The vocal content of conversations and narratives has long been studied for cues indicating lying and deception. Robust indicators, based on multiple studies, show a significant connection between vocal content and narratives with lying and deception. These include (1) the tendency to distance oneself from lying statements (a classic example of this is speaking in the third person about actions one has taken) or seeming less verbally involved, (2) issuing discrepant statements (lack of internal consistency), (3) providing fewer details, (4) exhibiting less logical structure and less subjectively plausible stories, (5) providing less context, and (6) making fewer spontaneous corrections or admitted lapses of memory (DePaulo et al., 2003).

As described in other sections of this report, experiments with scenarios attempting to mimic potential terrorist attacks have shown positive results for some of these indicators—particularly subjective plausibility and the lack of consistency across statements or conversations (Vrij et al., 2011a; Vrij, Leal, and Mann, 2011b). As discussed in a later section of this chapter in greater detail, Vrij and Granhag (2012) argue that vocal content in response to well-placed probes or carefully crafted questions is the most reliable indicator of deception.

While such formal analytic tools as manual coding schemes are available, they are both long and complex to administer and there is no evidence in the literature that such formal tools perform better than trained subjective assessments. Thus, the fast-paced, real-time challenge of detecting potential attackers lends itself to some combination of (1) automated analysis of lexical content (which would require accurate translation from audio to text and is currently possible only under the best audio conditions) and (2) the use of trained human interviewers and observers.

A later section discusses content-independent analysis of voice reflecting physiological state, making more explicit comparisons with techniques based on vocal content.

Threatening Communications

Communications regarded as threatening are worth discussing separately, albeit with some overlap with the previous sections. Analyzing communications for threats or other behavioral indicators of potential attack may draw insight from research on threatening communications to public figures made by violent groups and individuals. Given the low base rate of actual attempted attacks against public figures, much of this research has focused on uncovering characteristics of communications (e.g., methods, content) associated with "inappropriate approaching behaviors" thought to suggest the potential for violence (Scalora et al., 2002). While much of this research predates the Internet, such approaching behavior has been linked with behavioral indicators that may be observable from online communications. For instance, Dietz and colleagues examined threatening letters sent to public figures, including celebrities and members of Congress (Dietz et al., 1991a, 1991b). They identified both "risk-enhancing" and "risk-reducing" features associated with inappropriate approach. Examples of the former include using multiple modes (e.g., letter writing, phone) and repeated communications.

Interestingly, *explicit* verbal threats are not necessarily predictive about violent actions because potential attackers often do not publicly communicate their intentions, although they may communicate privately (perhaps online) to family and friends (O'Hair, Bernard, and Roper, 2011). Similarly, Dietz and colleagues (1991a, 1991b) found that overt verbal threats either were not associated with or were strongly negatively correlated with "approaching behavior" with celebrities and members of Congress, respectively. Similarly, people with certain specific motivations for inappropriate approach (e.g., those with delusions of having royal identity) may simply approach public figures directly, rather than writing threatening or otherwise observable communications (James et al., 2009).

Such findings suggest that examining explicit threats is insufficient and that analysis should address implicit behavioral indicators as well. For instance, O'Hair and colleagues use a communications-theory perspective that explores how and why media facilitates informational goals (e.g., interactivity, comfort level in engaging others). They propose online communication metrics that include patterns of messaging that may reflect behavioral insight about motivations and source credibility. To analyze implicit violence-related content of speech, Chung and Pennebaker (2011) suggest coding for constructs such as dominance values (i.e., seeking power over others) and affiliation motives (e.g., seeking to maintain internal group relations). They also categorize linguistic features for how threatening text reflects actual violent intentions. These features can include either deception (i.e., bluffing) or honesty (which may indicate delusional beliefs or that the writer intends to carry out stated threats). Similar analysis may also reveal symptoms of severe mental illness, which is particularly prevalent in those who stalk public figures (e.g., Meloy, 2011; Scalora et al., 2002; Mullen et al., 2009; James et al., 2009). Observable symptoms include delusional thinking, particularly of persecution or grandiosity (James et al., 2009; Scalora et al., 2002) and fixation on the potential target of violence (Mullen et al., 2009). As in so much of what we review in this report, such symptoms seldom have much predictive value. Very few people with mental illness, delusional thinking, and the like are stalkers.

Cautions and Shortcomings

Given the abundance and variety of online media types now available, a number of online communication patterns similar to those found in offline communications may prove to be useful correlates of threatening behavior. One caution is that much of the evidence linking written communications with approaching targets predates online communication. Meloy (2011) notes that little work has been done to compare the differential relationships with approaching targets for online and offline forms of written communication.

Assessing the Communication-Pattern Methods

Table 6.1 summarizes our discussion of communication patterns. It has columns for development status, upside potential, measurement requirements (which may suggest opportunities or limitations), and shortcomings and vulnerabilities.

Table 6.1
Considerations and Caveats: Detection and Analysis of Communication Patterns

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Online communication and activities	Extensive collection and analysis occurs today for commercial and intelligence reasons. Technologies and methods for analyzing such online activities are still relatively unproven in either academic or operational settings.	Given trends, even more and varied interactions will be available for collection.	Tools already exist. However, challenges for dealing with massive volumes of noisy data are formidable.	Methods have not been well validated in academic or operational settings. Low signal-to-noise ratio. Effects of encryption, using "code," using anonymizers, or moving offline.
Text analysis and natural-language processing	A considerable research base exists with numerous past applications. Even natural-language processing can be highly accurate in specific experimental settings.	Using operational data to train and to create baselines could improve detection of deception, hostility, or extremist patterns. Natural-language techniques, given training sets, could quickly analyze large amounts of data.	Online text is naturally occurring and publicly accessible, requiring only passive collection. Active elicitation of text or oral statements is possible in some security contexts, such as checkpoints or interrogations.	Context and cultural dependence. Inadequate testing in operational settings. Need for substantial data.

Table 6.1—Continued

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Speech analysis: lexical and vocal cues	This has been validated in laboratory settings, including those specific to counterterrorism.	Advances in protocols for rapid assessment of speech patterns and content would have wide applicability for screening, checkpoint, or other situations involving conversations with security personnel.	Such analysis currently requires skilled security personnel asking questions and making judgments.	Physiological drivers, such as anxiety and changes in vocal tone, are individual-dependent May be subject to counters, especially if criteria for judging are known.

Pattern-of-Life Data

It is possible to analyze patterns of communication, travel, purchasing, and other matters using existing records and databases, many of which are held by private industry. This section reviews several of these data sources and discusses methods for combining and making sense of this information to assist with the detection of potential threats. We address (1) mobile-device-tracking, (2) use of existing records, and (3) machine learning and big-data analysis.

Mobile-Device Tracking
Mobile devices (e.g., smartphones) are ubiquitous and indispensable; they are an important new source of data for individuals' personal information and relationships in real time, and, in the aggregate, to identify social relationships, groups, and networks. Mobile phones can leave electronic trails involving Global Positioning System (GPS), cellular service, and WiFi data. An increasing number of consumer applications ("apps") provide such location-based services as recommendations and directions for nearby restaurants, resulting in a large amount of potential data. Many commonly installed apps also have access to

personal information stored on the device or obtained through other online services (often without the user being aware of this, much less consciously consenting). Technologies for tracking and analyzing such mobile data may be useful for "Developing Intent" activities, such as revealing individuals' links with existing organizations, extremists, or people with violent inclinations. Mobile data may also uncover such evidence of attack-planning activities as trips to known training camps or repeated visits to known or potential terrorist targets.

Because people typically carry their mobile devices wherever they go (unless fearful of surveillance), tracking the devices (e.g., using WiFi, GPS, or connections to cell towers) is a useful proxy for following their owners' movements. Over time, these data can be aggregated to build a fuller picture. Identifying and predicting people's location and movements has already received a fair amount of academic attention. Analyzing millions of mobile call records, Kang et al. (2010) computed typical travel ranges at different times for individuals of different age and gender. Others have also estimated the predictability in people's whereabouts (Jensen, Meservy, Burgoon, and Nunamaker, 2010) and future locations (Burbey and Martin, 2008). Successfully predicting the patterns of people's daily lives may help understand motivations or logistical preparations, although such models might not be able to flush out anomalies, such as preparations for an actual attack.

Mobile-phone data can help in understanding the spread of information and attitudes. Madan, Farrahi, Gatica-Perez, and Pentland (2011) modeled individuals' exposure to diverse individuals and political information and how the diversity or lack of it affects opinion change. Measuring diversity of information exposure and political opinion change may also suggest potential for identifying radicalization or "Developing Intent" activities. Smartphone usage data (Chittaranjan, Blom, and Gatica-Perez, 2012) are related to the "Big Five" personality traits: openness, conscientiousness, extraversion, agreeableness, and neuroticism (Costa and McCrae, 1990). These may provide some insight into motivations. Extroverted people are more likely to receive calls and spend more time talking. Emotional stability is associated with more incoming text messages. Chittaranjan and colleagues also found gender differences, which would help identify individuals.

Going forward, experimental mobile data may also benefit from using more naturalistic settings, such as using actual mobile phones rather than experimental sensors (Montoliu, Blom, and Gatica-Perez, 2012). Such research may clarify the extent to which patterns of mobile data-based indicators are related to actual behaviors and are stable across individuals. Other mobile devices may also be useful for understanding patterns of behavior and movement: not just tablet computers, but also such fitness items as personal-activity trackers, which are wearable motion sensors that capture individual movements and transfer information wirelessly to websites or social media. These actions may provide further mobile-based information about patterns of life.

Cautions and Shortcomings. As social networking through mobile devices becomes more commonplace, individuals are increasingly communicating with others that they never meet in person (Lampe, Ellison, and Steinfield, 2008). In some cases, relationships that occur entirely through mobile devices or online may be substituting for more "traditional" forms of social contact (Deresiewicz, 2011). The substitution is imperfect, however, and communication patterns and links derived from mobile-to-mobile communication is increasingly "muddy" and divorced from both intent to meet and intent to act in the offline or "real" world. Clearly, this complicates drawing inferences about actual threat. And, of course, people can choose to not use a cell phone.

Existing Records

It may also be possible to develop profiles of individuals from disparate pieces of information about a person's experiences, behaviors, and relationships over time, and to provide context for assessing other incoming data (this is a type of information fusion, as discussed further in Chapter Seven). Behavioral indicators in the Developing Intent phase may also prove useful when creating profiles of individuals.

Building individual profiles could use whole-life data (e.g., from school, criminal, civil, legal, interrogation, medical, travel, financial, consumer, and social/public communication). Consumer data are increasingly available on such matters because of electronic payments becoming the norm, leaving an electronic trail. General purchasing behaviors are stored by dedicated consumer data firms (e.g., Axciom)

as well as retailers (e.g., Target) and may help understand planning or other actions. Information on the purchasing of restricted material (e.g., firearms, explosive precursors) can be shared and flagged. Criminal data may be used to track activities indicating logistical preparations or planning for violent attacks, or that may suggest extremist motivations. Data systems for rapid intelligence have been used by police and security organizations. Online activities also provide a wealth of behavioral data. Social networking services store and analyze data on user activities. A Facebook "Data Team," for example, mines user data to better understand social relationships and interactions (Simonite, 2012). The team has produced reports showing how, for example, novel information is spread via social networks through distant contacts (Bakshy, 2012). Finally, public behaviors can be captured and stored (including by surveillance cameras, which appear in many public areas).

Integrating such information well would ideally exploit analytic techniques for all-source information fusion, some of it in real- or near-real time. Big-data analytic tools are increasingly available from such providers of cloud computing as Google and Amazon, while social-media analytic companies such as Topsy provide metrics of trending topics and sentiment. Techniques such as probabilistic unsupervised topic models may be used to extract topics from data (Farrahi and Gatica-Perez, 2012).

Cautions and Shortcomings. Clearly, extracting and combining data across multiple sources and owners within and outside the United States presents its own administrative and database challenges. Perhaps even more difficult is knowing what combinations of indicators should raise alarm signals. While particular purchasing patterns may be appropriate red flags in isolation (explosive precursors, bulk ammunition, etc.), many available records contain somewhat cryptic signals that are often discovered only in retrospect, as with subsequent discovery of Timothy McVeigh's troubled record in school and in the Army (Smith, Damphousse, and Roberts, 2006). After an event, a set of indicators over time can sometimes paint a relatively clear pattern of warning signals that existed beforehand, but the critical challenge is producing analytic systems that can notice these patterns *before* an attack takes place (perhaps in response to queries made at the time of screen-

ing or related monitoring). The information-fusion issue is discussed in the next section and extensively in Chapter Seven and Appendix D.

Machine Learning and Big-Data Analysis Drawing on Online and Other Activity

As data on human behaviors are increasingly digitized, the volume of potentially analyzable data increases accordingly, and manual analysis becomes impractical or impossible. It will be increasingly important to analyze such data without the benefit of prior hypotheses or known points of comparison. Doing so can help to detect outliers or anomalies, discover previously unknown associations and rules, and continually monitor data streams as a preventive measure.

"Supervised machine-learning techniques" rely on known data sets to train the algorithms, which learn and apply rules to classify data, identify relationships, and discover concepts. For example, one might wish to build a classifier to categorize online threats as either legitimate or spurious. Supervised learning requires a training data set where some threats were known to have led to violent attacks, and others were known to have not led to violent attacks. In other words, the data would be "labeled" as either having led to violent attacks or not. The algorithm would then be applied to the training data and tested on a separate set of data.

Techniques designed to analyze data without the aid of such known comparisons are referred to as "unsupervised learning." Researchers from Stanford and Google used thousands of images from YouTube to demonstrate the feasibility of building high-level feature detectors (e.g., faces) without providing labeled data (Le et al., 2012). It is easy to envision this technique being applied to learn and detect human bodies and, potentially, those that present violent attack indicators, such as hostile affect or carrying a weapon.

Such machine-learning techniques have been applied to national security and law enforcement issues, such as using data mining to uncover fraud (Li, Yen, Wu, and Wang, 2012) and using classification methods to predict deception in computer-mediated communication (Zhou, Burgoon, Twitchell, and Qin, 2004). Predictive policing efforts draw on incident reports, geographical data, and other informa-

tion to feed information into algorithms that generate crime forecasts (sometimes with hot-spot maps). Companies such as Palantir have used such techniques to build software that is widely used by law enforcement, intelligence agencies, and militaries worldwide (Vance and Stone, 2011). Quantitative approaches, such as artificial neural network models, appear to be promising ways to predict national security problems and can be applied in real time (Bueno de Mesquita, 2011). Video or image analysis and machine-learning techniques (Le et al., 2012) could be employed to discover, for example, activities in martyrdom videos, such as shaving heads and praying.

Cautions and Shortcomings. A major drawback of machine-learning techniques is that they require a large amount of data for model building and testing. An SRI International engineer with experience in machine-learning and national security issues suggested in a July 2012 interview with the authors approximate rules to assess where learning technologies could help a problem. First, if one does not have a strong understanding of the phenomena (e.g., indicators of violent attacks), then a lot of data are necessary (e.g., thousands to millions of instances)—data with the occurrence to non-occurrence ratio of the phenomena being relatively good (e.g., 1:4 rather than 1:1,000,000). That is demanding a great deal. Second, if instead one does not have a large amount of data (e.g., only hundreds of instances), a strong understanding of the phenomena is needed.

At least in the public domain, sufficiently large databases of violent attacks and other events do not exist for topics such as terrorism (according to the SRI International engineer we spoke with and others) or threatening communications and their relationships to actual behaviors (Chung and Pennebaker, 2011). Commonly used terrorist databases lack the necessary information. One innovative method for obtaining large, labeled data sets is to "crowdsource" the work of collecting and labeling individual pieces of information. The effectiveness of crowdsourcing has been demonstrated in other domains unrelated to terrorism, notably creating a dataset of emotional facial expressions (McDuff, Kaliouby, and Picard, 2011).

Assessing Pattern-of-Life Approaches

Table 6.2 is our assessment of the various approaches focused on records-based whole-life information.

Table 6.2
Considerations and Caveats for Pattern-of-Life Data

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Mobile-device tracking	Algorithms to predict individual movement patterns, preferences, etc., have been developed and validated in laboratory and experimental settings, but can benefit from more naturalistic validation.	Mobile devices will continue to add connectivity features that enable tracking (e.g., location and motion sensors, Near Field Communication chips).	Mobile device tracking may require device-owner permissions or cooperation of communications network providers.	Not traveling with or turning off device will defeat methods based on mobile-device whereabouts. Mobile-to-mobile communication is often divorced from "real-life" behaviors and intent.
Pattern-of-life data	Validating techniques to analyze large amounts of pattern-of-life data may be difficult in academic settings. Commercial data sets and analytic tools are increasingly available.	Pattern-of-life data may allow integrating disparate data types to build fuller behavioral profiles on individuals of interest. Accessing and integrating data is an issue.	Measurement does not require active or voluntary consent. However, access to various databases held by commercial or private sources may be necessary.	Pattern-of-life data may be vulnerable to "cover" activities and behaviors. Databases and algorithms for detecting threatening patterns are in early development.
Machine learning and big-data analysis	Machine-learning techniques have been extensively used and validated in experimental and some applied settings. Such techniques have been used in national security and law enforcement.	Machine-learning and big-data analysis may "discover" unknown patterns or activities hidden in large amounts of data, but massive amounts of data are needed for training.	Measurement does not require active or voluntary consent. A large amount of data or a strong hypothesis regarding relevant activity is required.	Learning techniques are probabilistic and vulnerable to noisy data. Current systems do not understand how to associate behaviors of multiple threatening individuals.

Data on Movement and Physiological State

Behavioral science has identified a host of nonverbal behaviors associated with emotional and psychological state as well as intent and motivation. These can be roughly categorized into gross motor movements (including the specifics of whole-body movement) and internal physiological changes with outward signs, including the "micro-behaviors" of facial expression. Each can potentially provide information about intent to deceive security officials or to carry out a hostile action. However, as detailed below, these indicators also suffer considerably from nonspecificity, context-sensitivity, and individual variability—factors that limit their potential utility in detecting pre-incident indicators of attack.

This section is long because so many strands of research exist and because many of them are quite controversial.

Kinetics and Gross Motor Movement

Existing technology can collect data for kinematic patterns (movement). Surveillance and reconnaissance platforms (e.g., tower cameras or systems on unmanned aerial vehicles) can monitor individuals as they maneuver before an attack. Video recording and transmitting devices can view individuals before attacks and collect information on individuals who frequent sites, providing a baseline for identifying individuals engaged in pre-execution activities. For example, "gait signatures" may be compared against information in a database analogous to that of the controversial early-in-century Defense Advanced Research Projects Agency (DARPA) program on Total Information Awareness (TIA) (Pugliese, 2008). Existing recordings of terrorism incidents (e.g., suicide bombings) may also provide baseline data for training new analysis tools. For example, Cohen, Morelli, and Scott (2008) proposed a method to model and flag potentially hostile intent gestures (including gait) from CCTV feeds for manual observation. This method, however, has yet to be tested experimentally.

Fewer tools exist for gait analysis than for data collection. DARPA has funded some biometric technologies, including HID (human identification at a distance) and VEW (video early warning) projects

(Pugliese, 2008). A major hurdle for gait analysis that seeks to identify individuals or their intentions by comparing against some normative standard is establishing good context-dependent baselines. Without baselines and careful analysis that recognizes contextual issues, certain sets of people may be misidentified or their intentions miscast. These may include people from other cultures or people with walking disabilities (Pugliese, 2008). Related to gait analysis of hostile intentions is recognition of potentially threatening body postures or poses, or "dismount threat recognition through automatic pose identification." Freeman (2012) used a machine-learning approach to "detect behaviors and postures that precede threatening actions/activities." This approach used a Microsoft Kinect camera, along with a training algorithm for classifying data. Results showed that the algorithms were better able to determine threats than to correctly identify postures.

Such machine-learning approaches are illustrative of potential computer science-based solutions to automated gait analysis, which may include, for example, algorithms, affective computing, and feature extraction. Incorporating emotion may be one way to increase the future utility of such computer science technologies. Affective computing may need to select from various psychology and neuroscience findings and theories of emotion (e.g., appraisal models). Often these models provide subsystems of monitoring and interpretation of stimuli, which can be computationally modeled. One further possibility, as suggested above, is to improve machine capabilities for detecting and recognizing emotion (Dittrich and Atkinson, 2008). For instance, improvements are possible when distinguishing between emotional states that differ in arousal, such as anger and sadness (Karg, Jenke, Kuhnlenz, and Buss, 2009).

Methods are being developed to analyze the gaits of people who may be carrying weighted objects, such as IEDs or firearms (Greenemeier, 2011). These methods could be applied to such existing technologies for cameras as the Microsoft Kinect and the Nintendo Wii for motion capture (Savva and Bianchi-Berthouze, 2011). Human observers and analysts may also be employed to detect these potentially threatening indicators (Blue Ribbon Panel, 2011).

A number of factors influence human performance in threat detection, only some of which are mutable. Training, of course, is one factor. There may also be relevant individual differences (e.g., stress or cognitive load, gender, expertise, motor experience) or bias in observers' detection abilities. Threat-detection-program setup or organizational factors can also play a role, and situational features can introduce bias. For example, as shown in a study conducted by Fessler, Holbrook, and Snyder (2012), people who are either seen or known to be carrying weapons may be perceived as taller, larger, and more muscular. Participants of the study were asked to estimate the height and rate the size and masculinity of men pictured in still photographs. The men shown brandishing a dangerous weapon were perceived to be taller, larger, and more masculine than other weaponless men of the same size.

Emotion—in targets as well as observers—plays a significant role in individuals' ability to detect or interpret gait or other body movements. People are most sensitive to detecting emotions associated with gait when the human walkers are expressing anger, as compared with walkers expressing other emotions or moving neutrally (Chouchourelou, Matsuka, Harber, and Shiffrar, 2006). Figure 6.1 shows the average detection performance broken down by emotion in that paper.

Observation of merely a single stride can be highly accurate (reaching 95 percent), suggesting that gait can be "an additional modality" for recognizing affect (Karg et al., 2010). Performance, however, varies by individual, suggesting a relationship with identity verification. Also, men and women differ in their ability to identify individuals based on gait observation. Compared to men, women may be more accurate at determining actions, and faster at recognizing emotions, from movements such as walking (Alaerts et al., 2011). Furthermore, the ability to recognize emotions from body movements appeared to be correlated to emotion recognition from facial cues, suggesting a generalized ability to recognize emotions and even biological motion that varies across individuals. As discussed in the same article, some specific features of gait detection are particularly difficult, such as recognizing an individual's gender. Participants in one study were not above chance at recognizing gender from motion, and in another study participants conflated angry motion with men and sad displays with women (Johnson,

Figure 6.1
Greater Sensitivity to Anger
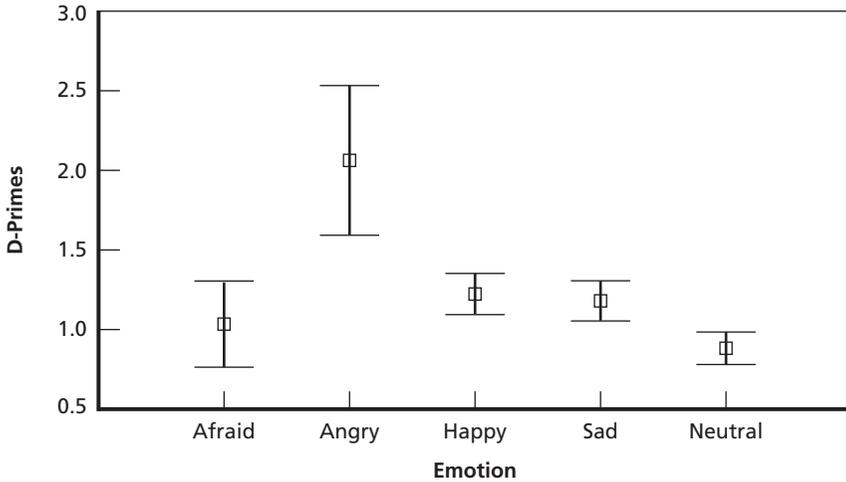
SOURCE: Chouchourelou, Matsuka, Harber, and Shiffrar, 2006, p. 68. Used with permission.
NOTES: The y-axis refers to sensitivity of detection. Large d-prime values indicate increased sensitivity.
RAND *RR215-6.1*

McKay, and Pollick, 2011). Yet even if a walker's gender is difficult to identify explicitly, it may be that gender implicitly influences perceptions of emotional gait. Halovic and Kroos (2009) found that fear and anger were identified significantly (p < 0.01) faster in female walkers than in male walkers. Figure 6.2 shows some results from their paper.

Observers may also be able to detect deceptive or clandestine movements. Research on deceptive motion has tended to focus on situations in which an actor attempts to deceive potential observers regarding (1) the nature of his or her actions (e.g., pretending to lift boxes as though they were different weights [Runeson and Frykholm, 1983]) or (2) whether people are truly performing an action (e.g., fake passing a ball versus actually passing it [Kunde, Skirde, and Weigelt, 2011]). Research has shown that people are better at detecting deceptive movements if they are themselves experienced in those deceptive actions (Cañal-Bruland and Schmidt, 2009). In the 2009 Canal-Bruland and

Figure 6.2
Gender Differences in Speed of Gait-Based Emotion Recognition
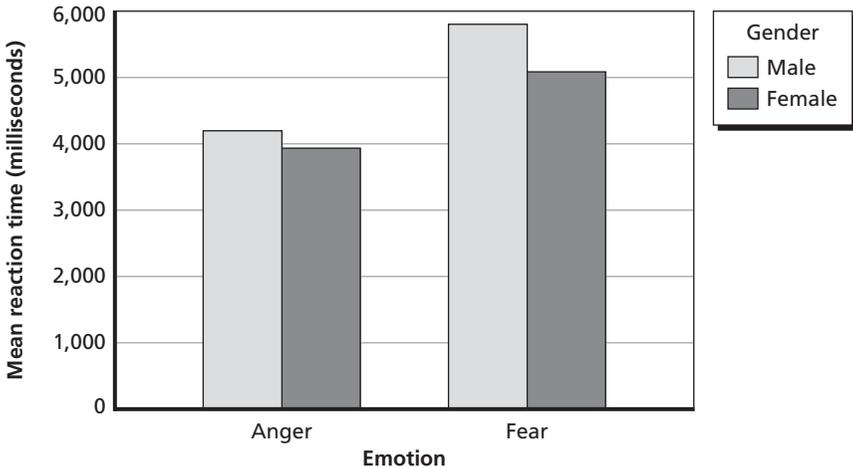
SOURCE: Adapted from Halovic and Kroos, 2009, p. 5. Used with permission.
RAND *RR215-6.2*

Schmidt study, veteran handball players and novices were asked to pre-dict whether a simulated player shot or faked a shot. Skilled handball players significantly outdid novices in discriminating shots from fakes. People are also more likely to recognize intentionally deceptive actions by observing significant kinematics. However, observer expertise does not help in determining deception when the body movement is inci-dental to the intended deception (Sebanz and Shiffrar, 2009):

> . . . studies have not investigated situations in which the body is consciously used as an instrument for deception. Rather, the focus has been on non-verbal signals that leak out without the individual's awareness. . . . Such a passive perspective on the body does not capture situations wherein movements are designed to be deceptive, such as when people fake injuries. . . .

Analysis of kinetics and gross motor movements should be appli-cable to a wide variety of security contexts, many of which involve people moving in relatively confined spaces of interest. Indicators of

emotion and action (e.g., lifting weighted objects) are fairly well understood, although validation in more naturalistic settings is needed. Furthermore, these indicators are not specific to attacks—for example, people who are anxious or upset are common at security screenings or checkpoints. Active elicitation of specific motor movements may help measure these indicators in naturalistic experiments to further understand gross motor movements.

Cautions and Shortcomings. One challenge for improving gait analysis (both machine and human) is that current detection systems and protocols are often built by (and algorithms trained on) "acted affect" rather than on naturalistic behaviors (Karg et al., 2010). Naturalistic (rather than laboratory) observation of moving individuals who are expressing anger or hostility are needed to ameliorate this problem. Savva and Bianchi-Berthouze (2011) offer an example of a way to use actual actions (using a Nintendo Wii) to capture and explore affective body movement, rather than relying solely on gait. Threat detection may also benefit from greater detail on human motion itself. Roether, Omlor, Christensen, and Giese (2009) propose algorithms to unpack specific features of how we understand emotional human motion.

Physiological State and Reactions

Observing physiological state and physiological changes holds promise for detecting deception and other behaviors, but the science base notes myriad difficulties and ambiguities.

Polygraphs

The best-known approach to using physiological indicators is polygraph testing. It has been extremely contentious for decades, and continues to be.[*] The most definitive review was accomplished by the National Academy of Sciences in 2003 (National Research Council, 2003). Most work subsequent to the 2003 review has echoed or embellished the original findings, maintaining that physiological responses to conversational probes are highly context-dependent and display dramatic

---

[*]   See the self-published Maschke and Scalabrini (2005) for a particularly harsh critique by authors who advocate against use of polygraphs.

variability within and across individuals, making their use questionable in courts of law (Vigluicci, 2009). Research has reiterated concerns about questioning techniques used during polygraph tests that may extract false confessions (Kassin et al., 2010; Porter and ten Brinke, 2010). Only a small subset of nonverbal indicators is (weakly) correlated with lying (Vrij, 2010).

Despite these problems, enthusiasm for the methods continues in law enforcement and intelligence communities, who argue that the methods are useful—as part of larger investigative processes[*]—to deter lying, loosen tongues, and generate information (including confessions). In such investigations, the guilty party is also relatively likely to be among those tested, raising the "base rate."[†] Polygraph methods, then, have proven value in forensic psychiatry (Grubin, 2010).

The value of polygraphs for national security screening is less clear-cut even than its use in criminal matters, in large part because the base rates are typically quite low. Honts and Schweinle (2009) highlight the role of base rate with a Bayesian "information-gain analysis," originally introduced in connection with assessing eyewitness testimony (Wells and Olsen, 2002), which describes the value of added information[‡] as a function of base rates of deceptive intent.[§] Figure 6.3, taken from their paper, shows results for a national security screening application. Information gain is the vertical axis; the base rate (expressed as percentage) is the horizontal axis. The results distinguish between information gain when assessing deception versus assessing truth-telling. Unfortunately,

---

[*]   These larger processes might include nothing more than an extra round or so of questioning, or they might involve trickery, psychological pressure, physical discomfort, and repeated rounds of interrogation. The National Academy study contains some case histories that are illuminating, both positively and negatively (National Research Council, 2003).

[†]   The base-rate issue is discussed more in Chapter Seven, since it is generic.

[‡]   "Information gain" is defined as the difference between post-analysis assessment of the probability of guilt and the base rate of guilt.

[§]   We discuss Bayesian combination and fusion methods more fully late in the chapter and in Appendix D. Although the Honts-Schweinle indictment of signal-detection theory/receiver operator characteristics (ROC) approaches is controversial (Rosenfeld and Penrod, 2011, p. 117) and seems to depend on how those approaches are used, their emphasis on the information-gain depiction of issues is new and has distinct advantages, in our view.

Figure 6.3
Information Gain Versus Base Rate for Polygraph Testing
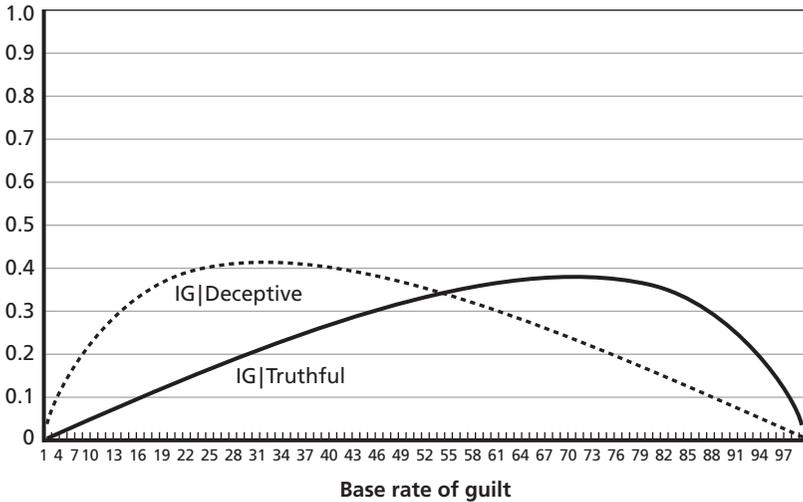
SOURCE: Honts and Schweinle, 2009, Figure 5. Used with permission.
NOTE: IG = information gain.
RAND *RR215-6.3*

base rates for national security screening will typically be low (near the bottom left), so little benefit is to be expected. However, if initial screening or other classes of information could increase the base rate to something more like 10 percent or more, the polygraph testing would have significantly more value. Here and elsewhere in the report, we conclude that candidate methods for exploiting behavioral responses appear at first not to have much potential, but that the potential can change significantly with use of combined information, as discussed in Chapter Seven.

Countermeasures, of course, are another problem, as discussed more generically in Chapter Seven. The empirical science on the issue is insufficient to justify strong conclusions but is rich enough to call into question a good deal of conventional wisdom (see the extended discussion and review in National Research Council [2003]). Drugs, for example, might seem to be an obvious countermeasure, but empirical studies have had mixed results about the effects of both drugs and

mental training on polygraph testing. Physical countermeasures, or a combination of physical and mental measures, are probably most effective. The use of countermeasures, moreover, can sometimes be detected by professional examiners, so that the counter becomes an indicator itself.

Methods Related to Polygraphs

Changes in such physiological indicators as galvanic skin response measuring sweat-pore activity and heart rate have long been components of polygraph tests involving sophisticated (and often lengthy) questioning paradigms. However, recent research indicates that more advanced physiological parameters such as respiratory sinus arrhythmia (RSA) may reveal clues regarding deceptive intent if measured during basic tasks unrelated to the deceptive behavior per se (Aikins, Martin, and Morgan, 2010). Since an individual need not be actively engaged in deception, there has been some optimism in the use of physiological parameters in standoff or screening environments. Single physiological parameters such as RSA are notoriously sensitive to individual physiological differences and competing physiological demands (Brownley, Hurwitz, and Schneiderman, 2000) and therefore tend to be very poor indicators of outcomes in isolation. However, the *combination* of several physiological parameters together holds more information and offers more potential predictive ability about future events (Gruenewald et al., 2006).

One such effort is called Preliminary Credibility Assessment Screening Systems (PCASS), developed by the Applied Physics Laboratory of Johns Hopkins University. It uses three sensors and something akin to a personal digital assistant. As discussed in an National Research Council workshop (Pool, 2011, p. 13):

> Two of the sensors are electrodermal sensors, which measure the electrical conductivity of the skin, and one is a photoplethysmograph, which is attached to a finger and used to measure changes in blood flow The signals from the sensors are fed through an analog-to-digital converter and sent . . . for analysis.

PCASS is specifically intended to detect deception by combining the signals collected. According to the workshop report (Pool, 2011), laboratory experimental evaluations showed performance significantly better than chance for both detection and minimizing of false alarms. PCASS has been used in combat, but without rigorous field evaluations (due to the dangers involved). Informal assessments vary from quite positive to much more skeptical, but the need for field evaluations is clearly recognized by the scientists involved, who cite (p. 16) unpublished 2008 work of Sujeeta Bhatt and Susan Brandon (previously of the Defense Intelligence Agency).

DHS is conducting experiments to determine the optimal combination of physiological indicators using standoff technology in efforts first known as Project Hostile Intent (PHI) and later Future Attribute Screening Technology (FAST), both of which are discussed later in this chapter. Sensors for detecting behavioral and physiological phenomena are an area of very active development, and include thermal, hyperspectral, laser Doppler, radar, and other detection modalities (Bornstein et al., 2010; Fein, Lehner, and Vossekuil, 2006). As many of these technologies are in the prototype stage, references in the literature are often relegated to the description of the basic features of notional or prototype detection systems (e.g., Derrick et al., 2010). As is apparent from these descriptions, conversations with experts (Middleton, 2011), and from our own review of physiological indicators, any assessment of the utility of a single detection modality must be considered in a systems perspective as part of a multi-modal detection suite; and thus the cost-benefit trade-off of any particular technology should be considered as part of a broader portfolio analysis (see Chapter Eight).

Voice Analysis

Considerable work has addressed voice analysis technologies, such as layered voice analysis and voice stress analysis, primarily in attempts to detect deception, to detect truthfulness, and to discern inaccuracies, high stress, high cognitive effort, anxiety, and overall physiological state. Measuring vocal pitch is easily accomplished; only a laptop and a microphone are necessary to evaluate observable "micro tremors" in the voice that are indicative of deception. For instance, Villar, Arciuli,

and Paterson (2012) used standard voice recordings and freely available audio editing software to analyze vocal pitch, then simply took the averages of vocal pitch in a given speech sample.

Vocal pitch and other nonlexical features of speech can be subsumed under the general category of "vocal stress" and are measurable via a range of commercially available devices, each of which uses a different combination of frequency, pitch, and other parameters to assess vocal stress. In part because of convenience and low cost, such techniques have been used by DoD, the Federal Bureau of Prisons, the Intelligence Community, and law enforcement agencies, including the Los Angeles Sheriff's Department. The techniques have reportedly not been embraced by the more cautious and skeptical intelligence communities (Pool, 2011, p. 11–12, referring to discussion by Philip Ruben). An earlier review conducted for the Air Force Research Laboratory (Haddad, Walter, Ratley, and Smith, 2002) concluded that such methods can—like polygraphs—be useful in helping to obtain confessions during interrogations. The work for the Air Force Research Laboratory also reported that such methods can detect stress at greater-than-chance levels and competitively with polygraphs and, as mentioned above, are less expensive, easier to use, and less constrained (Hopkins, Ratley, Benincasa, and Grieco, 2005). A 2003 review and meta-analysis covered upward of 1,300 articles about cues of deception and identified both vocal tension and higher vocal frequency as robust predictors of lying (DePaulo et al., 2003). There is some indication that these aspects of speech are much more difficult to control than other indicators of deception (Villar, Arciuli, and Paterson, 2012). However, more recent work, reviewing research over 30 years, concluded that the voice-stress technologies performed, in general, no better than chance (comments by Ruben in Pool [2011, p. 11] citing work by Bhatt and Brandon that appears not to have been published in the public domain). Ruben went on to say (without citations) that questions exist about the underlying physiological hypotheses. The perceived shortcomings had been part of what motivated the PCASS approach described above.

As with polygraphs, then, the techniques remain scientifically very controversial, and their value, which many still report, seems to depend on interrogator art, as in persuading those being interrogated to

believe that the technologies work. The methods have been described in such strong language as "charlatanry" and "pseudo-science" (Eriksson and Lacerda, 2007).[*] We found useful the full report (rather than its paraphrasing) of an experimental study done with prisoners in which prisoners were asked about drug use in prison (Damphousse, Pointon, Upchurch, and Moore, 2007; Damphousse, 2008). Only roughly 15 percent of drug users (as determined by blood tests) were detected by the voice measurements, meaning that the test was not very sensitive. More than 90 percent of those found to be nondeceptive were telling the truth about whether they were using drugs. That might be regarded as a good result, but it means that about 10 percent of those not being deceptive were misclassified. The ratio of false positives to true positives was roughly 9:4. The Damphousse study also saw no ability to distinguish between stress and deception. Its most positive result bore on something often mentioned by practitioners: The study showed that subjects were only about one-third as likely to be deceptive if they knew that the voice test and a computer program was being used, and therefore that their deception might be detected. This deterrence factor is known as the "bogus-pipeline effect" in the related literature. Another analysis also showed voice analysis results near those expected by chance, and with high false-alarm rates (Harnsberger, Hollien, Martin, and Hollien, 2009).[†]

Confusing Considerations. The literature on this subject is confusing and appears to be contradictory because some of it (particularly headlines) obscures key technical issues and logical differences. It seems that if subjects are being deceptive *and* are afraid of being caught and punished, then voice stress analysis will have a high detection rate (sensitivity). Also, it is not easy for subjects to "control" micro tremors

---

[*] The online version of the Eriksson-Lacerda article was withdrawn by the journal after the threat of a lawsuit for defamation by an affected manufacturer, but the article itself was neither withdrawn nor repudiated, despite claims to the contrary in some sources.

[†] Some very recent research, although in a different domain and of uncertain significance, is interesting. It applied voice analysis in the context of financial-analyst conference-call discussions with business managers and concluded that evidence of affect gained by voice analysis "contains useful information about a firm's fundamentals" above and beyond that from quantitative earnings data and linguistic analysis (Mayew and Venkatachalam, 2012).

of voice, suggesting resistance to countermeasures. However, if subjects are not especially worried (as when stakes are low), then there is no stress to detect, explaining some of the negative experimental results in laboratory conditions. It is unclear that the tests distinguish between stress and deception-related stress (Haddad et al., 2002, p. 19), no small thing given that subjects of questioning may be stressed for many reasons. Further, while subjects may not be able to "control" micro tremors, the value of micro tremors as an indicator can be drastically reduced if they occur randomly and frequently in the base population[*] or in the course of an individual subject being asked both control and target questions (e.g., curling toes, biting the tongue), by analogy with results obtained in polygraph experiments (Honts, Raskin, and Kircher, 1994) :[†]

> . . . the spontaneous use of countermeasures by untrained subjects has been found to be ineffective. . . . However, other research has shown that training in simple physical maneuvers, such a biting the tongue or pressing the toes to the floor, can be effective in defeating polygraph tests by enhancing physiological reactions to control questions. . . . [Honts and co-workers] reported that 60% of their decisions were incorrect when subjects had been trained to unobtrusively bite their tongues and press their toes to the floor when control questions were presented during the test. Using similar training and stronger incentives to pass the test, [Honts and co-workers] failed to correctly classify any subjects who were using countermeasures.

To add to these cautions, a scientific review of vocal stress analysis through currently available means concluded that such analysis still suffers from many of the same problems as other methods to assess

---

[*]   For example, we would expect an individual with hostile intent to be much more difficult to detect in a rowdy crowd or in a population of angry people going through an irritating checkpoint.

[†]   Some sources claim that Honts and co-workers say variously that voice-stress testing can be defeated by a "tack under the tongue" or a "tack in the shoe." We were unable to find such statements in the referenced articles and assume that the authors are extrapolating from Honts's earlier work on polygraph experiments.

deceptive or hostile intent (Haddad et al., 2002). That is, vocal stress can be caused by many different stressors—not just the stress of lying. Furthermore, vocal stress is both context and individual-specific in its baseline state as well as in its response characteristics (range, sensitivity to emotional provocation, etc.). Optimally, analysis of vocal stress to detect hostility or deception would first involve measurement of the individual (or a group of like individuals) in a "normal" nondeceptive state (see also Porter and ten Brinke, 2010). That is typically not possible when screening with clandestine, standoff, involuntary measurements, where no prior vocal data on the population is available. This increases the potential for total error (either false positive or false negatives, depending on where detection thresholds are set) in trying to identify potential attackers in a crowd.

In sum, detection of deception or hostile intent through automated vocal analysis is still in its infancy and suffers from nonspecificity; i.e., it is easy to detect stress but harder to detect more specific emotional or motivational states. A recent paper summarizes the status as follows, which captures our own sense from the literature (Elkins, Burgoon, and Nunamaker, 2012):

> Our voices are encoded with emotional information. While it is complex and difficult to develop software to classify emotion from the voice, it is possible. . . . It is unrealistic to rely completely on the voice to detect deception and hostile intent for all people and all situations. But, by exploring the vocal variables used by the software, we are able to correspond and fuse them with other detection technologies for higher prediction reliability and accuracy.

> Implementing an unreliable and invalid detection technology could place the country's security in jeopardy by failing to detect actual threats. Just as deleterious, however, would be to dismiss technology, such as vocal analysis, before it has been thoroughly examined. This would deprive DHS of a valuable tool for detecting threats and securing our homeland.

Electroencephalograms (EEGs)

Neuroscience has come into its own in recent years, but most studies are done in laboratory settings with individuals wired up to machines. Technologies have recently been developed allowing some physiological features to be observed without such wiring up, and sometimes at a distance, even covertly or surreptitiously (for example, using heat-sensitive cameras to detect capillary dilation and blood flow to the face and head to infer underlying patterns of central nervous system activation).* Much more seems to be feasible than one might expect intuitively. Also, there is some evidence (described below) of unique value in indicating *deception or imminent action* by an individual *if* baseline information is available for that specific individual ahead of time. Most of the technologies are in a relatively early stage of development, but some seem to have potential. Measurement of physiological signals closer to the central nervous system (i.e., the brain) holds more promise for detecting guilt and behavioral intent. An example is the work of Meixner and Rosenfeld (2011), which used electroencephalograms (EEGs) to measure response to specific stimuli relative to the individual-specific baseline.

In experimental settings, EEG responses to stimuli associated with a simulated crime proved useful in determining guilt or innocence (Meixner and Rosenfeld, 2011). Specifically, an event-related potential called P300 that indicates brain activity in the parietal area of the brain shows more pronounced amplitude if an individual views particularly salient stimuli. Thus, the accuracy of this detection method depends partially on prior intelligence of a planned attack, as it requires the display of something like key dates, attack modalities, or targets as stimuli to help distinguish suspicious individuals from innocent civilians. Using this technique, false positive rates tend to be low (< 5 percent),

---

* More of the ambitious undertakings are now becoming publicly available knowledge, as with DARPA's recent announcement of progress in its Cognitive Technology Threat Warning System, which combines a 120-megapixel, tripod-mounted, electro-optical video camera; cognitive visual processing algorithms that can be run on laptops to identify potential targets and cue images for operator review; and an EEG cap that monitors the operator's brain signals and records when the operator detects a threat (Defense Advanced Research Projects Agency, 2012). DARPA reports low false-alarm rates.

and false negatives low as well (5–10 percent). These neural responses to high-salience stimuli related to an attack are essentially unintentional physiological "tells" of prior involvement in planning a violent attack.

Figure 6.4 shows the waveforms retrieved in Meixner and Rosenfeld (2011). In this study, some participants were given information to withhold (the guilty group) and others were not (the innocent group). Figure 6.4 shows that the average Probe P300 amplitude (peak–peak) was seen to be in the guilty group when withheld information was displayed (solid line depicts P300 response to withheld information). However, the innocent group did not display increased P300 amplitude when the same stimuli were displayed.

Cautions and Limitations. To be sure, this technique has limitations. First, it probably requires specific credible intelligence about an impending attack so that relevant stimuli can be displayed, although guesses can sometimes be made about the potential details of a planned attack, as is common in criminal interrogation by skilled investigators. Furthermore, since P300 is generally responsive to the salience of stimuli, if individuals have lived near the attack target or were born on the displayed day, this will also produce a P300 spike—a potential source of false positives. Finally, an increasingly strong body of research indicates that individuals with aggressive tendencies (including general externalizing tendencies, impulse control problems, and reactive aggression tendencies) show a *reduced* P300 spike relative to normal controls. Such individual differences would decrease effectiveness of this tool for screening out individuals with hostile intent (Patrick, Bradley, and Lang, 1993).

It would seem that, at present, the most pragmatic approach is to flag individuals with physiological states far from the norm of the contextually determined population and to then check against evidence of nervousness or unusual calm (as in earlier discussion).

Facial Expressions

Much scientific work has concluded that humans share at least some universal facial expressions indicative of underlying emotional and motivational states (Ekman, 1970; Ekman and Rosenberg, 2005).

Figure 6.4
Average P300 Response over Subjects to "Guilty" Stimuli
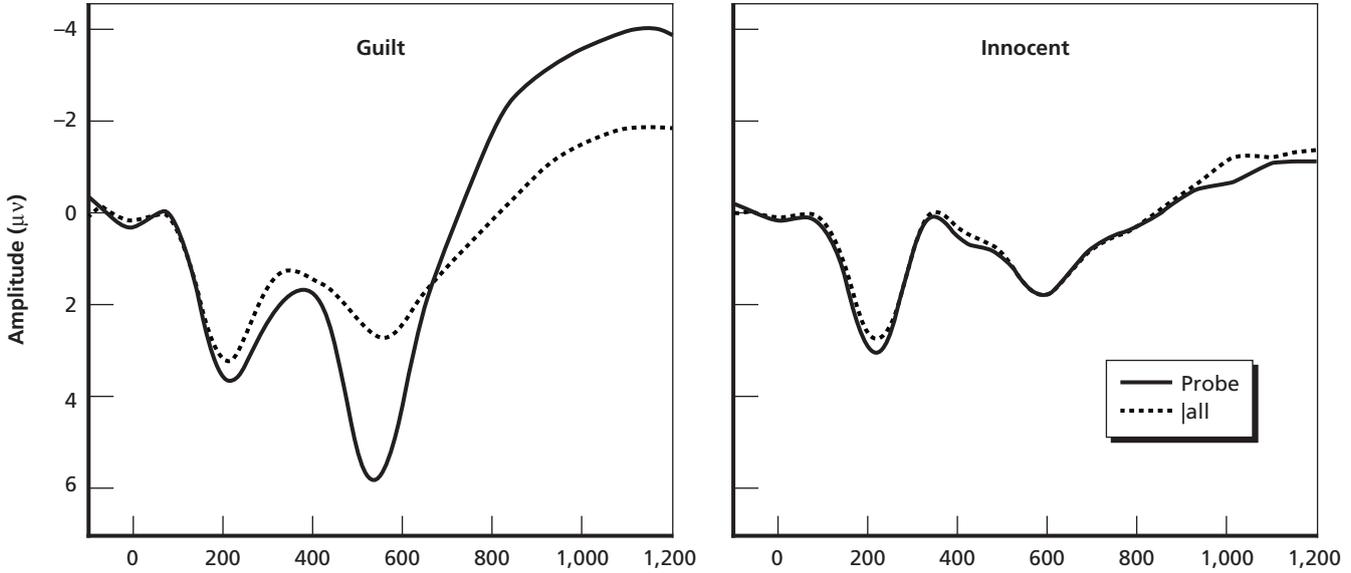
SOURCE: Meixner and Rosenfeld, 2011. Used with permission.
NOTES: The dotted curve is for "all irrelevant items"; the dark curve is for the items on which the subjects have and wish to conceal knowledge. Those are just irrelevant to the innocent subjects.

While there has long been some academic disagreement,[*] particularly with respect to cultural differences (Scherer, 1970; Russell, 1995), proponents of universal emotions have defended with rigorous data analysis (Ekman, 1992a; Ekman, 1992b; Ekman, 1993). Cultural differences seem relegated to the secondary dampening or accentuation of emotional responses, the categorization and perception of emotional states (Jack et al., 2012), and enculturated "display rules" that cause minor differences in predominant facial expression tendencies determined by major facial muscle groups (Matsumoto, 1990). A strong expression of the science seems to be that

> The seven fundamental emotions—anger, disgust, fear, happiness, sadness, surprise, and contempt—are displayed on the face with some fundamental features that are generally recognizable on all humans (barring neurological impairment).

For our purposes, the most promising domain of facial expression analysis is the detection of facial micro-expressions. Micro-expressions are involuntary expressions of fear, anger, or other emotions that display on the face for milliseconds, despite the best efforts of individuals to dampen or hide these expressions (Ekman, 2003). Whether the relevant behavior is early in the cycle of attack (e.g., smuggling weapons or traveling on forged documents), or closer to the actual attack (e.g., hiding anger, nervousness, or fear near security officials), facial micro-expressions potentially hold vital information about attackers and their intent. These micro-expressions can be detected via movements in the facial muscles that are coded as "action units," as displayed in Figure 6.5.

Drawing heavily from the work of Ekman and his colleagues (Ekman, 1970; Ekman, 1992a; Ekman, 1992b; Ekman, 1993), DHS is

---

[*] A popular-level summary of controversy regarding some of Ekman's work is Weinberger (2010), which draws on material from an unpublished JASON study and the claims of some of an inability to reproduce Ekman's work. For a journal-quality discussion, see Vrij and Granhag (2012) to the effect that the questioner-subject relationship is crucial and that assessment of unstimulated facial expressions (and other physiological observations) is ineffective. See Frank and Svetieva (2012) for a response. Controversy continues.

Figure 6.5
Example Facial Action Units (AUs) for Detecting Emotion

AU 6	AU 9	AU 12	AU 14
Cheek raiser	Nose wrinkler	Lip corner puller	Dimpler

AU 15	AU 20	AU 23	AU 24
Lip depressor	Lip stretcher	Lip tightener	Lip pressor
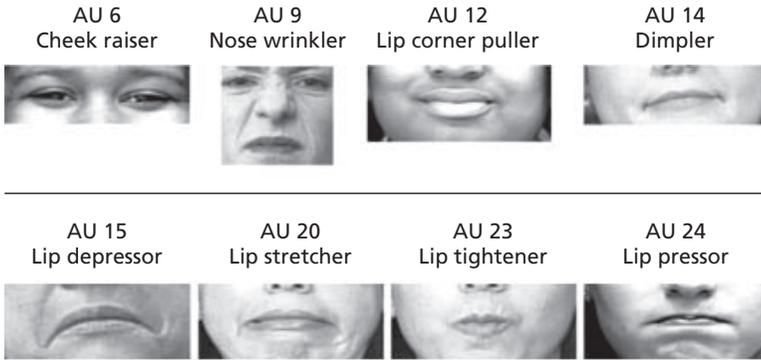
SOURCE: Sayette et al., 2012, p. 873. Images from Cohn-Kanade Database, copyrighted by Jeffrey Cohn. Used with permission.
RAND RR215-6.5

currently examining how both close-up and standoff detection of facial expression (or use of facial expression analysis in screening scenarios) might aid the early detection of impending attacks. Some existing research has shown that so-called "micro-expressions" of emotion (lasting only milliseconds) are "leaked" despite the best efforts of individuals to conceal them in social situations. While some individuals seem to be naturally more attuned to picking up on these micro-expressions (Posamintier and Abdi, 2003), individuals can also be trained to do so (Kanade, Cohn, and Tian, 2000).* Also, micro-expressions can be detected via frame-by-frame analysis of video footage (Ekman, 1970). Working with both Ekman and his colleagues, DHS has a two-pronged research program to determine the efficacy of facial expression for determining deception or hostile intent. This involves both

---

* We are aware that the existence of deception-detecting "wizards" is also debated, as discussed in a popular article (Weinberger, 2010). We are familiar with the relevant journal article (Bond and Uysal, 2007), a terse response (O'Sullivan, 2007), a more detailed response (O'Sullivan, 2008), and several increasingly grumpy iterations in the form of journal comments. O'Sullivan and Ekman note the important distinction between testing college students detecting low-stake lies versus testing, e.g., Secret Service agents in their ability to detect high-stake lies. A broader article of interest is Vrij (2010).

experimental research and the analysis of years of video footage from airport environments in which actual attacks (or attempted attacks) took place. While there has been recent controversy and disagreement about the use of facial expression for this purpose, only some of the controversy revolves around scientific validity versus ethical and legal concerns. There is severe criticism about using such behavioral cues to infer *intent*.

Cautions and Shortcomings. To some extent, facial expressions suffer from the same problems of nonspecificity as physiological signals, but less so. While changes in heart rate or blood pressure are driven by many non-emotional physiological demands and also may show similar changes across different emotional states, the facial expression of emotion is a biological system specifically for communicating emotional and motivational state. It is thus "closer "to actual motivational state and intent than peripheral physiological signals. However, the link between facially expressed emotion and actual behavior (or behavioral intent) differs dramatically across individuals. An attacker with a psychopathic profile might be more inclined to show micro-expressions of "duper's delight" (Ekman, 1981) while passing through a checkpoint undetected (hence, his subtle expression of delight). A non-psychopathic attacker might instead show micro-expressions of fear, which a perfectly harmless nervous traveler might also show.[*]

While the link between micro-expressions and deception is well evidenced, the usability and utility of micro-expressions in security-related settings is another matter. Coding emotional expressions for use in scientific studies currently involves a painstaking process of frame-by-frame analysis in which hours of labor are required to analyze seconds of data. Such a burden of analysis and time lag is too great for this technique to be usable in real time at checkpoints or other screening areas. However, Ekman has developed a training system, the Micro Expression Training Tool (METT), that appears to increase the capacity of individuals to detect facial expressions and micro-expressions, with demonstrated evidence of effectiveness in clinical populations

---

[*]   Video illustrations of duper's delight can be found with Google searches, but the web links may not be stable.

(Russell, Chu, and Phillips, 2006; Russell, Greene, Simpson and Coltheart, 2008).[*]

Furthermore, while recognition of emotional expressions based on automated algorithms and computation (rather than human observers) is still in its infancy (commercially available platforms can easily be fooled), this is an active field of development, and improved algorithms are likely to yield greater accuracy and robustness (e.g., Polikovsky, Kameda, and Ohta, 2009). Furthermore, as with many pre-incident indicators of attack, emotion-recognition algorithms that fuse multiple parameters (facial expression, speech, and gross motor movement) seem to perform much better than inferring emotional state simply from facial expression (Castellano, Kessous, and Caridakis, 2008; Caridakis et al., 2007).

Of course, checkpoints and other security environments are dynamic locations where it is difficult to capture high-resolution video (or audio or physiological data) for individuals, but such detailed information is necessary for the reliable detection of hostile or deceptive intent.

Video Data for Observing Kinetic or Other Indicators

Previous sections discussed behaviors and indicators of interest. This one discusses one way to observe those indicators.

Video image capture is becoming increasingly ubiquitous in many settings around the world, including long-distance observational assets in operational environments, as well as continuous monitoring through CCTV at government institutions, transit points, commercial settings, and even private residences. As a result, significant attention has been focused on the use of automatic detection of hostile, deceptive, or otherwise suspicious patterns of behavior from video data that signal either an imminent attack or advance preparation for an attack (e.g., surveillance behavior, planting explosives). Such approaches have also been applied to the security screening and interviewing context; a

---

[*]   Some of this is familiar to television viewers of the series *Lie To Me,* which was inspired by Ekman's work. The show, of course, is entertainment and takes considerable liberties for the sake of interest.

group of researchers at the University of Arizona has made significant headway in developing automated multimodal systems to simultaneously analyze patterns of kinetic movement in combination with vocal and speech patterns to detect deceptive intent (Jensen et al., 2010; Burgoon et al., 2009; Meservy et al., 2005a, 2005b). This is a limited example of information fusion.

Detection of behavioral state (and possibly intent) from video images faces many of the same challenges as detection of intent from the other behavioral (or behavior-related) modalities described in this report. That is, it is quite possible to design high-performing detection systems in controlled laboratory settings and "clean" data sets— for example, a single camera at a set distance in continuous and even lighting, a quiet room to aid the detection of vocal onset and other characteristics of speech (Jensen et al., 2010). However, "real life" operational settings introduce a large number of challenges, including distinguishing human actors from complex and dynamic visual backgrounds, classifying behaviors into distinguishable and meaningful categories, and fusing information across multiple camera feeds (Ko, 2008). Additionally, setting and context (including social and cultural context) affect the nature of behavior patterns that must be detected, which means that not only the background but the core set of threatening behaviors is a moving target. As a result, artificial intelligence (AI) approaches are at the forefront of current development efforts (Ko, 2008), and data mining or unsupervised learning approaches show promise (Hospedales, Gong, and Xiang, 2009; Ke-Xue, Guo-Hui, and Ya-Li, 2006; Peng et al., 2012; Zhongfei, 2002). Such fully automated approaches are very much in the development stage, and combined human-machine interfaces (i.e., using automated detection to "zero-in" on a video image for screening by a live human operator) may be the best interim solution (Cohen et al., 2008).

Forensic Data for Observing Kinetic, Physiological, or Other Indicators

Forensic techniques are another mechanism for observing the range of indicators mentioned above, and some others as well. Such techniques are constantly improving and highly pertinent to gaining knowledge

about the capabilities and behavior patterns of terrorist and insurgent attackers. Many forensic investigation tools involve highly technical analysis, including complex biological procedures involved in DNA matching, and the physics of deriving insights from destroyed structured and blast patterns. We do not cover such physical methods in this report, but a few observations are appropriate here.

Forensic investigations conducted after terrorist or insurgent attacks can result in significant new insights regarding the tactics, operational steps, targeting preferences, and capabilities of attackers. Forensic information can be gathered in the following ways:

- reviewing CCTV or other camera feeds (see previous section)
- interviewing bystanders or survivors of the attack as well as security forced involved in resolving the incident
- analyzing the severity and patterning of wounds incurred by casualties (both dead and wounded)
- examining weapons, planning documents, and any other material left by terrorist or insurgents at the site of the attack or nearby staging areas and safe houses
- examining the bodies of dead attackers.

These sources can provide a wealth of information about behaviors of terrorist and insurgent attackers. The unclassified literature mostly addresses attacks in the small number of arenas that do not directly involve U.S. interests (McCorkill and Griffin, 1998). Most recent forensic investigations of terrorist and insurgent attacks around the globe remain in the sensitive, limited distribution literature. What follows is our own cut, drawn from published accounts and logic, at a list of the generic types of information that might be extracted regarding terrorist and insurgent attackers from forensic investigations, including the techniques for information extraction described above:

- deployment and positioning of attackers, including feints designed to draw attention or resources away from the primary attack
- level of training and aptitude with weapons

- tactical style and techniques (e.g., rapid movement or changing firing positions)
- missed opportunities for detection or interdiction before the attack
- disguises or impersonation/camouflage (e.g., concealing attackers in large crowd) and other techniques used to approach target without detection or interruption
- demographics (age, ethnicity, gender) of attackers as well as number involved
- chronological reconstruction/timeline of attack steps
- probing and testing events used to determine vulnerabilities in security
- If attack was failed or impact mitigated, why and how?
- trends in attack style, targeting preferences, and weapons (comparison with past attacks, projections for future attacks).

Testing and Validation Attempts

Much of the relevant work to test and evaluate national security methods using behavioral indicators is not available to the general public, but some is.

The SPOT Program

One of the continuing failures in U.S. planning has been the failure to conduct independent, rigorous, peer-reviewed research on the efficacy of methods and performance of technologies. Responding to scientific criticism, this failure has been decried by the Government Accountability Office (GAO, 2010) and various members of Congress (see, e.g., statement of Paul C. Broun in U.S. House of Representatives, 2011), much of it with respect to DHS's Screening of Passengers by Observation Techniques (SPOT) program.

The SPOT program is intended to provide behavior detection officers (BDOs) with a means of identifying persons who may pose a potential security risk at Transportation Security Administration (TSA)–regulated airports by focusing on behaviors and appearances that deviate from an established baseline and that may be indicative

of stress, fear, or deception. SPOT may make referrals for additional screening that may include use of new or experimental methods.

In an effort to conduct sound experimental validation tests, DHS conducted a randomized trial study of the effectiveness of SPOT in identifying illegal behaviors.[*] According to congressional testimony, the study involved 43 airports that instituted a procedure of random selection for secondary screening in parallel with SPOT. The data set included a total of 71,589 random selections and 23,265 SPOT selections. As discussed in testimony, SPOT-initiated selections reportedly resulted in the detection of illegal behaviors at a significantly higher rate than random selections. In the experiments, only 2.8 percent of the general public exhibited the most common single behavior on the list of SPOT behavioral criteria.

The GAO reported in some depth on the SPOT program (GAO, 2010), noting that

> . . . approximately 14,000 passengers were referred to law enforcement officers under SPOT from May 2004 through August 2008. Of these passengers, 1,083 were arrested for various reasons, including being illegal aliens (39 percent), having outstanding warrants (19 percent), and possessing fraudulent documents (15 percent). The remaining 27 percent were arrested for other reasons such as intoxication, unruly behavior, theft, domestic violence, and possession of prohibited items. As noted in our May 2010 report, SPOT officials told us that it is not known if the SPOT program has ever resulted in the arrest of anyone who is a terrorist, or who was planning to engage in terrorist-related activity. More recent TSA data covering the period from November 1, 2010, to April 18, 2012, indicates that SPOT referred 60,717 passengers for additional screening, which resulted in 3,803 referrals to law enforcement officers and 353 arrests. Of these 353 arrests, 23 percent were related to immigration status, 23 percent were drug-related, 9 percent were related to fraudulent documents,

---

[*]   We benefited from a lengthy interview with Mr. Larry Willis of DHS regarding these experiments and related issues.

22 percent were related to outstanding warrants, and 28 percent were for other offenses.[*]

A later report, GAO (2012a), which includes testimony from DHS's Stephen Lord, refers (footnote 16) to a validation study that is not generally available; the footnote states:

> See DHS, SPOT Referral Report Validation Study Final Report Volume I: Technical Report (Washington, D.C.: Apr. 5, 2011). DHS's study defines high-risk passengers as travelers who knowingly and intentionally try to defeat the security process, including those carrying serious prohibited items, such as weapons; illegal items, such as drugs; or fraudulent documents, or those who were ultimately arrested by law enforcement.

A good many details about SPOT effectiveness were deleted from the public document because it is deemed sensitive security information by TSA (see subsequent footnotes).

The SPOT program has been bitterly criticized by numerous critics, and its deployment prior to scientific validity testing has been criticized by the GAO. The desire for prior testing is readily understandable given the high costs of deployment and the potential for the SPOT system be either useless or worse if it infringes on privacy, dignity, or civil liberties. The other perspective, however, is illustrated by testimony of Paul Ekman, on whose work some of the SPOT program's indicators are based (Ekman, 2011):

---

[*] Critics of the SPOT program often focus on the finding that it has not resulted in the arrest of a terrorist or would-be terrorist, even though it is reflecting the extremely low incidence of terrorists attempting to get through the screening rather than a shortcoming of the program. The program, after all, seeks to detect behaviors that suggest the need for further scrutiny, not detecting terrorists per se. That the people detected in trials were deceptive for reasons other than terrorism was inevitable. Whether DHS/TSA should be using methods like the SPOT program given the low base rate of terrorists and the potential of such programs for abuse such as illegal profiling is beyond the scope of this report and a matter of strong dispute. An apparent example of racial profiling by TSA officers at Logan Airport is described in news accounts (Schmidt and Lichtblau, 2012).

I have also been told by critics of SPOT that TSA should have first done observational research in airports, and the type of experimental check-point study carried out by Mark Frank and colleagues at Buffalo. . . . That would be a great plan if Al Qaeda and associates agreed to a three year vacation, during which the American people would not need the layer of security provided by SPOT.

TSA was not groping in the dark when it initiated SPOT. It reached out for the best evidence available that would allow them to introduce this layer of security without delay. They came to me and my colleagues, based on their perusal of the scientific literature; I did not reach out to them to sell them anything. We were able to provide relevant information because our research showed that hot spots are useful clues that are not lie-specific but are present in all high stake lies when there is a threat of severe punishment. And finally, keep in mind that these behaviors do NOT trigger an arrest. They trigger a conversation, usually around 30–90 seconds in length, during which the Behavior Detection Officers attempt to ascertain why this individual showed the behaviors they did. At times they uncover malfeasance, at times they find an innocent reason, at times they find a stressful but not illegal reason (e.g., a philandering traveler sneaking off to cheat on their husband or wife).

We do not know whether the field-test trials of SPOT have been peer reviewed and published in a formal and objective process (even on a classified basis), but the work was done by the American Institutes for Research. Nonetheless, the lack of transparency has undercut confidence and provided fodder for critics, including those in Congress. It is not a subject that we can resolve in this report. We do note, however, that first-class peer-review and publication processes within an appropriately cleared classified community have many precedents over the years. Indeed, this is common with some federally funded research and development centers (FFRDCs) and national laboratories. Also, studies by the National Academy of Sciences, the Defense Science Board, JASON, and others are often done on a classified basis or have classified components. Unfortunately, their very existence is sometimes not

acknowledged in public-domain discussions, although it certainly can be if the government chooses to do so, as may occur naturally or as the result of congressional testimony.

Project Hostile Intent (PHI) and Future Attribute Screening Technology (FAST)

Another DHS effort, Project Hostile Intent (PHI), looked at multiple additional indicators, including micro-expressions as studied by Ekman and colleagues (King, 2007). The goal of PHI was to develop remote, non-invasive, automated sensor technology capable of detecting deception in individuals with hostile intentions in real-time assessments; to understand the associations between behavioral cues, deception, and hostile intent; to conduct experiment to capture participants' behavioral and physiological responses to security questioning during a high-stakes deception scenario, in which participants were required to lie about their intentions regarding a future action; to extend research in the field of deception detection by identifying cues most strongly associated with deception by an individual with the intention to commit a future hostile action; and to conduct a series of analyses to develop a baseline understanding of the relationship between behavior (facial expressions and body movements) and deception during a hostile intent scenario.[*]

PHI later became the FAST program described in Burns (2008) and Middleton (2011). The program seeks to improve the screening process at transportation and other critical checkpoints by developing physiological and behavior-based screening techniques that will provide additional indicators to screeners to enable them to make more informed decisions. FAST is not intended to provide probable cause for law enforcement processes, nor would the technology replace or pre-empt the decisions of human screeners. Research on FAST is performed at laboratories such as Batelle, Draper Labs, and the Naval Research Laboratory.

---

[*]   Additional research addressed privacy issues of some of the DHS experiments themselves. A brief discussion is given in Willis, 2008.

The overall goal of the FAST project is to determine whether technology can help identify and interpret cues or signatures without the need for operator-induced stimuli. If successful, that would allow security personnel to remotely diagnose intent to cause harm. It would eliminate (or at least reduce) the need for face-to-face interaction and reduce costs.

At the outset, the program sought to evaluate five sensors (Burns, 2008):

1. A remote cardiovascular and respiratory sensor to measure heart rate and respiration, which allows for the calculation of heart rate, heart rate variability, respiration rate, and respiratory sinus arrhythmia.
2. A remote eye tracker, which is a device that uses a camera and processing software to track the position and gaze of the eyes (and, in some instances, the entire head) of a subject. Most eye trackers will also provide a measurement of the pupil diameter.
3. Thermal cameras that provide detailed information on the changes in the thermal properties of the skin in the face will help assess electrodermal activity and measure respiration and eye movements.
4. A high-resolution video camera that allows for highly detailed images of the face and body to be taken so that image analysis can determine facial features and expressions and body movements, and an audio system for analyzing human voice for pitch change.
5. Other sensor types such as pheromones detectors.

Experiments have been conducted, but we did not find results in the public domain. Numerous articles express extreme skepticism about the very feasibility of the FAST objectives. Popular-media and blog-level articles often deride the program as being an ill-advised attempt at mind reading.

We mention the FAST program here because it is a significant effort that includes testing of remote measurements.

Assessing the Kinetic and Physiological Approaches

The strengths and weaknesses of various detection approaches were discussed earlier, but here we elaborate on the shortcomings and limitations of physiological indicators. The primary problems with using physiological state or physiological changes as pre-incident indicators of terrorist or insurgent attack are twofold (see also Chapter Four): (1) The physiological states and reactions can be due to a variety of reasons, most of them benign, and (2) the states and reactions associated with attack and deception differ markedly across individuals and settings. Both imply the likelihood of a high rate of false alarms.

The nonspecificity of physiological states is due to the states being associated with general system demands and emotional states, such as perceived demand (Blascovich, Mendes, Hunter, and Lickel, 2000), fear of negative outcomes, or even physiological exhaustion—rather than something as specific as deception or intent to detonate a suicide belt or engage in a mass shooting. This is partially due to the fact that individual differences in developmental pathways mean that similar physiological end states can result from diverse upstream processes (Cacioppo, Uchino, and Berntson, 1994; Brulliard and Hussain, 2011). In practical terms, this means that the false-alarm rate would be quite high if relying on such indicators for detection.

Another aspect of this is that most measurable physiological parameters lie far downstream of the central nervous system and are distant echoes of the complex cognitive and motivational states within the brain. Higher-quality lie detector tests involve sophisticated protocols with multiple lines of questioning linked to multiple physiological parameters to better distinguish deceptive intent from general nervousness (or individual-specific oddities) (Grubin and Madsen, 2005). Polygraphs, however, cannot be applied to a large population, but rather only to a small number of people selected for special observation and questioning. Further, their most dramatic value (e.g., detecting knowledge of a plot or deception regarding a specific target) requires using stimuli that depend on the tester having some knowledge of that plot or target.

The other major challenge is that different individuals show very different physiological responses when attempting, e.g., to deceive or

attack. Individual differences or response patterns in a physiological state may reflect multiple different emotional states, motivations, or impending actions. In short, everyone is "wired differently," with different thresholds for physiological response as well as different baseline physiological states (Brulliard and Hussain, 2011). As a result, one man's sweaty palms and nervous lip-biting before an attack are another man's cold stare and blissful smile. To elaborate, literature reviews of pre-incident indicators (Smith, Damphousse, and Robert, 2006) and expert opinions (interview with Naval Postgraduate School professor Nadav Morag, 2012) report two "opposite" indicators of impending attack: nervousness, sweating, etc. (indicators of extreme sympathetic nervous system activation), and intense calm, dissociation, and flat affect (evidence of extreme parasympathetic nervous system activation) (Speckhard, Jacuch, and Vanrompay, 2012).

An interesting corollary is that an individual's physiological signals are likely to be much better cues of attack *timing* than anything else—but only when individual potential attackers have already been identified and baselined. In that case, moment-to-moment shifts in blood pressure, heart rate, and other parameters could be helpful in determining the exact moments when insurgents or terrorists are about to commence their violent attack (Scarpa and Raine, 1997; Raine, 1996). Such indicators of attack timing may come too late for effective intervention, and they also assume precise real-time measurement and interpretation, which is not yet feasible. While laser-Doppler technology and other related approaches hold promise for future applications (Bornstein et al., 2010), end-organ (blood pressure, sweating, etc.) and even EEG-based measurements are still quite sensitive to movement artifact and typically require the direct application of sensors for accurate and reliable measurement. This makes such technological tools potentially useful in interrogation or intensive questioning environments but less useful where standoff distance or early and clandestine detection is paramount.

A Recent Review Across Physiological Indicators. To end our assessment, we draw on recent research (on lie-catching more generally, not just from polygraph testing). This discussed the need for combining information but deplored the lack of adequate scientific experiments

and noted candidly that expert human skills must be part of the lie-catching method (Porter and ten Brinke, 2010):

> In observing the absence of a single cue or behavioural channel that consistently reveals deception, a holistic approach with concurrent attention to multiple channels of a target's behaviour (ideally videotaped for review) and changes from baseline behaviour is recommended whenever possible. Among the best-validated cues to be considered together include: illustrators, blink and pause rate, speech rate, vague descriptions, repeated details, contextual embedding, reproduction of conversations, and emotional 'leakage' in the face. While advocating a reliance on empirical evidence, we observe that few studies of high-stakes deception yet have been conducted. Further, some manifestations of lying are highly idiosyncratic and difficult to address in quantitative research, pointing to the need for keen observation skills, and psychological insight. A recurring theme is the need for the field to devise innovative approaches for studying high-stakes lies to promote ecological validity. Ultimately, such work will provide a strong foundation for the responsible application of deception research in forensic and security settings.

Critiquing a Key Underlying Assumption and Setting New Goals. One of the most important conclusions to which many scientists have been coming in recent years is that the assumption that hostile intent betrays itself in stress is highly questionable. Indeed, a recent study argued that it is not worthwhile for researchers to look further into detecting malign intent by, e.g., just observing subjects as they go by, whether by facial expression or otherwise. Instead, it is argued, much greater emphasis should be placed on finding ways to stimulate observable reactions by increasing "cognitive burden," rather than stress (Vrij and Granhag, 2012). This line of argument is based on substantial research discussed in the article, which supports the idea that lying is "more difficult" than telling the truth. Lying may come easily in a sense, but not if it requires putting together a convincing narrative and answering questions phrased in ways that require "thinking." This stream of recommendations emphasizes the crucial role of the ques-

tioner and his or her skill, which is also consistent with the experience of many practitioners in the long-time use of polygraphs. In a partial rebuttal, Frank and Svetieva (2012) support the importance of the questioner and questions but are less sanguine about the effectiveness of methods such as asking questions that require "thinking backward." Also, the authors strongly dispute the relevance of many "laboratory" assessments of emotion detection because, they argue, those studies are typically conducted with students or other people who are merely asked to pretend to be terrorists. That, it is argued, is quite different from the context of real terrorists in very high-stakes circumstances. Frank and Svetieva conclude with the observation that "whether we researchers like it or not, real world law enforcement or intelligence interviews feature highly elevated emotional and cognitive loads."

Informal Criticisms That Cannot Easily Be Assessed. One of the many frustrating aspects of our literature review was that many of the continuing disputes are not pulled together so that convergence is possible on both points of agreement and points of disagreement. Authors write in their various favored journals but seldom in a forum expecting clarification, iteration, and convergence on points of agreement and disagreement. This, arguably, is just a standard problem in science that is not solved by normal peer review. Resolutions, when they are achieved, often come about with studies of the National Academy of Sciences or independent analysis by organizations such as federally funded research and development centers (FFRDCs).

To illustrate, a well-written *Nature* article (Weinberger, 2010) summarizes some of the many disputes about the ability to detect "intent to deceive" but does so in a journalistic manner of contrasting what various people say, not necessarily with peer-reviewed publications to support their statements. One of the important studies referred to, by DoD's JASON organization, is not published (a summary was apparently provided to *Nature* as background for the article). Several skeptics of the work of Paul Ekman are quoted as being unable to reproduce his work, but the failure-to-reproduce studies have apparently not been published. Further, some of the articles quoted suggest to us that part of the dispute stems from potentially irresolvable matters of perspective. For example, the author of the JASON study

is quoted (in Weinberger, 2010) as saying "The scientific community thinks that it is extremely important to go through the process of scientific verification, before rolling something out as a practice that people trust." That statement initially seems reasonable, but it seems to assume that there is an accepted and feasible process for scientific verification, which appears not to be the case. It also seems to ignore the fact that operators who favor the tools in question think of them as tools to be used as part of more complex processes, not as something that can be trusted on their own.

Table 6.3 is our assessment of how the approaches discussed in this chapter stand in terms of maturity, potential, measurability, and vulnerability to countermeasures.

Distinctions Helpful Amidst Controversy

As discussed earlier, many of the behavioral-indicator methods are highly controversial.

Table 6.4 is another depiction comparing approaches, this one intended to draw some contrasts and highlight distinctions that are not always clear in the often-emotional debates about methods.

The second column of Table 6.4 reminds us that if subjects are worried about security officials watching for behavioral indicators, they may be deterred in some respects. They may go more out of their way to avoid the observation they fear; they may be more stressed when questioned; and, in some cases, they may be more inclined to cooperate or even tell the truth when questioned. There is a long history of subjects being more worried about their deceptions being detected than laboratory studies suggest would be objectively warranted given the methods' imperfections. The remaining columns of Table 6.4 show distinctions in value as a function of defense-system intent. If the purpose is merely to indicate that an individual merits a somewhat closer and not very troublesome look than the average person being screened, then the methods all have potential or actual value, in some cases even with automated methods. Using the methods to justify more vigorous and extended checking, undertaken with some prejudice and perhaps

Table 6.3
Detecting Hostility or Deception from Movement Physiology and Movement

Domain	Status	Upside Potential	Measurement Requirements	Shortcomings and Vulnerabilities
Kinetics and gross motor movement	Indicators have been validated for human observation and automated analysis in laboratory and experimental settings, including some operational settings (e.g., for gait of individuals carrying weighted objects).	Gross motor movements may reveal action, intent, or deception. On-foot motions may be unavoidable in such proximal security settings as checkpoints. Gross motor movement may be passively observed, but also actively elicited.	Some security contexts may not allow for sufficient physical movement to be interpretable (e.g., interrogation).	Masking with deceptive movements. Sensitivity to context and individual differences. Nonspecificity: triggering by diverse emotions and motivations.
Physiological state and reactions	Indicators have been validated in laboratory and experimental settings, with some experimental paradigms simulating elements of counterterrorism and some (facial) cutting across culture. In some cases (e.g., voice stress and facial indicators), automated recognition shows potential but currently has high error rates.	Internal physiological reactions are relatively automatic and difficult to control (e.g., micro tremors in speech or micro facial expressions). Probing of various sorts (even seemingly random conversations) can trigger reactions. Certain elements of facial expression are very difficult to alter voluntarily, including micro-expressions.	Currently, measurement requires direct application of sensors or the physical observation of facial flushing, sweating, etc. Some (e.g., facial) require lighting and proximity with currently painstaking coding feasible only for high-value interrogations. Success requires exceptional "natural" talent or training, but limited available data suggests training is effective. Measurements are most valuable when comparing against an individual's baseline, which is only feasible in voluntary monitoring or interrogation context.	Differences across contexts and individuals. Nonspecificity. Influence of drugs and training (e.g., to dampen or obscure differences between baseline and signals). Masking, in some cases (e.g., sunglasses or plastic surgery) Some differences exist (perhaps not critical) across culture. Masking (e.g., sunglasses, plastic surgery, or Botox for facial), but this may also be an indicator.

Table 6.4
Some Comparisons of Behavioral Methods

Method	Deterrence or Cost Imposition	Flagging for Further Routine Screening		Flagging with Prejudice for Extended Checking and Detention	Tool in Interrogation	Basis for Arrest or Conviction
		Automatic	Human			
Polygraph	Yes	No	Yes	Maybe	Yes, but	No
Voice stress analysis	Yes	Yes	Yes	No	Yes, but	No
Facial expression	Yes	Technology not well developed	Yes	No	Yes, but	No
EEG	Yes	Technology not developed	Yes	Maybe	Yes, but	No
Text or speech content	Yes	Maybe	Yes	Maybe	Yes, but	Maybe
Gait analysis	Yes	Yes	Yes	Maybe	No	No

including significant delay and detention, is quite another matter. Here, *none* of the methods are generally valid (i.e., always valid). However, in some cases they may help trained operators spot someone worthy of careful attention even though some false alarms will occur. For example, someone displaying high vocal stress or marked use of language correlated with deception is a natural candidate for further inquiry, even though the causes may be benign. There are analogues from the long history of experience with polygraphs. For example, an entirely innocent person may fail a polygraph test because of idiosyncratic fear of the test or suppressed guilt about unrelated and even trivial matters. Because of this, it has proven extremely important that operators be well trained and professional and that the testing process follow protocols minimizing inappropriate anxiety and unreasonably intrusive questioning.

We conclude, then, that *in some cases* (but not "on average") the methods may be valuable for flagging suspicious individuals. With respect to the small subset of individuals who might be interrogated in more detail with the intent of extracting confessions, there is solid evidence about the methods' value *as part of a more general process that includes skilled officers*. At the same time, there is also clear evidence of false confessions and other significant side effects. The "Yes, but . . ." in Table 6.4 is meant to indicate that, yes, there is value, but great care must be taken and, historically, there have been serious lapses.

The last column in Figure 6.4 is crucial: We see no basis in the scientific literature for using evidence from most of the methods as the basis for arrest or conviction. There are good reasons for existing legal constraints on this matter. An exception is that the *content* of text and speech (including threatening communications) is sometimes significant evidence, and is allowable in criminal trials. The primary point here is that the behavioral indicators are or have the potential to be useful indicators for identifying people meriting a closer look, but they are not robust enough for anything more.

Even this statement would be resisted by some because the objective evidence of the methods' independent validity remains controversial and less than scientifically persuasive, and because claiming that the methods are valuable when used as a tool by a trained professional as part of a larger process is inherently dissatisfying scientifically. This is a domain in which there are enduring tensions analogous to those that allow police officers to take some measures based on subtle indicators, and even intuition, but not to take other measures with anything less than "probable cause." We cannot resolve such issues in this study.

General Observations. Regardless of the sensor used or specific behavior being measured, a common set of issues pervades the use of detectable behavioral indicators to predict terrorist or insurgent attacks. These issues directly bear on the advisability of investment in detection technologies and tools.

One crucial axis of variability across indicators and detection measures is the degree to which they are applicable across attacks. Some indicators (and thereby their associated detection measures) are common to a wide variety of attacks (for example, a high percentage

of attackers approach their targets on foot, at least in the final stages, making gait analysis applicable to many attack types). Another crucial axis of variability, covered in greater detail in Chapter Seven, is whether indicators can be detected through clandestine, passive measurement or whether they require some sort of active intervention (for example, a "random intercept" interview while waiting at screening checkpoints). Indicators and associated detection methods also vary in the degree to which they are detectable at standoff distances or while surrounded by various sorts of "noise" (such as the presence of other individuals). As mentioned earlier in the chapter, indicators also vary in their discriminant validity—some indicators indicate emotional or motivational states (such as anger or deception) that are common during attacks or attack planning but also common to a wide variety of other behaviors. Indicators also vary in their context-specificity and cultural variability. Very few indicators provide good predictive capacity for attacks on their own. Thus, information fusion (as discussed in the latter part of Chapter Seven) becomes important. How much it can accomplish depends on the false alarm rates of the separate indicators, and the techniques used for fusion, which may include subjective judgments.

To render a definitive assessment of which sets of behavioral indicators and associated sensor systems are most worthy of further investment, it would be advisable to develop a formal set of metrics with which to assess currently available (and emerging or "under development") indicator-detection technologies with respect to the criteria described above. This formal assessment, however, is beyond the scope of this report.[*]

---

[*] Other methods not much discussed in this report because they are not behavioral include bomb-detection systems that work at a safe distance, such as standoff non-imaging radar, and provide method and automatic detection systems to protect high-value assets, such as buildings, airports, and checkpoints, from personnel-borne IED threats.

Cross-Cutting Issues

Introduction

This chapter is about cross-cutting issues that loomed large as we sought to make sense of our review of the science base. To our knowledge, the literature contains no agreed framework for thinking about this, so we constructed a first-cut framework ourselves. We characterize a system in terms of its detection effectiveness and its effectiveness in limiting false alarms and their consequences. We then see these as depending on seven characteristics of the system, as shown in Figure 7.1. We discuss them in the numerical order shown in the figure, which is not strictly left to right so as to keep the figure clean (no line-crossing). The characteristics are (1) layering; (2) the sensitivity and selectivity of subsystems in the absence of behavioral stimulation or countermeasures; (3) behavioral mechanisms for stimulating responses; (4) countermeasure resistance; (5) observation distance, covertness, and automaticity; (6) information fusion; and (7) mitigation of false-alarm consequences when they occur.

Given such system characteristics, it is possible to discuss alternative systems with a depiction such as the radar/spider plot in Figure 7.2, which compares four notional systems represented by the contours. The notional systems are labeled as primitive, baseline, advanced, and very advanced. The dimensions representing the characteristics from Figure 7.1 appear as axes or rays. Each is scaled from 0 to 100, allowing a subjective measure of how well the system does with regard to the particular dimension, in comparison with what might be "optimal" in terms of effectiveness and cost (which are affected, of course, by

Figure 7.1
Factors Affecting Overall System Effectiveness
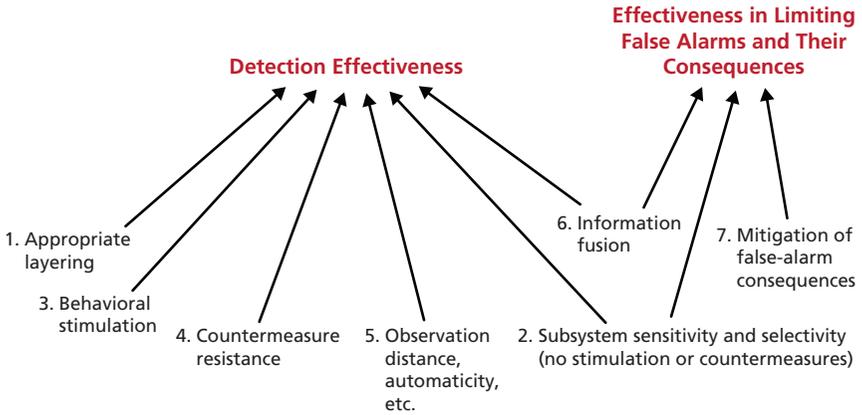
Figure 7.2
A Notional Framework for Characterizing an Overall System
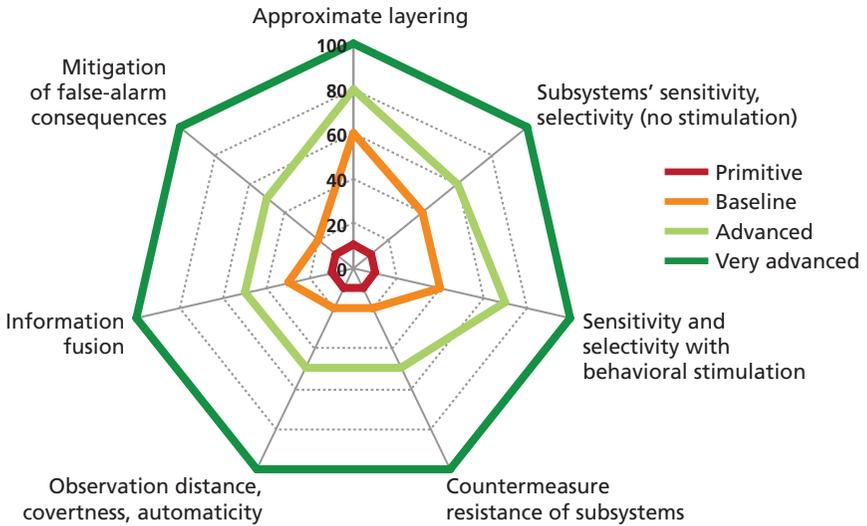
the projected state of technology, feasibility, etc.). Again, this is highly notional (the scales could be defined better in future work), but the qualitative framework provides a useful way to structure discussion. Merely for the sake of illustration, Figure 7.2 assumes that the baseline (say today's system) is more advanced in some respects than others, with layering having been taken seriously for some time, but with information fusion, for example, being relatively primitive[*] and too little work having been done to minimize consequences of false alarms when they occur. Progress would correspond to systems with contours farther and farther toward the extremity of the radar/spider plot.[†]

The following sections discuss all of the dimensions above, but, before proceeding, we note again that our study focused primarily on the science of whether possibilities for better detection methods exist, without prejudice as to whether they should be acceptable in American society. Some of the detection methods raise important issues of privacy, ethics, and law, as discussed in a recent National Academy of Sciences study of privacy and civil liberties (Perry and Vest, 2008). Although not analyzing such issues, much less the difficult tradeoffs, we comment on some of the issues in the course of the chapter, particularly in regard to reducing false alarms and mitigating their consequences when they occur (which they assuredly will).

Appropriate Layering

Layering and Screening

Since no single foolproof detector is plausible, good security-system designs exploit the potential leverage of layering, i.e., of using a sequence

---

[*]   We were struck by how little information is available to airport security, for example. The information certainly may include basic passport data, whether the person is on a watch list, and something about the person's recent travel history, but that is a small fraction of the information that *could* be available if that were desirable.

[†]   The area within a contour is not a sound measure of the option's overall effectiveness because not all of the dimensions are necessarily of equal importance and their value need not add linearly. The figure, however, is sufficient for our purposes in indicating that overall progress involves progress along all of the dimensions shown.

of detection measures. Mathematically, the leverage can sometimes be dramatic. Suppose that an attacker must penetrate three independent layers, each of which has only a one-third chance of detecting him. With three layers, the defense has about a 70 percent likelihood of detecting him. The same benefits can be achieved by applying multiple detector methods at a given point in time, turning a single layer into the equivalent of multiple layers.

The mathematics of layered defenses has been developed in the past for ballistic missile defense (Wilkening, 1999), defense in-depth generally, and even cyber defense against worms (Albanese, Wiacek, Salter, and Six, 2004). One treatment deals with countermeasures, game-theoretic considerations, and common-mode failures (Willis, Bonomo, Davis, and Hillestad, 2006). Another technically rich discussion includes subtleties of inter-layer correlations (LaTourette, 2012). For many reasons, such as those discussed in the article, how much layering is enough (and when additional layering is even counterproductive) depends on many considerations.

If the first layer cannot detect the intruder but can distinguish between high- and low-risk individuals (initial screening), then those classified as high-risk can be required to go through additional checks (secondary screening). Or, more typically, everyone will go through additional checkpoints, but those in the high-risk category will be exposed to more scrutiny, which takes time and resources and may also delay and otherwise inconvenience those affected.

The problems that arise in such approaches include (1) low detection probabilities, (2) high false-alarm probabilities with numerous negative consequences, (3) correlated probabilities, and (4) countermeasures (which affect all of the previous three). Detection probabilities based on individual behavioral cues are typically modest. Further, many people may exhibit similar behaviors even though they have no malign intent, thereby creating false alarms. As noted, the effectiveness of successive detection layers is often correlated, as when lax management affects multiple parts of a system, or when—because of weariness or anxiety—an innocent individual triggers multiple behavioral indicators. Finally, of course, a would-be attacker will try to "beat the system" by avoiding defenses or by employing countermeasures such

as by training to avoid behavioral cues. As discussed in LaTourette (2012), interactions among non-independent layers can either reinforce or degrade performance, including through deterrent effects and so-called shirking effects.

In our study, layering comes in when, for example, behavioral indicators are used for screening, to flag individuals meriting further scrutiny.

If behavioral indicators are used to classify people into those who are and are not regarded as representing more than normal risk, and should therefore be subject to further scrutiny, this can be referred to as negative screening—i.e., screening to identify people of concern because they do not "pass" all measures of "normalcy." Screening and the related issue of profiling[*] continue to be quite controversial, as discussed by the GAO (2010), Congressional Research Service studies (Elias, 2009, 2011), and news media.

Screening/profiling based on or seemingly based on national origin, age, or apparent racial or ethnic grouping has generated some of the most heated debates (for a serious popular-level debate with informative discussion, see Harris and Schneier, 2012). Using a purely mathematical approach to analyze the advisability of racial or ethnic profiling, Press concludes that weak profiling (rather than what he calls democratic screening, which is when everyone is screened) is optimal (Press, 2009).[†] It is also is important to consider the secondary effects of such screening processes, as well as the possibility of malign actors thwarting simplistic screening procedures through the recruitment of operatives not displaying the screened characteristics (e.g., lack of an overtly Muslim or Middle Eastern appearance and use of female or child operatives).

---

[*] The terms "screening" and "profiling" are not generally differentiated, but in some contexts "screening" is more descriptive and objective, while "profiling" *infers* characteristics, i.e., is more extrapolative. Sometimes, "profiling" refers to making decisions, such as about whether to interrogate, based on racial or ethnic characteristics, which is illegal in many jurisdictions. Other times, "profiling" has no such negative meaning.

[†] Press defines "strong profiling" as screening in which the probability of being selected for secondary screening is at least proportional to a "prior" (the prior probability). A weaker version selects for secondary screening in proportion to the square root of the prior probability.

Screening that discriminates on the basis of protected characteristics is illegal, but using behavioral cues is acceptable unless it has disproportionate impact on the protected categories. Ongoing research includes empirically based research in both laboratory and field settings, as well as research and analysis based on mathematical models and simulations. Those below are intended to be illustrative.

A Different Kind of Screening: The Trusted Traveler Concept

Although it is outside the scope of this study, we should at least mention a different kind of screening, the purpose of which is to identify people who can be excluded from some or all subsequent checking. This is the idea behind the "Trusted Traveler" concept, insights about which were published early in the previous decade (Shaver and Kennedy, 2004; Robinson, Lake, and Seghetti, 2005). Robert Poole championed what he called a "risk-based" approach to screening (Poole and Passatino, 2003; Poole, 2009). The subject was reviewed analytically (Jackson, Chan, and LaTourrette, 2011), assessing the value of using background checks to sort individuals into high- and low-risk categories for differential attention at checkpoints. The authors use receiver operator characteristics (ROC) curves and relatively simple mathematical models to lay out the different conditions under which screening into low- and high-risk groups would yield increased or decreased detection rates. The conditions at issue involve the false-positive and false-negative rates of primary screening, the increased time that is "freed up" and available for secondary screening, the baseline detection rate of secondary screening, and the base rate of attempted terrorist attacks. Together, these determine whether a two-stage screening process yields increased rates of detection. If attackers do not try to get into the trusted-traveler program (and measures can arguably be taken to help deter them from doing so), then the increased time and resources available to scrutinize those not in the program would presumably raise detection rates significantly. Aspects of the trusted traveler concept are now being implemented by TSA.[*]

---

[*] Some additional recent papers are worth citing: Morosan, 2012; McLay, Jacobson, and Kozba, 2006; Jacobson, 2012; Stewart and Mueller, 2012.

Sensitivity and Selectivity

Basic Concepts and Terms

In discussing the effectiveness of a detection system and its subcomponents, it is necessary to define some technical terms. Table 7.1 uses classic terminology to show the four cases that apply when a given individual is tested by a detector that reports either positive or negative (rather than, e.g, "maybe," as discussed in the section on information fusion). The results, then, are referred to as True Positive (TP), False Positive (FP), False Negative (FN), and True Negative (TN). The intent is to minimize false negatives (failure to detect) *and* minimize false positives (false alarms). These errors are commonly referred to as Type I and Type II errors, respectively.

Less-than-perfect detection systems show a tradeoff between sensitivity and false alarm rate. That is, it is usually possible to increase the system's sensitivity (minimizing Type I errors), but only at the expense of raising the rate of false alarms (raising Type II errors). For example, if an initial screening tries to distinguish between high-risk and low-risk individuals, it is certainly possible to toughen the screening criteria so that more people will be regarded as high-risk, but then more innocent people will then be subjected to the subsequent scrutiny. Further, the false-alarm rate may rise dramatically and overwhelm the system. This is especially so when the "base rate" is low, i.e., when the fraction

Table 7.1
Classic Matrix of Detection Outcomes

		Testee's Actual Character	
		Positive	**Negative**
Detector's Verdict	**Positive**	True Positive (TP)	False Positive (FP) (Type II Error)
	Negative	False Negative (FN) (Type I Error)	True Negative (TN)

of people being screened who have malign violent intent is quite low (e.g., 1 in a million rather than 1 in 3).

Several terminologies are used to address the same basic concepts. Unfortunately, authors do not always use the words to mean the same thing. In this report, we use "sensitivity" to mean the detection rate, the probability that a deceptive subject will be detected. "Specificity" is the probability that an innocent subject will be classified as innocent.[*] These are complementary measures. They measure how probable it is that someone deceptive and someone innocent will be classified as such, respectively. Taken together, they provide a good measure of simple detection-system capability, although there are many further subtleties.

Another measure that is often used is "accuracy," which in simple cases, such as that depicted in Table 7.1, is the fraction of calls (positive or negative) that are true. The term "reliability" is sometimes used, but inconsistently. It is best used as a measure of whether results are consistent across trials. The false-positive rate is the complement of specificity, i.e., 1 – Specificity. The false-negative rate is the complement of sensitivity, i.e., 1 – Sensitivity. These are just the fractions of positive and negative calls that are true, respectively (National Research Council, 2003, p. 122). Table 7.2 summarizes some of the relevant expressions mathematically.

The tradeoff between sensitivity and false-alarm rate has been studied for years, often in papers referring to signal detection theory (SDT) or ROC-curve. The mathematics is relatively straightforward, but not intuitive (Fawcett, 2006). Figure 7.3 is one example of a so-called ROC-curve, adapted from Appendix H of National Research Council (2003). If the sensitivity of the system (y-axis) is 0.8, so that it detects 80 percent of attackers, and if the test applied has a system accuracy called the accuracy index A of 0.9 overall (see page 44 of the NRC study), then both the false-positive and false-negative rates are about 0.2 (20 percent) (actually, 0.17 and 0.2). What is not so clear from this, however, is that if the attackers are only a very small fraction of those being screened (e.g., 1 in a thousand or million), then almost

---

[*]   See also National Research Council, 2003, Chapter 4.

Table 7.2
Mathematical Expressions

Variable	Equation
Sensitivity	$$\text{Sensitivity} = \frac{TP}{TP+FN}$$
Specificity	$$\text{Sensitivity} = \frac{TN}{FP+TN}$$
Accuracy	$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$
False positive and negative rates	$$\text{False positive rate} = \frac{FP}{TP+FP}$$ $$\text{False negative rate} = \frac{FN}{TP+FN}$$

all positives will be false positives. We dramatize this in Figure 7.4, which shows the false-alarm index (number of false alarms per alarm) as a function of the base rate (fraction of true positives in population) and detection-system accuracy. The takeaway is that for populations of interest, the false alarms can readily dominate the system. This means that initial screening could be useful if it increased the base rate of threats were something more like 0.1 (1 in 10), rather than 1 in a thousand or million. As the table at the right of Figure 7.4 shows, the false-positive index is still high (0.2–0.6, depending on system accuracy), but drastically better than with even lower base rates.

The Role of Context. Context matters when pondering how to trade off detection rate and false-alarm rate. As occurs in the immediate aftermath of a crime, it is possible temporarily to ratchet up the sensitivity of screening (e.g., at a security checkpoint on a road leading away from a terrorist incident) despite the price paid in increased false alarms. That price can be met with temporary additional resources (the additional security personnel focused on that area). Experience with

Figure 7.3
Illustrative Tradeoffs Among Sensitivity, Accuracy, and False Alarms
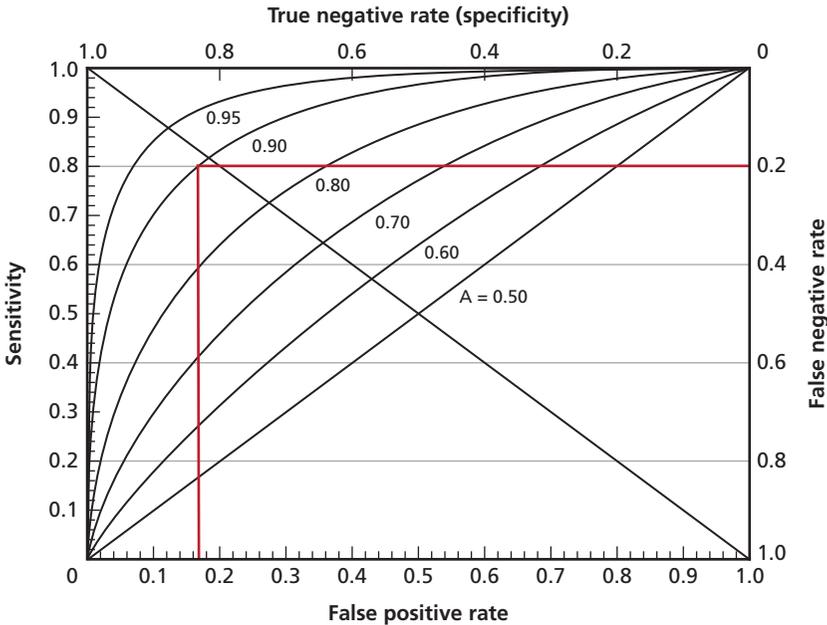
SOURCE: Adapted from National Research Council, 2003, Figure H.2, pp. 340ff.
NOTES: The curves are based on a number of illustrative assumptions. In particular,
the probabilities of positives with deceptive and nondeceptive groups are assumed
to be described by Gaussian distributions with identical means.
RAND *RR215-7.3*

polygraph work supports this concept: Data from polygraph testing
(and related interrogation and discussion) can significantly increase the
likelihood of detecting the guilty party and may be especially justified
when the guilty party is more likely to be among those tested (a higher-
than-normal base rate) (Honts and Schweinle, 2009).

By analogy, security personnel on the watch for potential attack-
ers could use less discriminant behavioral (and other) cues during a
period of high alert if they had additional resources. Furthermore,
during such a period, the public would probably also be more forgiving
of inconvenience.

Figure 7.4
False Positive Index Versus Base Rate and Accuracy

False-positive index (cell value)

	Accuracy		
Base rate	0.9	0.8	0.7
0.001	208	452	634
0.01	21	45	63
0.1	1.9	4.1	5.7
0.5	0.21	0.45	0.63
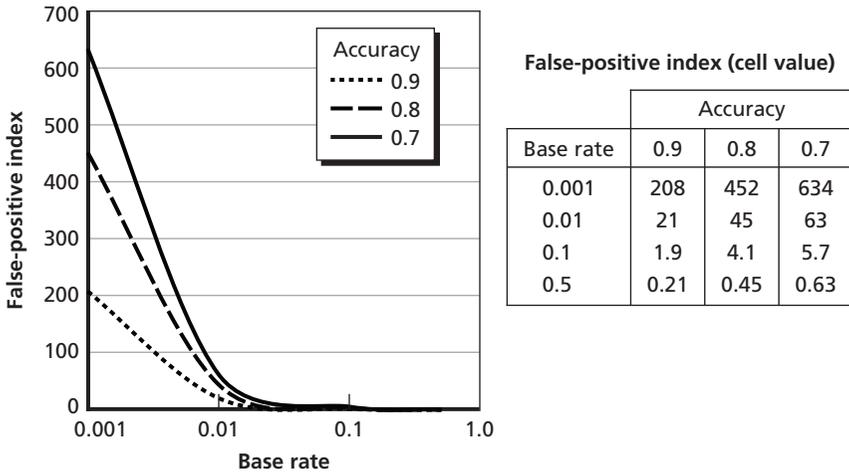
SOURCE: Calculated from data in Appendix I of National Research Council, 2003.
RAND *RR215-7.4*

Effectiveness of Screening Without Stimulation

Considerable disagreement exists about whether screening with the various behavioral indicators works when security officers observe subjects as they pass by or move through routine lines. As discussed in Chapter Six, TSA's SPOT program has been described by DHS as highly effective in trials, but that evaluation, although serious and ambitious, is not in the public domain. Many outside scientists are skeptical and have published discouraging experimental results, which have in turn have been criticized as unrealistic. We cannot resolve the matter here.

Improving Effectiveness with Behavioral Stimulation

One of our early conclusions when beginning our review was that we should be sure to distinguish between the value of behavioral indicators with and without stimulation. That hypothesis was corroborated by the literature. Although some things can be accomplished by observing unstimulated behavior (e.g., observing a gait that may suggest decep-

tion or observing general nervousness and an apparent desire to avoid checkpoints), stimulating behaviors appears much more valuable in other instances. It is a contrast, for example, to observations and even mild questions that allow subjects to proceed relatively comfortably and to perform as they have rehearsed. A recent review article expressed this as a general principle in using behavioral indicators, seeing it as key to progress (Vrij and Granhag, 2012). This view is consistent with long experience by Israel's airport security personnel. The point can be exaggerated, and the unstimulated behaviors can sometimes be valuable in themselves or in combination, as discussed earlier (see also Frank and Syetieva, 2012, which responds to Vrij and Granhag, 2012), but this section focuses on stimulation or probing.

General Considerations

Probing refers to the intentional stimulating of verbal or behavioral responses to assist in detection activities. Police forces and security organizations have much experience in such matters. The first issue becomes how much additional information is gained from such probing rather than passive observation of individuals' "natural behavior." The information gains then have to be weighed against (1) the resources expended (perhaps at the expense of more comprehensive passive observation, background checks, etc.); (2) such negative consequences as inconveniencing, insulting, or unnecessarily raising the anxiety of innocent people; and (3) the potential for the probing activities to reduce future cooperation or even cause some of those affected to be radicalized due to a sense of humiliation and unfairness. The third of these is especially salient when the probing is done by, say, an occupation force or the security forces of an authoritarian government.

Conceptually, behaviors might be stimulated by diverse means:

1. verbal questioning, or even interrogation
2. anxiety-raising changes of procedure or process (e.g., being questioned in an isolated room, being sent into a clearly different line)
3. subliminal stimuli
4. tests, such as with polygraph equipment.

Probing may be polite, intrusive or even "in-your-face," or unobtrusive. Examples of the latter are being studied. Probing, however, can also be quite intrusive. Much more complex probing is possible, such as directing people in unanticipated directions, to unusual lines, or to anxiety-inducing machinery. A reasonable hypothesis is that even well-trained attackers would be more likely to "lose their cool" in such circumstances, especially if the probing risks an operational delay. To our knowledge, however, this tack has not been pursued nearly as far as it might be, nor did we encounter public scientific literature on experiments to test such procedures.

Given efforts to stimulate, behaviors can be observed in several different ways. We highlight only a few, drawing from the literature.

Physiological Responses to Probing
Broad Considerations

As noted in Chapter Six, measurement of biological responses to probing seems to be most effective when closest to measuring actual brain activity. For example, using EEG measurements to measure the P300 amplitude wave to detect involvement in a simulated terrorist plot in response to known items associated with the plot (target location or individual, weapons used, etc.) has shown false positive rates as low as 5–10 percent and false negative rates as low as percent in experimental settings (Meixner and Rosenfeld, 2011). However, such methods require the cooperation or coercion of individuals and the use of expensive monitoring equipment.[*] Nonetheless, the successes shown in the Meixner and Rosenfeld work are a remarkable advance. Rosenfeld and students have also published a number of articles on countermeasures and countermeasure-resistant methods to P300 testing.

---

[*]   Furthermore, individuals with aggressive tendencies tend to show P300 spikes with reduced amplitude (Patrick, 2008). That is, while this detection method requires distinguishing between individuals based on P300 spike amplitude associated with specific stimuli, the very individuals who are most aggressive tend to show a tendency toward a smaller P300 spike in general, regardless of the context. This would probably reduce the effectiveness of a P300-based stimulus-response detection technique if applied in a large-scale population context (e.g., at airport checkpoints)—increasing false-alarm rates and, possibly, reducing valid detection rates.

Somewhat "further from the brain," researchers have also attempted to use the response to provocation of peripheral physiological signals (heart rate and heart rate variability, blood pressure, etc.) to detect hostile or deceptive intent (Aikins, Martin, and Morgan, 2010). However, peripheral physiological responses are notoriously nonspecific (Porges, 1995) and rife with individual differences in response profiles (Cacioppo et al., 1992, 1994). For example, the onset of aggressive behavior itself may be accompanied by either increased or decreased heart rate, depending on whether the aggressive act is reactive or instrumental (Scarpa and Raine, 1997). Instrumental aggression is common among those with psychopathic or sociopathic tendencies, as discussed in Chapter Four. See Dutton, 2003.

Verbal Probing and Human Observation

Another approach to probing is verbal. Verbal provocation and the human assessment of verbal and behavioral responses without the use of sophisticated or expensive biological monitoring equipment can be quite effective in some circumstances. Israeli airport officials have used such techniques for many years, apparently with great success. Israeli security authorities engage passengers designated as high-risk with numerous rounds of aggressive questioning or detainment (sometimes by different people, sometimes by an officer overseeing the process) to observe their verbal and behavioral responses.

Current research indicates that basic subjective assessment of the plausibility of reasons given for traveling, or being at a certain location, along with the consistency of stories over time across interviewers together provide the best clues about hostile or deceptive intent (Vrij, Granhag, Mann, and Leal, 2011a; Vrij, Leal, and Mann, 2011b; Vrij and Granhag, 2012). To elicit these stories (especially multiple iterations of the stories), passive observation is not enough. Potential attackers must be asked to provide reasons for their presence in a particular area, intent to travel, and so on. These provocations may produce both linguistic and other behavioral cues that can lead to detection and interdiction.

These are checks that non-expert personnel can perform. For example, Israeli bus drivers are trained in both behavioral profiling

and deliberate verbal probing in order to help detect potential suicide bombers boarding buses. There are quite a few documented cases in which this training and simple verbal provocation has either helped thwart or mitigate the damage from attempted attacks. The following two examples illustrate the value of such verbal probing:

> . . . the driver, 52-year-old Menashe Norial, followed the security-awareness training procedures and stopped the bus before the bus stop to have time to examine the waiting passenger's appearance and behavior. . . . He opened the door, and in accordance with security-awareness procedures, asked the passenger several questions about his destination. The passenger didn't respond and started to board the bus. Norial thought that he was a "weird" person, maybe on drugs, like some of the youth traveling to the festival. He asked again for the passenger's destination and asked what was in the bag, but the young man did not respond and took another step onto the bus. At that point, Norial became suspicious and decided to take action. He pulled the handbrake, turned, and stood up against the passenger, ready to tackle him. . . . Norial pushed the young man down the stairs and out of the bus, holding the bomber's hand to keep him from operating the switch. At the last step, Norial pushed the would-be bomber away from the bag and threw him down on the ground. . . . Police arrived at the scene, arrested the terrorist, and disarmed the explosive device. (Butterworth, Dolev, and Jenkins, 2012, pp. 10–11)

> . . . he [the driver] still exercised a high level of security awareness and decided to question the suicide bomber. He called to him and asked him for his destination. The suicide bomber responded, "To the hospital." The driver recognized an Arabic accent and the high level of tension in the bomber's voice. Increasingly suspicious, he told the terrorist to wait on a bench next to the Number 12 line bus stop. . . . The security guard started to walk toward the bomber and reported on his radio that he was approaching a potential suspect for a security check. He realized that the suspect was alternately standing up and sitting down. The security guards

acted almost simultaneously." (Butterworth, Dolev, and Jenkins, 2012, pp. 58–59)

In both cases, simple verbal probes led to detectable verbal and nonverbal behavior that disrupted or lessened the impact of ongoing attacks.

Dealing with Countermeasures and Adaptation

Much of the literature on detecting behavioral responses assumes that the targets of observation are behaving normally, but it is obviously essential to worry about countermeasures. Indeed, vulnerability to countermeasures should be a prime consideration in evaluating investment programs.

As noted in a recent study,

> It should not immediately be assumed that the newest and most advanced technologies—the highest wall, the most sensitive surveillance—will best protect society from terrorist attack. . . . It is only through fully exploring an adversary's counter technology behaviors that vulnerabilities in a nation's defenses can be discovered and the best choices made to protect the nation from the threat of terrorism. (Jackson et al., 2007, p. 23)

RAND has reviewed countermeasures taken by various groups to avoid detection—including Palestinian terrorist organizations, Jemaah Islamiyah, the Tamil Tigers, and the Irish Republican Army. Documented countermeasures (including adaptations) include changes in patterns, style, and media for communication, disguise, false documentation, new weapons innovation, switching target sites, modifying attack duration, monitoring of monitoring devices and personnel, relying on the capabilities of more advanced affiliate organizations, destruction of forensic evidence, and punishment of informants to decrease the effectiveness of human intelligence efforts (Jackson et al., 2005). In each of these cases, the state security apparatus responded with its own adaptations, too, including some that required technological development and refinement.

Not surprisingly, much information is readily available in the public domain on how to defeat security systems. Indeed, both individuals and groups seem eager to find and communicate such information, perhaps by analogy to communicating how to "jailbreak" a smart phone or, in earlier decades, how to make long-distance telephone calls for free (a famous prank of the young Steve Jobs). As one example, an ongoing project called "CV Dazzle" provides open-source information on how to thwart automatic facial detection algorithms (cvdazzle.com).

Countermeasures, of course, lead to iterative patterns of innovation and counterintelligence on both sides over time. Thus, investment in tools and technologies for detecting malign intent must be weighed against their vulnerability to countermeasures by malign actors. This is less straightforward than it might seem because even if countermeasures exist, they may not be used (or may not be used well). Worst-casing would eliminate measures that could be quite valuable. On the other hand, basing decisions on optimistic assumptions (absence of countermeasures) has the potential to waste enormous resources and divert attention from more promising methods.

Observation Distance, Covertness, and Automaticity

There is clear value in being able to observe behavioral indicators from a distance and, sometimes, covertly and even automatically. These are different "desirables." Viewing from a distance can allow observers to look at an entire crowd rather than just an individual at a checkpoint, for example. Subjects may not know that they (or, for example, their communications) are being viewed, even if they know that observations take place. Other observations may be more fundamentally covert. Automaticity refers to the desired ability to collect and analyze a great deal of data with machines, inserting humans only by exception. This, of course, could greatly increase efficiency. These matters are discussed further in Chapter Six.

Having discussed various screening issues, what follows is different in character. Perhaps the strongest of our cross-cutting observations relates to the crucial role of combining information.

Combining Information: From Heuristics to Information Fusion

This section is longer than the preceding ones because we concluded that progress in information fusion should be a critical element of future work, but we did not find an abundance of directly relevant literature. Thus, we sought to add value to the existing literature by providing our overview of where some promising areas may lie.

Initial Observations

The Spectrum of Combining-Information Methods

Before proceeding, we note that combining information is a matter of degree and level. Even in something as narrow as polygraph testing, operators have long used information from several physiological measurements to make an overall estimate. Also, from time eternal, security officers have combined information when, e.g., noticing that an individual is *both* looking around furtively *and* seemingly attempting to move amidst crowds of people. In this chapter, however, we are dealing largely with higher-level aspects of information fusion, which may combine information from heterogeneous classes of data (e.g., gait, behavior, and past arrest history), and perhaps do so in a future version of a "fusion center," rather than at a security checkpoint itself.[*]

The Basic Challenge

A crucial step in assessing possible malign intent is reaching an overall assessment based on combining diverse indicators and overlaying inference (since the indicators are not able to reliably and selectively detect intent in most cases). The best that can be hoped for is a stronger sense of relative likelihood (even if small), so as to know when to look harder at an individual, or even take preemptive action. To illustrate, suppose that an individual attends one radical meeting each week for six months

---

[*]   Existing DHS fusion centers have recently been discussed in scathing terms in a bipartisan Senate report by Senators Carl Levin and Tom Coburn (U.S. Senate, 2012). This report's references to fusion centers have in mind future centers that would have very different classes of information and analytic tools available to them. We did no research on the current centers.

and also accesses radical websites.[*] How likely is it that he is contemplating membership in a radical group? How do we combine knowledge from the two indicators? How do we "fuse" the indicator reports to modify our likelihood assessment? In this context, fusion is the process of combining information from various sources (similar and disparate in character) with the intention of obtaining a better composite of that being studied (see a review discussion in Perry, Signori, and Boon [2004]). Fusion may be accomplished with simple methods or much more sophisticated mathematical processes, as described briefly below and, in more detail, in Appendix D. What follows in the remainder of this section is a brief discussion of fusion methods that we thought representative of the classes of fusion methods in use for other applications and which could be used to assess hostile intent.[†]

Heuristic and Simple-Model Methods

Diverse simple methods are in use today in many domains, and for screening at many airports in the United States. Four of them are discussed below. Opportunities exist to greatly improve them over time, and it is possible that they eventually could provide a good fraction of what would otherwise be attempted with more expensive and demanding approaches.[‡]

Checklists, Negative and Positive

Checklists are perhaps the most familiar of decision aids. They can be either negative or positive. One kind of negative checklist has a set of indicators, any of which, if observed, triggers additional screening, monitoring, or both. A positive checklist might have a set of indicators

---

[*]   In this report, we use "radical" more or less synonymously with "extremist," or "potentially violent extremist." In other contexts, a "radical" may be a perfectly legitimate and respected figure that just happens to be seeking more than incremental change.

[†]   We do not treat data mining here because we do not see it as information fusion, but others sometimes do, and it is important in any case, for reasons discussed in both Chapter Six and Appendix D.

[‡]   As noted in a classic paper by psychologist Robin Dawes, simple, and even linear decision aids have a track record of being remarkably effective in comparison with expert predictions (Dawes, 1979).

such that if *all* criteria are met, the individual is deemed not risky and put in a fast-track line.

Security guards often use what amounts to a mental negative checklist. A positive-checklist approach is familiar to anyone who has to show identification and be on an expected-visitor list to gain a non-escort badge in a high-security office building. For the context of this report, a positive checklist might require that a person at checkpoint have ID, appear visually to be the same as the pictured person, have a clearance on record, and not be on a watch list.[*] Numerous variants of checklists exist. Most of these are essentially simple examples of an index approach, as discussed next.

Risk Indexes or Scoring Methods

Index methods (scoring methods) typically characterize a risk level by summing a number of indicator scores (perhaps with a rule such as "risk is high if the sum is greater than" some level), or by computing a simple product, as with treating risk as the product of a likelihood and a consequence. If the score exceeds some threshold, then risk is considered, in different settings, to be significant enough to justify more tests (medicine), more caution (granting of credit), or further screening and monitoring for detecting possible terrorists.[†] We highlight score-based methods here because, if fusion techniques are to be operationally feasible, they may need ultimately to be rather simple even if they come from a more substantial research base. This is especially so when considering decision aids for security personnel at ordinary checkpoints or monitoring stations, rather than special high-capability teams operating at a regional or national fusion center.

Scoring methods are used widely and are often referred to by government agencies and industrial organizations as part of best practices in assessing risk. They were used for decades in DoD's force planning

---

[*]   See the earlier discussion of TSA's "Trusted Traveler Program" in this chapter.

[†]   As an example, one such National Institutes of Health index asks for age, gender, total cholesterol, HDL cholesterol, smoking (yes/no), systolic blood pressure, and medication for high blood pressure (yes/no). It then reports the likelihood of a heart attack over the next ten years (National Cholesterol Education Program, no date). The underlying formula is based on statistical analysis of medical data over many years.

(Kugler, 2006). Score-based methods can be simplistic, moderately simple, or sophisticated, as with the Analytic Hierarchy Process (AHP) (Saaty, 1999). The individual indicators (typically risk factors) should be carefully chosen and should either be independent or have the correlations among them be accounted for in the scoring. Significantly, it is often necessary for good scoring methods to be nonlinear, as when evaluating a system that will be ineffective if any of several critical components are inadequate, or a strategy that will be unacceptable if it fails to address any of several conflicting objectives adequately. This is an issue in a version of portfolio analysis developed for defense and other types of strategic planning (Davis, Shaver, and Beck, 2008b; Davis, Gompert, Johnson, and Long, 2008a;).

Scoring methods have problems, as summarized in a recent paper (Hubbard and Evans, 2010), albeit a paper that gives only one side of the argument. The problems noted are that

1. Unaided subjective estimates of both probability and consequences are subject to cognitive biases.
2. There can be considerable variability in how qualitative labels such as "low" and "high" are interpreted, even among experts, and even though the illusion of communication exists.
3. Reasoning errors can arise when implicitly assuming ratio scales when in fact the values are more like ordinals (e.g., something with a score of 4 may not actually be twice as likely as something with a score of 2), or by range compression (e.g., assigning a score of 5 for anything "large," which eliminates differences between large and gargantuan).
4. The methods may be misleading because of statistical correlations among the inputs (e.g., two security defects may tend to occur together because of a common cause, such as poor management).

Because of such issues, authors such as Hubbard and Evans insist that risk assessments should be based on explicit probabilities and consequences, even if those must be subjectively estimated. They go on to

argue that various methods can be used to improve the quality of the subjective estimates.

The cautions are all valid, but the prescription is not uniquely right and can be counterproductive. If a problem is well enough understood so that concrete probabilities and consequences can be used, that is preferable. However, if a problem is complex, fuzzy, and poorly understood, tightening up analysis by using a "rigorous formulation" of only one element of the phenomenon may not be helpful. In this report, we are dealing with weak, ambiguous, and poorly understood relationships among behavioral indicators and security risks. Later in this chapter, we discuss Bayesian methods and extensions and conclude by arguing that more recent and "fuzzier" methods show more promise than familiar probabilistic approaches.

Scorecards and Score Sets

Although it would sometimes be convenient to reduce an assessment of risk to some composite score, it would often be better to provide a decision aid showing individual indicators or a set of aggregate indicators, as with a colored scorecard. This is especially so when the assessment "should be" nonlinear for one reason or another (e.g., the presence of an oddity in any of several indicators should be a concern, with normalcy of other indicators not canceling out the exception). If the user is not to be overwhelmed, however, the number of indicators should be modest or their presentation cognitively effective.

Conditional Indicator Sets

A more complicated method is to create conditional indicator sets. This could be a check list that reads: "If indicator a, c, and d are observed (but not necessarily indicators b, e, . . .), then the basis for concern is high." This results in a checklist in tabular form that lists the combination of indicators and the resulting assessment. This approach is often superior to a simple heuristic score based on linear weights. However, if the set of indicators is large, or if the indicators have degrees, then the combinatorics become onerous. For example, with five indicators, each of which has possible values of low, medium, or high, there would be 243 combinations to deal with. Still, in the era of computers and apps, what was once too complicated for routine operations can be doable.

Information Fusion

In this and successive sections, we discuss combining methods often referred to as information fusion. There are two aspects to fusion: (1) at a given time, combining several indicator reports from disparate sources, and (2) combining the result with previous reports to obtain an updated estimate of likelihood (Hall and Llinas, 1997). The second (updating) is the better understood, as discussed in Darilek et al. (2001). The first is more problematic because reports are often a mix of quantitative and qualitative information without an obvious mechanism for combining them.

A third aspect of the fusion problem is especially important for our context, and is a combination of the previous two: periodic updating of reports over time when the reports come from multiple sensors and sources that are often disparate in character. In all of this, it is important not to lose sight of the overall objective of the fusion process: to modify the likelihood that an individual or group is contemplating some hostile act—that there is some "basis for concern" (as shown in our conceptual mode, depicted in Figure 1.2).

Bayesian Updating

Perhaps the most common fusion method is Bayesian updating, which is based on Bayes' rule (Feller, 1950; Mood and Graybill, 1963; Raiffa, 1968; Stone, Barlow, and Corwin, 1999). Bayes' rule is a statement of conditional probabilities. It can be used to assess, for example, the likelihood that an individual is about to join a radical group given that (i.e., if) he is observed attending a group meeting.[*]

A simple example illustrates the method. Suppose that without any information from indicator reports, the likelihood that an individual is about to join a radical group is 0.01, or $P(A) = 0.01$. That means that the likelihood that he will not join the group is 0.99. These are referred to as the *prior* probabilities. Now suppose a report is received from an observer reporting that the individual attended a radical-group meeting. Suppose further that we have enough experience with both the observer and the history of previous meeting attendees to estimate

---

[*]   See Appendix D for an explanation of the development of this rule.

how significant attending such a meeting is. By applying Bayes' rule, we could update our assessment of risk. A single indicator report, for example, might raise our assessed likelihood from, say, 0.01 to 0.16. The latter is called the *posterior* probability. Although easy to implement in simple problems, and although a huge literature exists on Bayesian updating and its applications,[*] we believe that such methods are unlikely to go very far in our problem area, except for the simplest of instances. There are two basic problems with this approach:

- First, it is difficult to reasonably estimate all the conditional probabilities required of the approach, especially with multiple indicators and multiple values thereof.
- Second, the method does not allow us to express information such as evidence that an individual is engaged in two or more activities at the same time if, logically, he can only be involved in one. (Probabilities are calculated for the basic hypotheses only. That is, the method produces a probability distribution for the set of basic activities. It does not directly account for an individual who may be engaged in two or more activities at the same time. Nor does it allow us to account for the probability that an individual is or is not engaged in some combination of the activities.)

Belief Function Methods

Some relatively new methods are based on Glenn Shafer's *belief function* concept. Belief functions are considered a "less restrictive Bayes." In his book, first written in 1971, Shafer distinguishes between probability and belief. There is a fundamental shortcoming of trying to do evidential reasoning in terms of binary probabilities alone.[†] If we estimate the probability of an event occurring as P, then we are "forced" to assume that the probability of the event not occurring is $1 - P$. Sup-

---

[*]   Research on the use of Bayesian updating in intelligence analysis indicates that its utility may be highly variable across analysts, due to the possibility of amplifying bias (Poole, 2009, pp. 21–24).

[†]   A better way to say this is "in terms solely of probabilities of propositions being true." That is, one may use the apparatus of probabilities, but distinguishing between probabilities of necessity or provability, rather than probability of truth (Shafer and Pearl, 1990a).

pose, however, that we receive a report from a trusted agent that our suspect has just attended a radical group meeting and appears likely to join the group. His past reporting suggests that he is 70 percent accurate. From this, we can justify a "belief" (not a probability) of 70 percent that the suspect is about to join the radical group, but only 0 percent "belief" that he will not. That is, we have no evidence to support the proposition that he will not join and therefore, unlike probabilities (where we would assess 30 percent to the likelihood that he will not join), the two beliefs about the suspect joining or not joining need not sum to 100 percent. Together, then, these two constitute a belief function (Shafer and Pearl, 1990b).

For a given threatening activity, *A*, we have two hypotheses: Our suspect is engaged in this activity, or he is not. For example, we may assess whether an individual is or is not "Developing Intent," as discussed in Chapter One.

As reviewed in Appendix D, using Bayesian analysis, we can conditionally update probabilities based on the collected evidence. For belief functions, the updating uses Dempster's rule of combination, which produces updated estimates of activity likelihoods by multiplying the likelihood estimates and then normalizing (Dempster, 1967). The result is a normalized orthogonal sum of the belief functions for each of the reports.

Ultimately, there are three shortcomings in using Dempster's rule of combination when used to fuse indicator reports and activities:

1. There is no good way to deal with total conflict, that is, when two indicator reports contradict each other or when two activities on which there are indicators cannot exist at the same time.
2. The Dempster-Shafer approach allows us to express belief in the disjunction of any combination of the basic hypotheses (threatening activities, in this case). Hence we are able to assess belief that an individual is engaged in activity A "or" B, but not in A "and" B. Nor does it allow us to express belief that an individual is *not* engaged in an activity. This does not fit the nature of the problem we are working on or the model presented in Chapter One. For example, it is possible that some indicators

suggest that an individual is attempting to join a radical group, while others suggest that is not the individual's intention. The Dempster-Shafer method cannot handle reports of this nature.

3. The normalization rule used in Dempster's rule of combination has the effect of ignoring conflict and not accounting for evidence conflicting with the proposition having the highest belief score. As a result, the application of the rule may have the effect of producing inconsistent results (Sentz and Ferson, 2002).

The Dezert-Smarandache Theory of Plausible and Paradoxical Reasoning

The indicator reports that point to various threatening activities are likely to be imprecise, fuzzy, paradoxical, and highly conflicting. Combining the information from such reports can therefore be just as imprecise, and dealing with likely conflicting information is certainly problematic. We noted earlier that resolving conflicting information may not be possible in all cases with the Dempster rule of combination. However, several other recent methods of combination have been advanced in efforts to improve the situation. One is the Dezert-Smarandache theory (DSmT), which we concluded has promise (Smarandache and Dezert, 2009a, 2009b). We discuss the merits of DSmT in Appendix D.

Other Combining Methods

Several other combining methods may be of use in information fusion. In this section, we briefly introduce four of them: possibility theory, multi-attribute assessment, mutual information, and filtering.

Possibility Theory

The phrase "theory of possibility" was coined by Lotfi Zadeh (1978). Possibility theory is an uncertainty theory that deals with incomplete information, and is therefore well suited to the problem of discerning individual or group activity that may indicate hostile intent. Possibility theory states that any proposition not known to be impossible cannot be ruled out. A possibility distribution is taken to be a membership function of a fuzzy set (Dubois, 2006) of mutually exclusive values. As with all the combining methods discussed, we first start with some

measure of how likely it is that an individual or group is engaged in some hostile activity based on some indicator report. With the possibility distribution defined, Dubois and Prade (1994, 1998) introduce the *possibility measure* and the *necessity measure*. The two concepts are duals. There is no unique method of combination using the possibilistic approach. Rather, the method chosen will depend on what assumptions we make about the reliability of our sources of information. This "complication" is good substantively, in that this approach is the first of those that we have described that allows for taking such reliabilities into account. In earlier chapters, we discuss indicators and the likely means of observing subjects looking for these indicators. In some cases, we rely on humans to provide indicator reports from direct or indirect observation, and we also discuss technical means. Using possibility theory, there are two modes of combining reports from two or more disparate or similar sources: the conjunctive mode and the disjunctive mode. The former is used when all the sources agree somewhat and are reliable. The latter is used when the sources disagree such that at least one of them is wrong.

Multi-Attribute Assessment

The simplest (but perhaps not the most accurate) way to deal with the problem of fusing information is to create a weighted sum of the activity likelihoods included in the indicator reports.[*] As mentioned earlier, weights generally imply some notion of relative importance; in this case, the weights would be assigned to the reports—and ultimately to the sources. This, then, is another way to account for the reliability of the sources. However, it is better to consider the weights as reflecting the relative reliability of the reports. Regardless of how well we are able to assign weights that truly reflect the relative reliability of the various reports and report sources, a weighted sum is inherently flawed because the likelihood estimates need not be additive. Nevertheless, as a means of comparison, the method is useful at times. The objective of multi-attribute assessment, in this context, is to derive a single assessed likelihood. In this formulation, we assume that several indicator reports

---

[*]   Adapted from Perry and Moffat (2004).

consisting of the likelihood that an individual or group is engaged in one or more of the hostile activities we have identified. The methods we discuss to develop this single assessment derive from Multiple Attribute Decision Making (MADM) theory (Hwang and Yoon, 1981). Using MADM would be a way to assess the likelihood that an individual or group is about to engage in each of the threatening activities. In this case, the "attribute" would be the source, and the value of the attribute would be the reported likelihood estimate based on the observed indicator. The choice of one technique over another depends on the nature of the sources whose likelihoods are being combined and their relation to one another. There are three basic methods: (1) Simple Additive Weights (SAW), essentially the method discussed above; (2) Weighted Product, which is similar but generates a product instead of a sum; and (3) a Keeney-Raiffa multi-attribute utility method (Keeney and Raiffa, 1976). Multi-attribute assessment uses a nonlinear utility function, and the technique allows for the consideration of possible interactions between the reports, which could be important.

As mentioned earlier in this chapter, a number of related and important methods exist and could be adapted to the present purpose, but we do not elaborate on them here. These include the Analytic Hierarchy Process (Saaty, 1999), Value Focused Thinking (Keeney, 1992), and RAND methods for "portfolio analysis," which recommends multi-criteria assessment methods but recommends against combining scores into a single measure until, perhaps, strategic decisions have been largely made (Davis and Dreyer, 2009).

Mutual Information

This method examines the relationship among the threatening activities. It is less concerned with fusing indicator reports than were the previous sections. The question here is: What can we learn about other threatening activities given what we know about one or more particular threatening activities? For example, suppose that, based in several indicator reports, we conclude that it is highly likely that an individual suspect will join a radical group. Does that tell us anything about the likelihood that he will participate in target-identification activities? We refer to such questions in terms of asking about *mutual information*.

Mutual information is derived from information entropy (Cover and Thomas, 1991; Kullback, 1978; Shannon, 1948); it deals directly with independence among the activities. We construct a mathematical model that allows us to modify our estimates of the likelihood that an individual or group is about to engage in a threatening activity through our knowledge that the individual or group is likely engaged in some other threatening activity. Because one random variable informs another, we refer to this construct as *mutual information*. Mutual information is based on the concept of *relative entropy*. Relative entropy measures the difference in entropies as calculated with two probability distributions (Cover and Thomas, 1991). The difficulty associated with implementing this method is that it requires that we know the probability distributions on the activity likelihood variables to start.

Filtering
Filtering is a process that removes noise from a signal. When this is done repetitively over time, and perhaps with different sensors, it becomes an example of combining information. Applied to information fusion, the signal is the indicator report, and the noise is the inaccuracy associated with the uncertainties in the report and the errors introduced by the process itself. Of the various filtering methods, the most commonly used is the Kalman filter (Lewis, 1986). The combining process in a Kalman filter is essentially a sequential update of a state vector based on a prediction-correction process.

Summing Up Combining Methods
Information fusion is a critical component in detecting threatening behavior on the part of individuals or groups. Indicator reports are likely to originate in a wide variety of sources and sensors—some human and some technical, as discussed in earlier chapters. The common denominator among them all is an assessment of the likelihood that the individual or group observed is engaging in some threatening activity. This allows us to fuse a report from a remote heartbeat sensor with a human observation of some kind.

That said, the method best suited to fuse such information is far from settled. We have summarized a number of them in this chapter,

and elaborate further in Appendix D, but each has problems. The fused judgment may be better, but its false alarm rate may still be high, limited by the base rate and the rates of the indicators being fused. Considerable further research will be needed to understand which, or which combination, has the most promise, although we have offered some subjective judgments on the matter above.

Mitigating the Consequences of False Alarms

To complete this chapter on cross-cutting issues, let us touch again on the problem of false positive (or false alarms, or Type II errors). These impose costs in many dimensions. Given finite resources, it is important to minimize time (i.e., resources and money) wasted on unfruitful checking. Beyond that, however, there is cost when people's time is wasted, their fears raised, their dignity insulted, or their privacy invaded. There is cost to society when the security system demeans individual rights and dignities. There is also cost to society when procedures pivot on attributes such as gender, ethnicity, race, or religion. Finally, there is a security cost when large numbers of inefficiently handled false alarms reduce a security organization's credibility or legitimacy.

Technically, the issues are complicated by the need to consider not just the "average" effectiveness of the security system, but the distribution of results. If even a small number of people are severely inconvenienced or mistreated, that is a serious problem even if the average person has only minor inconveniences. If even a single violent extremist slips through security and is able to execute a significant attack, that is a serious problem even if the overwhelming percentage of attack attempts are foiled. We observe—in the scientific literature as well as in more usual discussions—a tendency to underestimate some of the considerations (such as direct economic cost), while focusing on others. We believe that it is important to recognize that

- The quality of the system, evaluated holistically, can be improved by (1) improving the likelihood of detecting someone with hostile

intent *or* (2) reducing the negative impact (costs) on those without hostile intent.

The second possibility is seldom discussed, which to us represents a gap, especially since behavioral science also has much to say about measures to mitigate the costs. We see three ways to reduce negative impacts, which we describe in the following paragraphs.

Improving the System's Efficiency

The direct side effects of false alarms (e.g., time lost, commerce interrupted, expenses incurred) go down as system efficiency improves, as with reducing the time required for secondary screening. Such efficiencies are most plausible when they involve computer-mediated sorting, comparing, and fusing of data, and least plausible when they require human-intensive actions, such as prolonged questioning or even interrogation. Many incentives and upsides exist for improving efficiency, however, so we do not deal further with the issue here.

Reducing Effects on Dignity and Perceived Violations of Civil Liberties

The scientific literature seldom discusses the profound side effects that occur when people are treated in ways that they perceive as unfair, offensive, humiliating, or in violation of their liberties. These issues are certainly recognized, as when TSA has gone to considerable effort to improve both reality and perceptions regarding full-body scanning,[*] but the scientific discussion is disappointing—especially when what mitigations work and can be accomplished efficiently is in part an empirical issue worthy of study and analysis.[†]

---

[*]   See a report from the Congressional Research Service for issues arising in airport screening and some efforts to allay concerns or mitigate issues (Elias, 2011).

[†]   What needs to be mitigated is, of course, culture-dependent. We largely have in mind consequences of false alarms in detection systems in countries with values akin to those of Western democracies. Intrusive interrogations in some countries are more commonplace. There are also distinctions between what is reasonable in normal screening and, say, in a targeting search as might occur in a theater of conflict.

Relevant literature exists in other domains. For example, the U.S. Joint Forces Command, in worrying about population-centric operations, supported work drawing on psychological research and business-world experience (Helmus, Paul, and Glenn, 2007). One point from the study is "Virtually every action, message, and decision of a force shapes the opinions of an indigenous population: how coalition personnel treat civilians during cordon-and-search operations, the accuracy or inaccuracy of aerial bombardment, and the treatment of detainees. Unity of message is key in this regard." The study also discussed the importance of anticipating that mistakes will assuredly occur and, therefore, of being proactive. We believe that much could be done with six classes of action:

- transparency
- destigmatization[*]
- explanation
- apology
- compensation (e.g., travel vouchers, cash, expediting travel by another route after a missed flight)
- prompt correction of errors (e.g., correcting watch lists).

Another potential domain could be psychological research on "personal control." Allowing people to feel a sense of control over their situation can increase satisfaction, psychological "comfort," etc. In our view, it should not be difficult for research and analysis to draw on other domains for insights and suggested doctrine.[†]

---

[*]   This refers to reducing the perceived significance of follow-up questioning by, e.g., randomly choosing people for such follow-up to reduce the stigma of such questioning. Related suggestions were included in the White House Commission on Aviation Safety and Security (1997), Appendix A, and remain useful. We thank Brian Jenkins for pointing this out.

[†]   The reader might think of personally familiar examples, such as how a high-quality hotel deals with mishaps, or how a professional police force deals with people of varied backgrounds and ethnicities.

Deterring Abuse

A third component of reducing the ill consequences of false alarms is again familiar and well-studied in other domains. This relates to avoiding abuse by those within the security system. For example, networked use of extensive personal information could be both invasive and injurious, leading to stolen identities, sullying of reputations, and the like. Further, methods such as probing, sequential screening, interrogation, and detention for questioning can easily include abuse for a variety of reasons. We see the need for pointed research and analysis, specific to the terrorist-detection problem, but informed by the extensive knowledge based in other domains on how to minimize the likelihood of abuse. Two elements of this are

1. monitoring of the monitors (or, more broadly those in the security system)
2. deterrence (e.g., enforcement of laws punishing severely those who misuse information or abuse authority when giving interrogations or holding individuals for questionings).

Identifying appropriate measures is inherently complicated by conflicting considerations and organizational resistance, but much has been learned over the years about how to "square the circle" on analogous issues such as deterring police abuse.

Why It Matters, Even If Detection Were the Primary Objective

In the absence of better ways to reduce side effects and abuse, there will be continual efforts to constrain or further constrain methods such as profiling, screening, sharing information, and information fusion—even when they have the potential to significantly reduce the likelihood of detecting and thwarting attacks. Thus, this section on "mitigating" negative consequences seems to us central to the problem, as well as important in itself. One subject for future investment in behavioral research, then, might well be on the mitigation challenge.

Summing Up

This chapter has covered considerable conceptual and technical ground. In our review of the major issues surrounding screening, detection, and data fusion, we found a few points particularly important to emphasize. First, it is important to consider multiple dimensions when designing a screening and detection system, including multiple layers of screening, the degree to which covert versus overt observation is employed, and the degree to which potential suspects are intentionally provoked to observe their behavioral reactions. In most cases, there are trade-offs between false positives and false negatives. Contextual conditions, such as the level of ambient risk, will affect decisionmaking on what measures are most acceptable and where detection thresholds should be set. Perhaps most important technically, because of weak signals, high false-alarm rates, and countermeasures, and because no "silver bullets" are on the horizon, *information fusion is likely essential for success.* It is not a panacea, and may or may not succeed, but information fusion seems to be the only hope. Finally, effort must be made to reduce the negative consequences of screening—both to protect society's values and because public cooperation highly important in the detection of threats.

Conclusions

Continuing Themes

Some themes have been important throughout our study. First, it is clear that most relevant behavioral indicators will have low detection rates and large false-alarm rates. Such problems are exacerbated by adversary countermeasures. Detection-system performance, however, can in a number of instances be improved by probing or otherwise stimulating responses.

Second, because of the weak signals and false-alarm rates, there is need for two classes of activity: (1) *pattern discovery* by man-machine study of data and (2) *information fusion*. It remains to be seen what either or both can accomplish, but we expect the gain to be considerable.

Pattern discovery often requires large data sets for training machines and extracting weak signal from background. Automated tools are essential (e.g., data mining, machine learning), but man-machine cooperation will probably remain optimal. The related state of the art has improved dramatically, but is still in its infancy.

Information fusion varies in degree, scope, and character. A checkpoint officer may use simple tools to achieve a significant degree of fusion (comparing passport with face, looking for signs of anxiety, and noting responses to questions), but regional or national fusion centers could draw on far more extensive data (some of it highly protected) and use much more sophisticated tools for fusion, some of it in real or near real time and some of it over a more extended period (perhaps while individuals remain under observation or tracking). Also, fusion may be passive or adaptive and interactive, as when initial fusion sug-

gests pointed questions to ask or other probing (e.g., isolating someone for special questioning), screening, or continued monitoring. Sophisticated fusion analysis may generate highly simplified but individualized and contextualized rules usable by on-the-spot officers. Fusion, then, is a large and multifaceted subject, and the scholarly literature on it is not at all well structured to dealing with detecting hostile intent as yet.

A third theme is that effective fusion will require networking on an extraordinary level to draw on information of disparate types and sources. As with fusion itself, networking can vary greatly in scope, the types of connections, accesses, and so on.

A fourth theme (covered in Chapter Seven) is the need to reduce the consequences of false alarms *both* by reducing false-alarm rates *and* by mitigating the negative consequences when they occur.

Many observations can be made that bear on these general themes. We offer a few of them here. Some are our attempt to put our findings in perspective. Some bear on where the challenges and opportunities lie.

Observations

Operator Initiative Versus Scientific Testing of Methods

In the course of our study, we frequently noted that "operators," whether in law enforcement or intelligence particularly, are currently well ahead of the science base in many instances. Many are skilled in using intuitive low-tech methods to observe behavior; some already exploit or experiment with advanced technology. The New York Domain Awareness Center illustrates that many items discussed in this report are operationally feasible, including degrees of fusion, networking, and probing (New York City, 2012).

It is also true that operators are sometimes more enthusiastic and less skeptical about technologies and methods than is justified by the science base—this is especially true when advanced mathematical methods are used to support predictive analyses. These can have a certain cache, while not being fully understandable. Contributing to the inadequate skepticism is the fact that operators (and indeed all of us)

are subject to psychological biases when making intuitive inferences, a problem that arises regularly in criminal law and other fields, such as medicine. For example, humans do not naturally account properly for base rates when estimating probabilities or assessing the significance of an additional increment of information.

As pattern-discovery and information-fusion methods evolve, it will be essential to vet decision aids with solid analytic reasoning and, where possible, to ground them empirically. This said, demands that advanced methods should not be deployed or employed until their validity is well established scientifically are indefensible: Science moves slowly, the experiments needed for testing are complex and only sometimes feasible, and security threats are a current reality. What *is* feasible and appropriate is to demand that the decision-aiding aspects of detection systems be well informed and updated by the substantial current knowledge about such matters. Using that knowledge can avoid many mistakes that would otherwise occur.* In addition, considerable investment should be made—and sustained—to improve the empirical base.

Knowledge in the Private Sector

The private sector has developed technologies that may be useful for preventing violent attacks that are sometimes at least as sophisticated as, and often more polished and ready for use than, what some government agencies possess. Microsoft's Kinect camera is perhaps the best example of this. Originally developed for Microsoft's Xbox video game system, the Kinect camera has been used, for example, for motion-capture applications to understand emotions in gait and facial expressions.

This suggests the value of an initiative to review the state of the private and commercial sectors for useful technologies that might uncover fruitful advanced technologies that could be used "as is." Also, it might better leverage academic research that has been conducted, and models and techniques that have been built, on top of those commercial technologies. Such a survey would, in some instances, have

---

* Examples include inappropriate kinds of profiling, poor lineup procedures in criminal law, and seriously misleading inferences based on failure to account for base-rate information.

to protect proprietary information and would need to go well beyond what is easily found in the public domain.

The Big Data Phenomenon

Researchers always want more data, but there are special needs with regard to improving the usefulness of behavioral observables. Our report took as general an approach as possible, but this meant lumping very different things together (e.g., complex insurgent attacks on hardened targets on the one hand, and lone-wolf mass murders on the other). Research programs need to assure that data being collected increasingly recognizes important distinctions, such as classes of attack, political and social context, and individual variations.

As one example, implanting and remotely detonating an IED on a military target in the context of a rural insurgency carries a very different behavioral signature than embedding and detonating suicide bombers in crowded locations within Israel. This is due in part to the mode of attack (and the expected fate of the attackers), but also the nature of the target environment and immediate possibilities for detection and interdiction or retaliation (military versus civilian, rural versus urban, etc.).

The principle illustrated by the example is general. We know from the existing science base that behavioral observables and their interpretation as indicators vary with cultural and contextual factors and across individuals. It follows that training sets for machine learning and rules of thumb that are developed for security personnel need to be based on large and substantial data sets that allow the distinctions to be recognized and better understood.*

One special problem that arises in discussing data and dissemination within the relevant scientific community is that some data are either sensitive for various reasons (e.g., privacy, a central concern in medical research) or proprietary. Although we did not look into such matter in this report, various possibilities were mentioned. One would

---

* Related issues are discussed in two National Academy of Sciences reports: Chauvin, 2011, and Fischhoff and Chauvin, 2011. Also, Chung and Pennebaker (2011) addresses such data needs for language-related indicators.

be for a trusted party to house sensitive data and independently perform all testing of new models and techniques within a specific detection-system domain. Another possibility would be to use a cryptographic protocol to control sharing, perhaps between particular agencies and private firms.

Information Fusion

We have given considerable weight to our discussion of information fusion (Chapter Seven and Appendix D). The challenges are formidable. The primary intent is to provide credible assessments of whether an individual or group being assessed merits special concern, such as additional screening, monitoring, or even intercept. This requires that (1) the activities and indicators used must be reasonably associated with hostile intent; (2) indicator reports can be mapped into measures of evidence (e.g., likelihood estimates); (3) the combining algorithm accounts for the nature of the activity set; and (4) reasonable intervention thresholds are established. We have reviewed a number of methods in this report (including in Appendix D), but very few of them have been carefully evaluated analytically or in laboratory experiments, much less operationally tested. Thus, much remains to be done. We can touch briefly here on the four classes just listed:

- *Activities and Indicators.* Are the activities and indicators we suggest adequate? Are they sufficiently complete? Are the activities well enough related to hostile intent so as to avoid unacceptably numerous false alarms? Are the activities plausibly observable? Is the set of activities and indicators readily expandable based on additional information?
- *Measures of Evidence.* To use the approach we introduced at the outset, in Figure 1.1, indicator reports must be translated into a measure of evidence that the individual or group is engaged in that activity (e.g., a measure in probability terms). Hence, information fusion in this sense is the combining of evidence to arrive at an overall assessment about whether a subject raises enough concerns to justify further actions. This is true of both qualitative and quantitative indicator reports. Indicator reports for example

are likely to originate in a wide variety of sources and sensors—some human and some technical.

- *Combining Algorithms or Rules.* Given a reasonable set of activities providing evidence of possible hostile intent, how do we combine the evidence—especially if the evidence is from disparate sources and spaced in time? We reviewed a number of possible approaches in Chapter Seven and Appendix D. What criteria should be used in choosing among them? We noted two important ones: (1) accounting for "fuzziness" (e.g., the evidence may be ambiguous, contradictory, or complicated, as with "we have evidence of either this, or this and that, but we're not sure which") and (2) operational and contextual considerations (even a trained behavior detection officer may need a checklist, whereas a fusion center could use more powerful methods if monitoring an individual over time, perhaps with probing and specialist intervention providing tailored information). We do not see a single method being universally appropriate.

    One more point is crucial here. Not all evidence is equally credible, and sometimes little of the evidence is very credible at all. Thus, a major issue in "combining algorithms and rules" is how to fold in credibility assessments. This is something dealt with in "possibility theory" (Appendix D) but it is not usually addressed at all.

- *Criteria for Intervention.* Ultimately, the assessments we discuss are supposed to determine whether follow-up action should be taken, whether in the form of secondary screening, monitoring, or even arrest. What should the criteria be, and what should be the follow-up, if any? There can be no general answer to this (deciding to require a secondary check with more probing, assigning a surveillance team to follow a suspect over days or weeks following some observed activities, or "taking down" someone intent on imminent suicide bombing are different matters), but developing related doctrine is a major challenge that should be informed by research.[*]

---

[*] As analogies, American police forces have largely shifted away from high-speed vehicle chases because accumulated evidence about the harm they cause and the feasibility of suc-

Our conclusions about information fusion in this report are based largely on theoretical and analytical considerations. More such evaluation is possible, but there is pressing need for more concentrated research testing the more promising methods with existing data on past hostile acts. The question would be this: "Would a given combining algorithm activity set and associated indicators have helped in thwarting hostile acts (not always, but often enough to be valuable)?"

Informing Investments

Although we cannot advise in specific detail on future investments based on this study, we can offer some insights about how to proceed. We have highlighted certain themes that we believe would be central to improved success exploiting behavioral information to detect potential attackers, such as suicide bombers. Also, Chapter Six identifies criteria that can be used to assess technologies and methods when thinking about resource allocation. Also, based on prior work with some analogous features, we recommend that a "portfolio analysis" approach be taken with a number of distinct criteria (objectives), particularly upside potential and vulnerability to countermeasures.

A standard difficulty encountered by organizations is that simple prioritization schemes often work quite poorly for complex decision-making because the results so often follow the most recent headlines or points of sensitivity expressed by senior leaders, rather than taking a more comprehensive and longer-term look. A better approach is portfolio analysis, good methods and tools for which have been developed over the past decade.* In this context, portfolio analysis refers to finding a good *mix* of investments so as to attend to quite a number of

---

cessful intercept without such chases. In contrast, police are largely assured the right to act in what they believe is self-defense, even with deadly force, despite occasional tragic errors.

* RAND has developed methods and tools for higher-level decisionmaking under uncertainty and disagreement (Davis and Dreyer, 2009; Davis et al., 2008a, 2008b; Davis, Shaver, Gvineria, and Beck, 2008c) and more mathematical methods to use in aspects of R&D investment where optimization is feasible (Chow et al., 2012). MITRE has developed other, partially similar methods in a variety of projects (Garvey, Moynihan, and Servi, 2012; Moynihan, 2005). Both approaches have been applied in multiple projects. A significant number of analogous commercial methods exist, although many are restricted to uncritical use of linear weighted sums.

semi-conflicting objectives, such as time scale of value, assuring that all critical components of capability are dealt with (e.g., deployability and sustainability, not just laboratory capability), and looking for high potential while at the same time attempting to limit risk and cost. Such analysis is much more difficult than financial investment because much of the input data for analysis are inherently subjective or based on preliminary assessments. Nonetheless, much can be done.

An especially big challenge for investment is developing portfolio options. It is by no means difficult to obtain long lists of discrete programs in which to invest, since many laboratories, companies, individuals, and agencies will have good suggestions. It is much more difficult to package them into sensible *composite options* so that investments attend properly to the multiple objectives while dealing with constrained budgets. One promising approach involves computerized generation of the many large combinations of discrete options available, followed by a screening analysis that can quickly eliminate most options as illogical or inappropriate, and that can find options that are near the "efficient frontier" (Pareto curve) by at least one set of assumptions (Davis, Shaver, Gvineria, and Beck, 2008b). The resulting set can then be assessed in more detail by humans, using a portfolio-analysis framework, but with full recognition that many subjective judgments must be made.

Finally, we note that core elements of sound resource-allocation analysis are missing. We encountered virtually no information that would directly inform using concepts such as "production curves," curves of diminishing returns, or tradeoff analysis. Such concepts are most meaningful when evaluating acquisitions rather than investments in R&D, but even R&D investment decisions need to be informed by approximate versions of such concepts, as suggested in Chapter Seven with our highlighting of upside potential as a criterion. R&D investment decisions also need to be informed by approximate estimates of eventual costs for acquisition and operations. Much could be done along these lines to better structure investment decisions despite the early development status of many of the proposed technologies and methods.

Takeaways

We found a number of important "takeaways" from our survey:

- Despite exaggerations found in commercial claims, studies, and the media, there is current value and unrealized potential for using behavioral indicators as *part* of a system to detect attacks. Unfortunately, analytic quantification of that potential is poorly developed.
- "Operators" are often well ahead of the science base, which is sometimes good and sometimes bad. It is very important that programs build in and sustain objective evaluation efforts, despite budgetary pressures and the tendency to see them as mere nice-to-have items. The evaluations should be subjected to objective peer review and adequate community scrutiny, even if security considerations would that such review should be accomplished within a domain of cleared personnel, with limited distribution, etc. For example, the federally funded research and development centers (FFRDCs), national laboratories, National Academy of Sciences, and other special national panels have conducted analogous evaluations for decades on a classified basis.
- Many serious problems and errors can be avoided by up-front review of procedures by experts familiar with the subtleties of detection and screening in conditions of high false-alarm rates and low base rates. Although full validation of techniques may take years (at a time when the dangers of attack are current), many problems can be avoided with existing knowledge. Some problems so avoided are quite significant to privacy, civil liberties, and the efficiency of travel and commerce.
- DHS and other security organizations are making efforts to experiment with and evaluate proposed methods—sometimes with laudable and ambitious scientific trials that have reported encouraging conclusions (which are difficult to judge without detailed access to data and methods).
- Operators, their agencies, and the scientific community have not done enough to understand how to mitigate the considerable bad

consequences of detection systems, which invariably have false-alarm problems. Much could be done.

- Information fusion is critical, not just desirable, if behavioral indicators are to achieve their potential. Fusion should occur not just within a given method (as within polygraph methods), but with information across activities and phases. Methods for accomplishing this are very poorly developed. This said, it remains to be seen how much can realistically be accomplished.

- Information generation, retrieval, integration, and sense-making will place enormous demands on both automated methods (e.g., including for "big data") and perfecting human-machine interactions: Machines can process vast amounts of data, but interpretation will continue to depend critically on human expertise and judgment. "Optimizing" should be for man-machine cooperation, not automation, despite what some technologists may be inclined to emphasize.

- Very little research has been done to understand how much is enough, or what the curve of diminishing returns looks like, but, subjectively, it seems that *major* improvements in detection are plausible with networked real-time or near-real-time integration of information. This would include not just integrating information of the CIA and FBI (much discussed since 9/11), but also in integrating (fusing) (1) proximate information at checkpoints with future versions of fusion-center information and (2) criminal, commercial, security-related, and even whole-life information. All of this is hypothesis. Developing a sharper understanding of payoff potential should be a priority task for objective research and analysis.

- Contemplating such steps raises profound issues of privacy and civil liberties, but the irony is that commercial organizations (and even political parties) are already far ahead in exploiting the relevant technologies and forever changing notions of privacy.

- Investment decisions about individual technologies and methods should be informed by a structured portfolio-analysis approach using the something like the dimensions of Figure 7.1.

Ideally, we would end by recommending further investment in specific technologies and methods. That, however, is not feasible because nearly all the technologies and methods that we studied appear from the literature to have at least some benefit (none are truly pseudo-science, although they might be judged so when viewed for stand-alone effectiveness), so the issue becomes one of cost-effectiveness when evaluated on a "system" basis (e.g., as suggested by Figure 7.2), including taking account of operational considerations and information fusion of various types.

Methodological Notes

Literature Review

The study's literature review, although undoubtedly incomplete, was extensive. It included (1) publications by individual researchers studying detection of stress, deception, and related matters; (2) publications describing related technologies, models, and methodologies in development or in the field; (3) papers on related cognitive, behavioral, and psychophysiological theories; (4) reports by organizations that are producing technologies or implementing technologies in security-enforcement settings; (5) and programs that focus on behavioral indicators and hostile intent. Since there were thousands of relevant items, our bibliography is necessarily much more selective, especially when we could list good review articles.

We conducted extensive searches for books and papers with Google, academic databases, and archived papers from journals in diverse disciplines and applied fields. In retrospect, we observe items in at least the following classes: psychology, neuroscience, public policy, law, behavioral science, culture, sociology, criminology, information theory, decision science, and pattern recognition, as well as classes such as terrorism, law enforcement, and security studies.

Some particular research of prior RAND studies proved especially helpful (Davis and Cragin, 2009; Hollywood, Snyder, et al., 2004; Jackson, Chalk, et al., 2007; Jackson, Chan, and LaTourette, 2011; and Perry, Berrebi, et al., 2013).

Although literature searches inevitably follow clues that lead in directions that were not originally anticipated, we did use numerous

search terms in our systematic initial efforts. Table A.1 shows such terms for the possible benefit of readers who may wish to do similar searches.

We followed up on the initial searches by reading key sources and tracking down important references cited therein. As the project proceeded, of course, we found additional sources through reactions of reviewers, interviews and discussions, and other means.

It is difficult to identify a small number of key references. We did, however, pay particular attention to certain studies of the National Academy of Sciences, such as National Research Council (2003) on polygraphs, Perry and Vest (2008) on privacy, and National Research Council (2008) on Cognitive Neurosciences. We were also aware of somewhat related studies by the Defense Science Board (2012), the Intelligence Science Board (Fein, Lehner, and Vossekuil, 2006), and other organizations.

Table A.1
Some Search Terms Used

Emerging technologies + counterterrorism	Actionable indicators airport terminal behavioral indicators violence
Public private partnerships + national security	Behavioral patterns terrorist violence
Psychophysiological testing	Biometrics + risk prediction
Polygraph testing	Neuroscience + terrorist
Credibility assessment	Physiological cues
Remote observation	Emotional facial expression
Remote sensors violent intent	Facial recognition methods
Voice stress analysis	Involuntary reactions fear terrorist
CCTV surveillance + violence + security	Physiological reactions violent intent
Pre-incident indicators + terrorist attack	Assorted theories and frameworks
Predicting terrorist behavior	Signal detection theory terrorism
Risk prediction	Social network analysis terrorist
Hostile intent	Terrorist attack stages
Hostile intent detection	Terrorist ritual pre attack
Deception detection	Terrorist affiliates actions post attack
Passive methods violent intent detection	Terror cells after attack
Observable behaviors pre-attack + terrorism	

Searches and Interviews

We tracked developments in and studies from government institutions, academic and research laboratories, and private organizations focused on security. Much of this consisted of Internet searches and correspondence by email. Table A.2 lists most, if not all, of the organizations. Asterisks indicate where we held interviews.

Table A.2
Searches and Interviews with People and Organizations

Defense Advanced Research Projects Agency	**Foreign and Multinational Public Organizations**
Defense Sciences Office*	UK Human Terrain Analysis Team*
Federal Bureau of Investigation	North Atlantic Treaty Organization (NATO)
Federal Bureau of Investigation, Special Agent in Charge of Counterterrorism, Los Angeles*	NATO Task Group on Psychosocial, Organizational, and Cultural Aspects of Terrorism*
FBI Behavioral Science Unit	**Academic and Research Laboratories**
Futures Working Group*	DHS Center of Excellence: Center for Defense Systems Research at the University of Texas, El Paso
Department of Defense	
Center for Technology and National Security Policy	Center for Homeland Defense and Security at the Naval Postgraduate School*
Office of the Secretary of Defense, Cost Assessment and Program Evaluation	Draper Laboratory*
Defense Threat Reduction Agency	International Center for the Study of Terrorism at Pennsylvania State University
Department of Homeland Security*	
Transportation Security Administration	Los Alamos National Laboratory
Science and Technology Directorate: Human Factors Division	Center for the Scientific Analysis of Emerging Threats*
Future Attribute Screening Technology (FAST)	National Research Council
Department of the Navy	Sandia National Laboratories
Naval Research Laboratory (NRL)	National Academy of Sciences
Threat Management Unit	SRI International*
Government Accountability Office	**Private Organizations and Public-Private Sector Efforts**
Los Angeles Mayor's Office	Hughes Research Laboratories, U.S.
Blue Ribbon Panel on Airport Security	NICE Systems, Israel
U.S. Postal Service	Park Assist, U.S.
Threat Assessment Team Task Force	Total Domain Awareness, U.S. (Microsoft/NYPD)
	Facebook, U.S.*
	Palantir, U.S.*
	Google, U.S.
	Shot Spotter, U.S

NOTE: Asterisks indicate where we held interviews.

References and Cases to Support Historic Examples

The following long table gives succinct descriptions of various attacks. An asterisk (*) indicates that law enforcement or intervening opposition thwarted the attack. A plus (+) sign indicates that the attack was primarily coordinated and carried out by a "lone wolf." The information here comes from a variety of public sources, including newspaper accounts, which vary in reliability. For each attack, we mention specifically only one or a few particular sources that might be useful to a reader interested in pursuing the cases. Information continues to emerge on many of the cases, some of it contradicting earlier accounts in the news and online media.

Table B.1
Historical Cases

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
The 1993 World Trade Center bombing (1993)	The attackers entered the United States (with no plans or orders) aboard a flight from Pakistan. One of the attackers was taken into custody when customs agents found a "terror kit" in his luggage, consisting of various videos and bomb-making manuals as well as various forged passports from the Middle East and Europe.
	With Osama bin Laden's endorsement, the World Trade Center became the prime target of the attack because of its association with free enterprise and global trade.
	Using an alias, one of the attackers rented a storage unit to store bomb-making materials. Also, the attacker ordered chemicals, including urea and nitric acid, to be delivered to the storage facility, paying more than $3,000 in cash.
	A day before the attack, one of the attackers reported a van (eventually used for the bomb) stolen from the Pathmark Plaza shopping center in New Jersey. Earlier on the same day, the attackers made a phone call in a last attempt to acquire even more compressed hydrogen tanks. (Illustrative source: Reeve, 1999.)
Oklahoma City bombing (1995+)	In the years prior to the bombing, Timothy McVeigh began checking-out and buying anti-government books and pamphlets. He would share the information and literature with "anyone who would listen."
	McVeigh frequented the Marietta Aggregates quarry in Marion, Kansas, to steal and test bomb components. Terry Nichols (McVeigh's partner) purchased 2,000 pounds of ammonium nitrate at a farm co-op. McVeigh, disguised as a biker, drove to a Texas racetrack and bought $2,775 worth of racing fuel.
	McVeigh wanted to target a federal institution and picked the Alfred P. Murrah Federal Building in Oklahoma City because he noted the huge glass facade, which would maximize injuries. McVeigh wanted also believed there would be widespread press coverage of his attack and thought an image of a devastated building would have "a profound effect on those who saw it." McVeigh drove past the building several times before the day of the bombing.
	At about 8:50 a.m. on the day of the attack, McVeigh parked the truck, loaded with explosives, right below the tinted windows of the America's Kids Day Care Center, locked the door, and walked away from the building. (Illustrative sources: Smith, Damphousse, and Roberts, 2006; Michel and Herbeck, 2001.)

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Would-be LAX bomber (The Millennium Conspiracy) (1999)* An individual associated with an organization.	Ahmed Ressam (the attacker) became friends with Raouf Hannachi, an al Qaeda member who had trained for jihad at a camp in Afghanistan. He told Ressam about the experience and jihad, encouraged him to train as well, and ultimately arranged a trip to the camp for him.
	Ressam traveled, using a fraudulent passport, to Pakistan. There, he contacted al Qaeda leader Abu Zubaida, who was in charge of the Afghan terrorist training camps funded and organized by Osama bin Laden. Abu Zubaida approved him and arranged for him to be transported over the Khyber Pass into Afghanistan.
	Ressam returned to Canada with $12,000 in cash he had obtained in Afghanistan to fund the attack, as well as chemical substances, and a notebook with explosives concoction instructions. He also obtained electronics with which he built detonators and timing devices.
	While he was on his way to Los Angeles International Airport from Canada, U.S. Customs searched Ressam's car, saying later that Ressam was acting "hinky." (Illustrative source: Smith, Damphousse, and Roberts, 2006).

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
9/11 (2001)	Osama bin Laden claimed in 2004 that the idea of destroying the towers had first occurred to him in 1982, when he witnessed Israel's bombardment of high-rise apartment buildings during the invasion of Lebanon (Al Jazeera, 2004).
	Hijackers were recruited in various ways, including from local universities and mosques. Bin Laden selected Nawaf al-Hazmi and Khalid al-Mihdhar, both experienced jihadists who had fought in Bosnia. Hazmi and Mihdhar arrived in the United States in mid-January 2000. In late 1999, a group of men from Hamburg, Germany arrived in Afghanistan, including Mohamed Atta, Marwan al-Shehhi, Ziad Jarrah, and Ramzi bin al-Shibh. Bin Laden selected these men because they were educated, could speak English, and had experience living in the West. New recruits were routinely screened for special skills. For example, Hani Hanjour already had a commercial pilot's license.
	In early 1999, bin Laden approved Khalid Sheikh Mohammed's going forward with organizing the plot. A series of meetings occurred in early 1999, involving Mohammed, bin Laden, and his deputy Mohammed Atef. Atef provided operational support for the plot, including target selections and travel arrangements for the hijackers.
	The terrorists took flight training before the attacks in Florida and Arizona.
	In July 2001, Atta met with bin al-Shibh in Spain, where they coordinated details such as final targets. Large planes with long flights were intentionally selected for hijacking because they would be heavily fueled. Bin al-Shibh also passed along bin Laden's wish for the attacks to be carried out as soon as possible. On August 29th, Atta gave the date for the attacks to Bin al-Shibh, who ordered active cells in Europe and the United States to evacuate. Bin Laden was told on September 6.
	Days before the planned attacks, hijackers sent notes to loved ones and engaged in religious practices.
	The ticket agent who served two of the hijackers said on the morning of the attacks he was suspicious of Mohamed Atta and Abdulaziz Alomar. Atta's demeanor and his angry-looking eyes made the ticket agent think twice.
	Early on the morning of September 11, 2001, 19 hijackers, coordinating with each other, took control of four commercial airliners after takeoffs from Boston, Massachusetts; Newark, New Jersey; and Washington, D.C. Hijackers told passengers they had bombs, but the FBI found no traces of explosives at the crash sites. The 9/11 Commission concluded the bombs were most likely fake.
	Immediately after the attacks, the Federal Bureau of Investigation started PENTTBOM, the largest criminal inquiry in the history of the United States.

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
9/11 (2001) continued	Shortly before the U.S. presidential election in 2004, in a taped statement, bin Laden publicly acknowledged al Qaeda's involvement in the attacks on the United States and admitted his direct link to the attacks (Al Jazeera, 2004).
	Many references exist on 9/11, including National Commission on Terrorist Attacks (2004) and, for bin Laden's speech, Al Jazeera (2004).
London bombings (2005)	The bombers joined and recruited each other while traveling to Pakistan to attend jihadist training. It was there that they learned the methods and techniques of bomb-making. Two of the bombers made videotapes describing their reasons for becoming what they called "soldiers." Their extremist views can also be traced back to blog writings and email communications years prior to the bombings.
	Relatives in Pakistan said that Shehzad Tanweer had boasted of wanting to die as a "holy warrior" and of the appeal of Osama bin Laden.
	The suicide bombers studied the layout of the underground system and planned their attacks to occur where the most civilian lives would be lost. Days before the attacks the bombers surveyed the locations they would eventual detonate the bombs.
	Tanweer, Mahammad Sidique Khan, Hasib Hussain, and Germaine Lindsay picked up the bombs from a house in the Burley area of Leeds, hiding them in large rucksacks.
	The attackers were in constant communication with each other up until the final decision to implement the planned attack. Khan postponed the event from July 6 because he had to take his pregnant wife to the hospital.
	The four men caught a train to London King's Cross railway station on the morning of July 7. Still communicating via mobile phone, the attacks split ways to their assigned positions. Once in position, the attackers detonated four bombs, three in quick succession aboard London Underground trains across the city and, later, a fourth on a double-decker bus in Tavistock Square. Fifty-two innocent people, and the four bombers, were killed in the attacks, and over 700 more were injured.
	Britain's security forces immediately increased security. Police sniper units began as many as a dozen al Qaeda suspects in Britain. The covert armed teams were ordered to shoot to kill if surveillance suggested that a terror suspect was carrying a bomb and he refused to surrender if challenged.
	Syrian-born cleric Omar Bakri Muhammad vowed in December 2008 that if Western governments did not change their policies, Muslims would give them "a 9/11, day after day after day."
	Some key references are London Regional Resilience Forum (2006), Intelligence and Security Committee (2006), and, for a journalistic account, Sciolino and van Natta (2005).

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Mumbai attacks (2008)	In October, U.S. intelligence agencies warned the chairman of the company that owns the Taj Mahal Palace Hotel that there would be a terrorist attack on the hotel. Security was increased, but was removed soon after.
	Group members received a huge cache of AK-47s and the explosive RDX, which were to be used for the attacks.
	Group members also obtained drugs, like cocaine, which they used in the attack to maintain stamina. Some of the attackers took steroids, also.
	The attackers had planned the attack several months ahead of time and knew some areas well enough for the attackers to vanish, and reappear after security forces had left. The attackers had used Google Earth to familiarize themselves with the locations of buildings used in the attacks.
	The attackers allegedly received reconnaissance assistance before the attacks. In the days before the attack, they were in communication with their supporters, and coordinated with each other via email and telephone positioning plans.
	Arriving ashore in Colaba on inflatable speedboats, the ten attackers were queried by local fishermen but told the fishermen to mind their own business before they split up into two groups heading in different ways. The fishermen reported this, but the police did not act.
	The men then carried out 11 coordinated shooting and bombing attacks across Mumbai. The men split up to attack different locations and breached security via multiple points of attack.
	While the attacks where ongoing, Deccan Mujahadeen sent messages to media outlets claiming responsibility for the attacks.
	The terrorists used Google Earth to plan the attacks, as in locating strategic positions for hiding from authorities.
	Several journalistic references are useful regarding the Mumbai attack (Moreau and Mazumdar, 2008; Sengupta, 2009). See also Rotella (2012).

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Fort Hood shooting (2009). An individual influenced by an organization.	Major Malik Hasan, a U.S. Army psychiatrist, became increasingly radicalized by perhaps 2005. He also alarmed medical colleagues who worried about his mental health. His email communications with a radical cleric also showed anger toward the American government.

Hasan entered the Guns Galore store in Killeen on July 31, 2009, and purchased the FN Five-Seven semi-automatic pistol that he was to use in the attack at Fort Hood after asking for the most technologically advanced weapon on the market with a high capacity. He visited the store regularly to buy extra magazines, along with hundreds of rounds of ammunition. In the weeks prior to the attack, Hasan visited an outdoor shooting range in Florence, where he allegedly became adept at hitting silhouette targets at distances of up to 100 yards (Brown and Granczyk, 2010).

At 1:34 pm local time, Hasan entered his workplace, the Soldier Readiness Processing Center, where personnel receive routine medical treatment before and after deployment. Hasan sat at an empty table and bowed his head for several seconds, after which he stood up, shouted "Allahu Akbar!" and opened fire.

Much is available about the Hasan case, with more emerging through court processes, but some contemporary journalistic sources were Esposito, Abraham, and Schwartz (2009), McKinley and Dao (2009), and Zwerdling (2009). |

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Christmas bomber (2009). An individual working for an organization.	The attacker's father reported to two CIA officers at the U.S. Embassy in Abuja, Nigeria, regarding his son's "extreme religious views," and told the embassy that Umar Farouk Abdulmutallab might be in Yemen. The suspect's name was added in November 2009 to the United States' 550,000-name Terrorist Identities Datamart Environment, a database of the U.S. National Counterterrorism Center. It was not added, however, to the FBI's 400,000-name Terrorist Screening Database, the terror watch list that feeds both the 14,000-name Secondary Screening Selectee list and the United States' 4,000-name No Fly List, nor was his U.S. visa revoked.
	Abdulmutallab had purchased his ticket with cash in Ghana on December 16 and obtained Pentaerythritol tetranitrate (PETN) with triacetone triperoxide (TATP), which he used to assembled the bomb.
	Abdulmutallab told authorities he had been directed by al Qaeda and coordinated with them the days before the attack. Indeed, he had spent several days with Anwar al-Awlaki, who arranged for him to work with the bomb maker who constructed the underwear bomb. Awlaki specified that the attack must be on an American target, but otherwise left the choice of target and flight up to Abdulmutallab (Savage, 2012).
	On Christmas Day, 2009, Abdulmutallab traveled from Ghana to Amsterdam, where he boarded Northwest Airlines Flight 253 en route to Detroit. Abdulmutallab spent about 20 minutes in the bathroom as the flight approached Detroit, and then covered himself with a blanket after returning to his seat. Other passengers then heard popping noises, smelled a foul odor, and some saw Abdulmutallab's trouser leg and the wall of the plane on fire. Fellow passenger Jasper Schuringa, a Dutch film director, jumped on Abdulmutallab and subdued him as flight attendants used fire extinguishers to douse the flames.
	Al Qaeda in the Arabian Peninsula, the organization's affiliate in Yemen, claimed responsibility for the attack, describing it as revenge for the United States' role in a Yemeni military offensive against al Qaeda in that country.
	For journalistic accounts, see Hosenball, Isikoff, and Thomas (2010), Shane and Lipton (2009), Schmitt and Liption (2009), and Savage (2012).

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Dallas Skyscraper (2009). A lone wolf captured in a sting operation.	Hosam ("Sam") Smadi, unaware he was under continuous surveillance, joined a social network, and then recruited individuals who said they would be willing to be jihadi warriors and attack America. The "sleeper cell" he created was composed of all federal agents.
	The agents in his "sleeper cell" had supplied him with inert chemical, so his bomb had not posed a real threat.
	Smadi initially wanted to target Dallas/Fort Worth International Airport, according to the arrest affidavit. On July 16th, he contacted one of the undercover FBI agents and said he changed his mind about the target, according to the document. Smadi allegedly decided the airport was not a viable target because the security was too strong. Smadi then allegedly told the undercover agent he wanted to target a larger building containing the bank.
	Smadi allegedly conducted reconnaissance of the building and told the undercover agent he had located a bathroom on the basement level that would be a good location to "plant a bomb." Smadi told the agent the bathroom had a locking door and a drop ceiling that could be accessed by standing on the toilet seat, according to the arrest affidavit.
	In communication with his "sleeper cell," Smadi went over final plans and location the day before the planned attack
	Smadi drove to Dallas to meet the undercover agent and got into the Ford Explorer that contained what he believed was a weapon of mass destruction. He then drove through downtown Dallas, entering the parking garage under Fountain Place building and parked the vehicle. Smadi attempted to ignite and detonate the explosive device by setting the device's timer and flipping its power switch before leaving the garage on foot.
	Smadi then walked over to the undercover officer and got into another vehicle. They drove several blocks away so that Smadi could remotely detonate the bomb via cell phone. The agent offered Smadi earplugs, but he declined, indicating he wanted to hear the blast. Smadi then dialed the cell phone, believing it would detonate the bomb. The phone number Smadi dialed rang to the phone of law enforcement officials, and he was arrested by the FBI Joint Terrorism Task Force.
	Some references on the case are U.S. District Court (2009) and Goldstein (2009).

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
CIA-team suicide bomber (double agent) (2009+)	Humam Khalil Mohammed had used the online persona Abu Dujana al-Khorasani and was an influential jihadi voice on the Web. In many of the posts under his online persona, Mohammed used elusive language filled with references to literature and the Koran to describe his support for violent opposition to the U.S.-led wars in Iraq and Afghanistan.
	He wrote in one posting, "When a fighter for God kills a U.S. soldier on the corner of a tank, the supporters of Jihad have killed tens of thousands of Americans through their connection."
	The bomber had been recruited by the Jordanian intelligence service and taken to Afghanistan to infiltrate al Qaeda by posing as a foreign jihadi.
	But in a deadly turnabout, the supposed informant strapped explosives to his body and blew himself up at a meeting December 30, 2009, at the CIA's Forward Operating Base Chapman in the southeastern province of Khost. The bomber was not closely searched because of his perceived value as someone who could lead American forces to senior al Qaeda leaders.
	See Oppell, Mazzetti, and Mekhennet (2010) for a journalistic account.
D.C. subway bombing plot (2010). An individual influenced by an organization's information, captured in a sting operation.	Farooque Ahmed was tracked after he began visiting Islamic extremist websites and attempted to recruit members of the community to participate with him in an attack against America.
	Ahmed took photographs and video of potential targets, including Metro stations and hotels. He told undercover agents the best time to stage an attack would be between 4 and 5 p.m.—basically rush hour—to have the most casualties and that he would use rolling suitcases instead of backpacks so he wouldn't attract undue attention.
	See Tavernise and Schmitt (2010) and CBS News (2010) for accounts at the time.

Table B.1—Continued

Attack	Examples of Phase Activities, Low-Level Activities, and Behavioral Indicators
Zachary Chesser (online Islamist radicalization) (2011+)	Zachary Chesser quit his job at a Blockbuster video store because "he objected to working at a place that rented videos featuring naked women." His parents described an increasingly hostile home environment in which Chesser would institute strict rules to enforce what he believed to be proper Islamic traditions. By August 2008, he had moved out of his mother's house in Virginia because, according to his father, "his mom's relationship with her live-in partner . . . violated his Islamic beliefs."

By fall of 2008, Chesser had become a full-fledged believer in the ideology of violent Islamist extremism and was searching for other like-minded individuals. He gravitated toward the Internet to find them.

In a series of posts entitled "Counter Counter Terrorism," Chesser outlined ways the violent Islamist extremist movement could win an ideological struggle—the so-called war of ideas—against the West.

Three weeks before his arrest in July 2010, Chesser authored a 25-page document entitled "Raising Al-Qaa'ida: A Look Into the Long Term Obligations of the Global Jihaad Movement." The piece gave suggestions and best practices for engaging Muslims who have not joined the violent Islamist extremist movement.

After arrest, Chesser wrote a letter to U.S. Senate Committee on Homeland Security and Governmental Affairs staff. In the letter, he promoted the idea of better understanding between violent Islamists and the government. He advocated for an online discussion board between counterterrorism policymakers and Islamist followers where debate and common ground could be found.

See, e.g., U.S. Senate Committee on Homeland Security and Governmental Affairs (2012). |

References and Cases to Support Indicator Tables

The following tables gives examples for major tables in the text. It draws on a range of publically available information, not all of which is equally authoritative. We also mention some specific references that readers might find useful in following up on the cases. We have typically cited relatively more scholarly sources or good journalistic accounts.

Table C.1
References and Cases for Developing Intent (Table 2.1)

Indicator	Examples	Instance Observed (Attack/Indicator)
Reveal hatred, prejudice, trauma, or shame	Early childhood trauma, bullying in high school; traumatic loss or separation from spouse or family members; open discriminatory statements; cruelty to animals, peers, siblings, etc.	Oklahoma City Bombing. Timothy McVeigh was a target of bullying at high school, suffered through the difficult divorce of his parents and the trauma of his mother leaving him, grew up in Buffalo during a very difficult economic period, and was heavily exposed to far-right attitudes.
Exhibit motivation for prestige, glory, status	Boasting or bragging about radical or criminal associations, violent acts, or intentions online or face-to-face.	
Approach experiences and decisions through ideological framework	Ideological statements to others in face-to-face communication or online, ideological content in school projects, "black-and-white" thinking.	Fort Hood Shooting. Emails between the shooter and a Yemen-based cleric (al-Awlaki) showed religious concerns, questions about martyrdom, and, anger toward the American government.

Table C.1—continued

Indicator	Examples	Instance Observed (Attack/Indicator)
Explore different organizations and strategies	Visits to political or radical websites, meetings, protests. Writing or engaging others on political or radical websites; attending meetings, protests.	CIA-team suicide bomber. Humam Khalil Mohammed wrote posts on the internet that used elusive language filled with references to literature and the Koran to describe his support for violent opposition to the United States.
Pure adherence to organizational standards	Attention to uniform or other ideological markers, strict following of religious or other group rules.	
Efforts to reinforce practices and beliefs in others	Demands on peers, new organizational recruits, family to follow religious or other group membership standards strictly.	Zachary Chesser. Chesser would institute strict rules on his family to enforce what he believed to be proper Islamic traditions.
Indications that individual is reinforcing commitment	Increasing exclusivity in social connections (connected vs. outside of radical organization), intensification of radical statements online or in face-to-face communication, denigration of individual outside organization or radical movement.	Christmas Bomber. The attacker informed his father of his radicalization and separation clearly enough so that the father reported his concerns to American intelligence.
New connections in social network to known terrorist elements	Facebook or other online networking (Twitter, Yahoo chat, etc.) associations to known operatives or recruiters, direct face-to-face social contact and meetings with these individuals, cell phone call records to known "red network."	
Seeking out, reading, or posting radical content	Visits to radical websites, collection or display of recruitment or ideological materials, chatting and sharing materials online or in face-to-face (political meeting) settings.	Zachary Chesser. In a series of online posts entitled "Counter Counter Terrorism," Chesser outlined ways the violent Islamist extremist movement could win an ideological struggle—the so-called "war of ideas"—against the West.

Table C.1—continued

Indicator	Examples	Instance Observed (Attack/Indicator)
Attendance at radical mosques or events	Reports by intel collectors embedded at radical meetings, protests, or other locations, interviews or interrogation reports from others attending such events/locations, CCTV or other imaging feeds monitoring events and locations.	The Millennium Conspiracy. The attacker became friends with Raouf Hannachi, an al Qaeda member who served as the muezzin at Montreal's Assuna Mosque. He began regularly attending the mosque.
Changes in behavior at school or home	Frequent unexplained absences, rapid shifts in mood, changes in patterns of interactions with family members regarding political issues, gender norms, or other radical group concerns.	

NOTE: Some suggested references include Vossekuil (2002), U.S. Senate Committee on Homeland Security and Governmental Affairs (2012), Hudson (1999), Pedahzur, Perliger, and Weinberg (2003), and Smith, Damphousse, and Roberts (2006).

Table C.2
References and Cases for Planning (Table 3.1)

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Seek information or guidance on construction of weaponry and explosives	Online searching for weapons construction manuals, conversations with weapons or explosives professional or retailers for firearms or explosive precursors, contact with traffickers and facilitators.	Fort Hood Shooting.
Visit known training camps and seek aviation or marksmanship training	Online registration, payment, and attendance at training programs; travel to known terrorist training locations.	The Millennium Conspiracy.
Acquire dual-use electronics, explosives, ignition devices	Acquire dual-use electronics, explosives, and ignition devices.	Oklahoma City Bombing.
Surveillance of target	Suspicious behavior near high-value human, infrastructural, or MASCAS targets (lingering, observing, photography, rapid departure if approached, etc.).	Dallas Skyscraper Bombing.
Dry runs to simulate and practice attack	Operational movements without weaponry or by alternate personnel.	
Actions that provoke or test security responses near target	Attempts to penetrate security perimeter, interactions with guards, attempts to approach forbidden or restricted areas.	CIA-Team Suicide Bomber. The bomber posed as an informant willing to provide information on the whereabouts of top terrorist leaders.
Release of information or discussions of how to most harm or influence target population (e.g., targets of symbolic value vs. mass casualty)	Strategic communication with operatives through public or private targeting directives.	9/11. Middlemen passed along the wishes and directives of Osama bin Ladin to the attackers.

Table C.2—continued

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Chemical or explosive hazards from experimentation	Injuries (burns, etc.) from experimentation and development, smells, smoke, explosions, or other physical signs from residence.	
Purchasing for explosive pre-cursors or retail sales of firearms	Purchase records or intelligence trail related to acquisition.	Fort Hood Shooting.
Clandestine movement or camouflage around potential targets	Attempts to conceal presence at location under cover of night or by blending in with crowds, attempts to fool or disable monitoring devices or security monitoring personnel.	

NOTE: Some useful references include U.S. Senate Committee on Homeland Security and Governmental Affairs (2012), National Commission on Terrorist Attacks (2004), and Smith, Damphousse, and Roberts (2006).

Table C.3
References and Cases for Pre-Execution Activities (Table 4.1)

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Isolation from nonterrorist elements of social network	"Going dark" and disappearing from family and peer social contacts.	
Increased communication with terrorist elements	Intensification of cell phone, online, or face-to-face contact with radical group members.	The Millennium Conspiracy. In the months before the attack, Ahmed Ressam attended one of the Afghan terrorist training camps funded and organized by Osama bin Laden.
Rituals or other actions to motivate self and co-attackers	Martyrdom videos, shaving head, praying, preparation of body with oils, wearing special outfits, jewelry, or other adornment.	
Inconsistent responses to questioning	Details of reasons for travel or presence at location differ across multiple conversations or interviews/ interrogations.	
Nonverbal signals of deception and lying	Evidence of stress and cognitive load.	Millennium Conspiracy. Customs officials checking Ahmed Ressam's car later said that he acted "hinky."
Hesitation near target	Slowed gait, increased attention, multiple changes in direction.	
Accelerated heart rate, sweaty palms, thermal indicators	Self-explanatory.	
Micro-expressions of fear, hostility, deception, detachment	Fear, anger, Duchenne smile ("duper's delight").	9/11. The ticket agent who served two hijackers was suspicious when they rushed to their flight out of Portland. Mohamed Atta's demeanor and his angry-looking eyes and the pair's first-class, one-way tickets to Los Angeles made the ticket agent think twice.

Table C.3—continued

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Indicators of instrumental aggression	Slowed heart rate, orienting and focused attention, filtering of extraneous signals.	
Kinetic (body movement) patterns indicating hostile intent, attempted clandestine movement, or carrying weapon/bomb	Changes in gait indicating burdened with weapon or explosives, gross motor movements (gait or otherwise) indicating anger or readiness for violence, patterns of footsteps indicating attempted evasion.	Christmas Bomber. Abdulmutallab spent about 20 minutes in the bathroom as the flight approached Detroit, and then covered himself with a blanket after returning to his seat.

NOTE: Some suggested sources include Mullaney and Costigan (2010) and Butterworth, Dolev, and Jenkins (2012).

Table C.4
References and Cases for Execution (Table 5.1)

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Sequence of actions in attack (i.e., 2 suicide bombers detonate, 3 run into nearby structure) portending actions in next stage of attack	Deploying to multiple different positions, employing a mixture of suicide bombers and shooters to penetrate checkpoints.	London Subway Bombings. The four men caught a train to King's Cross railway station. Still communicating via mobile phone, they split to their assigned positions and then detonated four bombs.
Specifics during and around attack target or individuals singled out) portending intent of group or clues about future attack types or locations	Targeting of specific types of individuals (e.g., Israeli tourists, soldiers, diplomats); destruction of targets with symbolic or strategic importance; destruction of easy access or "soft targets."	The 1993 World Trade Center Bombing. Bombers picked location because of symbolic association with free enterprise and American-endorsed capitalism.
Driving in car overly packed with fighting-aged males (sometimes with heavily weighted suspension due to munitions)	Excessively crowded or weighted/laden vehicles; large groups of fighting-aged males traveling together.	
Running checkpoints or security barriers	Aggressive or erratic driving behaviors near checkpoints; other unusual patterns of movement near checkpoints if on foot or other mode of transport.	The Millennium Conspiracy. Customs officials checking Ahmed Ressam's car later said that he acted "hinky."
Splitting into groups (signals multiple points of attack)	Deployment to different positions.	Mumbai Attacks. On arriving by inflatable speedboats, the men deployed to multiple assigned locations throughout breach security and carry out attacks.

Table C.4—continued

Activities/ Indicators	Examples	Instance Observed (Attack/Indicator)
Running into buildings or ascending structures with cover or line-of-sight (signals intent to engage in direct or indirect fire on targets)	Occupying positions with cover or line of sight to primary target.	
Shaping the population composition of the target area	Waiting for target groups to arrive or population concentration to increase , asking certain nontargeted individuals to leave or waiting until they depart. Waiting until ambulances or emergency support arrives to initiate secondary attack.	Oklahoma City Bombing. Timothy McVeigh purposefully waited to enter the building until the middle of the morning to maximize the number of people in the building at detonation.
Intelligence	Collecting final indicators on location of target before attack (while armed and ready) through verbal inquiry or line of sight.	
Interaction with security personnel	Final collection of information through interactions with security personnel; manipulating personnel by duping or feint attacks to help shape target (e.g., direct targets towards a central location).	CIA-Team Suicide Bomber. The bomber was not closely searched because of his perceived value as someone who could lead American forces to senior al Qaeda leaders.
Preparatory attack or feints	Softening checkpoints or primary target before primary attack; Decoy attacks to lure security personnel away.	
The real attack—detonation	Self-explanatory.	

SOME useful references are Mullaney and Costigan (2010), BBC News Special Reports (2008), Butterworth, Dolev, and Jenkins (2008), Heger (2102), and Oppel, Mazzetti, and Mekhennet (2010).

Table C.5
References and Cases for Aftermath (Table 5.2)

Indicator(s)	Activities	Instance Observed (Attack/Indicator)
Take responsibility or lay blame	Post Twitter feeds, and video and audio announcements. If attack unsuccessful or botched, "spin" to minimize loss of allegiance among adherents and maintain wider population of support.	Mumbai Attacks. While attacks where ongoing, Deccan Mujahadeen sent messages to media outlets claiming responsibility for attacks.
Calls to action	Use stories and images of attack to mobilize followers, including production and distribution of recruitment videos, flyers, and other materials. If attack is unsuccessful, spin to motivate subsequent efforts.	London Subway Bombing. Days after the attacks, planners stated more attacks were needed until Western governments changed policies.
Idolize attackers	Release still images or recordings of attacker to celebrate martyrdom (Pape, 2003). Develop recruitment videos about attackers. Compensate and appease family and friends (including praise for attackers and material gifts).	Hamas Suicide bombings. A major recruiting factor was economic (such as family support after suicide attack). Also, suicide bombers were idolized like "rock stars."
Communication with HQ about operational success or failure	Feed operational details into decision-making about future CONOPs, weapons purchases, etc.	
Silence (kill, threaten, etc.) those with protected information about the attack	Target operatives, collaborators, or facilitators who may have regrets or play into the hands of authorities.	
Clean evidence from safe houses and planning areas	Gather materials for future efforts, hide evidence that would give away tactics and procedures.	New York Subway Bombing. An associate of the attacker attempted to destroy evidence of homemade bombs created to detonate on Manhattan subway cars.

Table C.5—continued

Indicator(s)	Activities	Instance Observed (Attack/Indicator)
Report on troubles— including interdiction or interruption of attack.	Provide details on possible collection measures or resistance/ retaliation at target site.	
Develop new tools and CONOPs	Discuss different target types or sites, purchase or develop new weapons (including explosives and detonation devices), work on new disguises, transportation mechanisms, etc., as part of future CONOPs.	

NOTE: Some useful references are Ashworth, Clinton, Meirowitz, and Ramsay (2008), Hafez (2007), Wells and Horowitz (2007), Oppel, Mazetti, and Mekhennet (2010), and London Regional Resilience Forum (2006).

Information Fusion Methods

Chapters in the main report identified activities that might suggest a greater likelihood of malign intent than "normal," thereby suggesting the need for closer scrutiny of an individual or group. For example, an activity associated with "developing intent" is "recruitment and joining," i.e., an individual actively seeking membership in some radical group. Indicators (Table 2.1) might be attendance at meetings or visiting radical websites. If we observe an individual attending one meeting a week for six months and also observe that he has accessed radical websites, how likely is it that he is contemplating membership in a radical group? How likely is it that he is "developing intent"? That is, how do we combine knowledge from the two indicators? Fusion is the process of combining information from various sources (similar and disparate in character) with the intention of obtaining a better composite of that being studied (Perry and Moffatt, 2004).

Actually, fusion occurs at many levels of detail. A polygraph test, for example, is based on combining information from several physiological measurements. A security official may single out an individual with an unusual gait for secondary screening after observing that he or she was also sweating excessively. Implicitly, data are being fused. In this appendix, however, by "fusion" or "information fusion" we have in mind combining higher-level information, as depicted in Figures 1.2 and 1.3: (1) fusing activity-level information within a given phase and (2) fusing information across phases to form an overall (relative) likelihood of hostile intent. This chapter reviews various methods that might be used, at each of these levels, to fuse evidence about hostile intent.

Two types of fusion are usually mentioned: (1) fusion at a given time, combining several indicator reports from disparate sources, and (2) fusion over time, combining the result with previous reports to obtain an updated estimate of likelihood (Llinas et al., 1998). The second (updating) is the better understood (Darilek et al., 2001). The first is more problematic because reports are often a mix of quantitative and qualitative information, without an obvious mechanism for combining them. A third type of fusion is important for our context and is a combination of the previous two: updating of reports from multiple sensors and sources over time with those reports often being very disparate in character (heterogeneous).*

All of these aspects of fusion require mathematical algorithms that use subjective assessments of conditional probabilities, $P(A|B)$, where A is the activity and B is the indicator. For example, we ask, "What is the likelihood that an individual contemplates joining a radical group given that he was observed attending meetings of that group?"

An important question is "How do we know whom to monitor?" One important method for identifying individuals or groups to monitor is "data mining." We do not discuss it in detail here because we do do not see it as information fusion (although others do), but it is very important, for reasons discussed in Chapter Six. Data mining attempts to discover patterns in large data sets (see Witten, Eiber, and Hall, 2011). It draws on methods from artificial intelligence, machine learning, statistics, and database systems. For our context, "clustering" methods may be used to detect patterns of threatening behavior, such as membership in groups known to be hostile and participation in social networks with people known to be hostile. Another data-mining method, "anomaly detection," searches for unusual records, such as worrisome blogs by individuals or groups that support radical positions. As the report notes repeatedly, our focus is on detecting attacks rather than, say, broad "behavioral monitoring" of the population to

---

* A generalization is updating over items of evidence, which might or might not be ordered chronologically. As a now-familiar example, cold-case investigations may "update" assessments of an individual's guilt of a past crime by folding in DNA testing using a sample from long ago and a modern-day testimony of some prison inmate with alleged knowledge of the crime.

find patterns of behavior that might relate somehow to terrorism or some other subversive activity. Such behavioral monitoring raises major civil-liberties concerns (Perry and Vest, 2008).

In what follows, we describe the general form of the problem and then a number of specific methods of evidential reasoning drawn from the literature.

The Information Fusion Problem in Detecting Hostile Intent

Our concern is with fusing information relevant to *possible* hostile intent. We refer to the activities and phases described in Chapter One. For a given phase, we let $A = \{A_1, A_2, . . ., A_3\}$ be the set of activities associated with a phase such *developing intent*.[*†] Likewise, we let $\mathbf{I}_j = \{I_{j,1}, I_{j,2}, . . . , I_{j,m}\}$ be the set of indicators of activity A_j. Thus, the relevant conditional probability is $P_t(A_j|I_{j,1}, I_{j,2}, . . . , I_{j,m}) = P_t(A_j|\mathbf{I}_j)$; that is, we ask, "What is the likelihood that an individual or group is engaged in activity A_j at time t, given information from all possible indicators, $\mathbf{I}_j$?" The time component accounts for both updating an estimate over time and fusing several reports within the same time period.

Although this formulation is simple, its implementation is anything but. To illustrate, assume a single activity with three indicators, so that we get the model $P(A|I_1, I_2, I_3)$. For purposes of this first illustration, we omit the time dimension—assuming that all indicator reports arrive at the same time. If all three indicators appear (i.e., if one observer

---

[*]   For readers rusty in set notation, $S_i \in \mathbf{S} = \{S_1, S_2\}$ means that S_i is a member of the set $\mathbf{S}$ that has elements S_1 and S_2. The notations $A \cup B$ and $A \cap B$ refer to the union of sets A and B and the intersection of sets A and B, respectively. The empty set is denoted φ. It is also common to refer to propositions that are either true or not true, in which case $\overline{P}$ refers to the proposition being *not* true.

[†]   As described in Chapter One, the underlying model of indicators, activities, and activity classes called phases is only weakly hierarchical. A given indicator, for example, may have implications for more than one activity and phase. We ignore such subtleties in this appendix.

saw him attend a radical group meeting, another monitored his accessing the group's website, and a third witnessed his paying a membership fee), how likely is it that activity *A* is in progress? More generally, we would expect only some of the indicators to appear. Given three indicators, there exist possible reporting cases (e.g., all three report, only the first and third report, . . . , none report). We need a subjective assessment of the effect on the likelihood of *A* for each of these cases.

The problem becomes even more complex because the reports on indicators will often occur in various orders over time. For example, if we monitor activities of a suspect in the example above, we may get a report that he attended a radical group meeting on Tuesday and another report that he attended the same meeting the following week. It is also possible that after the first meeting he accessed the group's website. The more general model needed, then, is of the form $P_t(A|I_j) = f(P_{t-1}(A), I_j)$. That is, the current estimate, $P_t(A|I_j)$, is a function of both the previous estimate (the "prior") and the indicator report(s), I_j. This gives us an updating algorithm that accounts for both successive and simultaneous indicator reports.[*] The various fusion methods below can be considered alternative constructs of this function.

Bayesian Updating

Perhaps the most common fusion method is Baysian updating, which is based on Bayes rule (Feller, 1950; Mood and Graybill, 1963; Raiffa, 1968; Stone, Barlow, and Corwin, 1999).[†] That rule can easily be derived from the definition of conditional probability.

---

[*]   References to time relate to the time of a fusion estimate, which may use evidence about activities at quite a number of previous times. Sometimes, fusion at time *t* will include newly recognized evidence about old events (as when DNA analysis is made in 2012 using sample data collected years early at the place of a crime).

[†]   Bayes' rule is named for the Reverend Thomas Bayes (1701–1761), an English Presbyterian minister who first suggested using the rule to update beliefs. Being a Christian cleric, Bayes was interested in strengthening people's belief in God through evidence presented in the physical world (Hamburg, 1983). He first published his theorem in 1783 in an essay reprinted more recently (Press, 1989). The references we include above are but a few of the

If $P(A,B)$ is the probability that *both* A and B are true, i.e., the "joint probability," then it implies two relations:

$$P(A,B) = P(A|B)P(B) = P(B|A)P(A).$$

If we solve for the conditional probability, $P(A|B)$, we obtain

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}.$$

Next, we can replace the denominator by expressing the *marginal* probability $P(B)$ in terms of conditional probabilities using the notation $\overline{A}$ to mean the condition that A is not true. This is also called the negation of A. Since

$$P(B) = P(B|A)P(A) + P(B|\overline{A})P(\overline{A}),$$

we can use this in the equation above to obtain Bayes' rule:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B|A)P(A) + P(B|\overline{A})P(\overline{A})}.^{*}$$

If successive fusion estimates are made over time, and using I_t to indicate an indicator used in the estimate at time t, we can express Bayes rule using I rather than B and having the left side refer to time t and the right side refer to the previous time. The result is the algorithm for "Bayesian updating":

---

very good explanations of the theory. More recently, Bayes' theorem or rule has been used in decision analysis. Raiffa (1968) is one such book.

[*] $P(A,B)$ is the intersection of A and B, i.e., the *joint* probability that both A and B occur. Solving the conditional probability equation for the joint probability from equation we get $P(A,B) = P(A)P(A|B) = P(B)P(B|A)$. We also note that

$$P(B) = P[(A,B) \cup (\overline{A},B)] = P(A,B) + P(\overline{A},B) = P(B|A)P(A) + P(B|\overline{A})P(\overline{A}).$$

Substituting in the denominator of the conditional equation, we get Bayes rule as expressed in text.

$$P_t(A \mid I_t) = \frac{P_t(I_t \mid A)P_{t-1}(A)}{P_t(I_t \mid A)P_{t-1}(A) + P_t(I_t \mid \overline{A})P_{t-1}(\overline{A})} \, .$$

In this formulation, A is the event "the individual or group is engaged in a (potentially) threatening activity A," and $\overline{A}$ is its logical negation "the individual or group is not engaged in that activity A." $P_{t-1}(A)$ is our prior assessment of the likelihood that the individual or group is engaged in the activity A. $P_t(A \mid I_t)$ is the improved estimate of the likelihood that our suspect is about to join the group given related indicators. It is also referred to as the *posterior probability*. A more general model would relax this binary assumption (engaged/not engaged) and admit levels of engagement. We discuss this more below.

In general, our interest in an individual or group stems from some information causing us to concentrate some monitoring effort. For example, we might know that a particular meeting is by invitation only, raising the likelihood that an attendee might be interested in joining the organization, or we might know that many people drop in to such meetings, out of curiosity, but never come back or join the organization. The "priors" would be very different in those cases; that is, the "base rates" would affect the priors strongly.

As an example, suppose that a radical group's meetings are monitored over a few years and it seems that about 95 percent of attendees join. We would then expect an individual to be 95 percent likely to join if we see him in attendance. However, suppose that a highly trusted agent tells us that in very recent times the composition of the group has changed: "Real" radicals are going elsewhere, and attendees are virtually all just curious or perhaps vicarious adventure seekers, but not likely joiners. Suppose that he puts a 98 percent confidence level on that because, after all, there is some possibility of a radicalizing individual just "hiding in the crowd." If we take the agent's estimates as the basis for our time-t prior probability evaluation, then

$$P_{t-1}(A) = 0.95 \text{ and } P_{t-1}(\overline{A}) = 0.05 \, .$$

We can combine this with our prior estimate to "update" that estimate via Bayes' updating formula. Recalling that the "prior" estimate of $P(A)$ was 0.95, we obtain

$$P_t(A \mid I_t) = \frac{(0.02)(0.95)}{(0.02)(0.95) + (0.98)(0.05)} = \frac{0.019}{0.019 + 0.049} = 0.28.$$

Thus, the updated estimate suggests that seeing someone attend the meetings in question is not very good evidence—perhaps it implies a roughly 1 in 4 joining rate—but the estimate is still more conservative than the newest one based on the agent report.

This method is only an approximate inference, but it is at least systematic and in the right direction. It combines information in a not-unreasonable manner. So long as we do not interpret the "probabilities" too literally, the method might be helpful in identifying when someone merits more than normal attention.

Although easy to implement in simple problems, we believe that Bayesian updating methods are unlikely to go very far in our problem area, except for the simplest of instances. It is too difficult to reasonably estimate all the conditional probabilities required of the approach, especially with multiple indicators and multiple values thereof. And, as mentioned, it is difficult to reflect ambiguity or the absence of information. Finally, despite its ubiquity, Bayesian updating depends on subtle assumptions about the relative credibility of information and even the order in which updating is conducted. Other methods may hold more promise.

Belief Function Methods

Next we address information fusion methods that attempt to avoid the problems of the Bayesian approach, particularly assigning belief where there is no support. These methods are all based on Glenn Shafer's *belief function* concept. Belief functions are considered a "less restrictive Bayes." In his book (Shafer, 1976), Shafer distinguishes between probability and what he calls belief. In the real world, the value of proposi-

tions is often not binary, as in Yes or No, but something like Definitely, Probably, Maybe, Unlikely, Definitely Not, and also "Indeterminate" (i.e., we have no information). As a result, we need a richer vocabulary and mathematics. Evidence supporting one of the propositions may not say much about the others. If a juror says that the defendant is not proven guilty beyond a reasonable doubt, we don't know whether the jury believes that there was a 50 percent chance of guilt or a 90 percent chance of guilt.[*]

This illustrates the fundamental shortcoming of trying to do evidential reasoning in terms of binary probabilities alone.[†] If we estimate the probability of an event occurring as P, then—if everything is expressed in probabilities—we are "forced" to assume that the probability of the event not occurring is $1 - P$. For example, suppose we receive a report from a trusted agent that our suspect has just attended a radical group meeting. His past reporting suggests that he is 70 percent accurate. From this we can justify a "belief" (not a probability) of 70 percent that the suspect is about to join the radical group, but only 0 percent "belief" that he will not. That is, we have no evidence to support the proposition that he will not join and therefore, unlike probabilities (where we would assess 30 percent to the likelihood that he will not join), the beliefs of 70 percent and 0 percent need not add to 100 percent. Together, then, these two constitute a belief function (adapted from Shafer and Pearl, 1990). The remainder of this section briefly discusses the fundamentals of belief functions. Later sections cover information fusion methods that are based on belief functions.

For a given threatening activity, A, we have two hypotheses: Our suspect is engaged in this activity or he is not, or $\mathbf{A}_i = \{A_i, \overline{A_i}\}$. For example, the "Developing Intent" phase has three associated activities. This generates the set $\mathbf{P}_1 = \{A_{1,1}, A_{1,2}, A_{1,3}\}$, where $\mathbf{P}_1$ is the "developing Intent" phase and the $A_{1,i}$ are the three activities associated with the

---

[*]   It can be argued that assigning a subjective probability to a hypothesis is indeed assigning belief.

[†]   A better way to say this is "in terms solely of probabilities of propositions being true." That is, one may use the apparatus of probabilities, but distinguishing between probabilities of necessity or provability, rather than probability of truth (Shafer and Pearl, 1990a)

phase.[*] For ease of exposition, we drop the phase subscript from here on out.

In belief theory, the set of hypotheses, $\mathbf{P} = \{A_1, A_2, A_3\}$, is referred to as the *frame of discernment*. It is the set of possible states of proposition values (e.g., all can be true to none are true). The logical propositions are subsets of $\mathbf{P}$. For example, suppose indicator reports lead us to suspect that an individual has joined a radical group or that he is at least committed to the cause espoused by the group. This, then, is the logical disjunction (an "or") of the two propositions: A_2, "psychological convergence," and A_3, "joining." The set theoretic representation of the disjunction is $\{A_2 A_3\}$, a subset of $\mathbf{P}$. Note that the disjunction $\{A_1, A_3\}$ is true if A_1 is true, A_3 is true, or both are true.[†]

The Frame of Discernment

In general, we exploit the correspondence between propositions and subsets so that the logical notions of conjunction, disjunction, implication, and negation map into set-theoretic notions of intersection, union, inclusion, and complementation, respectively.[‡]

This leads us to examine all of the possible subsets of the frame of discernment: $2^{\mathbf{P}} = \{\varphi, \mathbf{P}, \{A_1\}, \{A_2\}, \{A_3\}, \{A_1, A_2\}, \{A_1, A_3\}, \{A_2, A_3\}\}$. At the logical level, this set represents the propositions derived from the frame $\mathbf{P}$. The set $2^{\mathbf{P}}$ is referred to as the power set because its cardinality (i.e., the number of elements in its set) is $2^{|\mathbf{P}|}$, or $2^3 = 8$ in this case. Applying this, we get $\mathbf{A} = \{A, \overline{A}\}$ if we assume that the only possibilities are engaging in the activity or not. The power set is simply $2^{\mathbf{A}} = \{\varphi, \mathbf{A}, \{A\}, \{\overline{A}\}\}$. Each of these propositions can be supported at some level based on evidence obtained from indicator reports (sensors

---

[*] The three activities are $A_{1,1} =$ Motivational development reflecting inherent characteristics and experiences; $A_{1,2} =$ Psychological convergence; and $A_{1,3} =$ Recruitment or joining. See Figure 1.3.

[†] If A and B are propositions in the frame, then $\{A, B\} = \{A \cup B\}$, the union of the two propositions or, in logic, the disjunction of the two. We deal with the conjunction later.

[‡] To illustrate using two propositions or sets, A and B: The set notation $A \cup B$ is equivalent to the logical notion $A \vee B$. The former is the union of the two sets and the latter is the logical disjunction of the two. The first is read "A union B" and the latter is read "A or B."

and sources). We denote the propositions in the power sets as C, and next we formally define the level of support, $m(C)$, for the proposition C.

Definition: If **P** is a frame of discernment, then a function, $m = 2^{\mathbf{P}} \rightarrow [0,1]$, is called a *basic probability assignment* number when

- $m(\varphi) = 0$ and
- $\sum_{C \in \mathbf{P}} m(C) = 1.$

In this formulation, $m(C)$ represents the belief committed to C *only*. No information concerning the support levels for the subsets of C is available from this assignment. As in our earlier example, if C represents the proposition "an individual has joined a radical group or he is at least committed to the cause espoused by the group," then regardless of the support level for C, we draw no conclusions about whether he has joined the group or has undergone psychological convergence. Suppose the support level for C is 0.3. Then $m(C) = m(\{A_2, A_3\}) = 0.3$.

Note the difference between this formulation and the probability approach. In general, if $C = \{A_2, A_3\}$, then

$$P(C) = P(A_2) + P(A_3) - P(A_2 \cap A_3).$$

If A_2 and A_3 are mutually exclusive, we impose the identity $P(C) = P(A_2) + P(A_3) = 0.3$. This restriction does not apply with belief functions. Hence, unlike the Bayes formulation, it is possible to start with no support for any of the activities. Consequently, we need not impose the requirement that $m(C) + m(\overline{C}) = 1$.

Combining belief functions allows us to deal more directly with conflicting evidence through the use of *focal elements*. A *focal element* is a subset of **P** that has some (nonzero) support; that is, if $C \subset \mathbf{P}$, then C is a focal element if and only if $m(C) > 0$. In our example, the proposition $\{A_1, A_2\}$ might be considered logically inconsistent, in that A_1 implies that the individual or group is just beginning to develop the cognitive and emotional underpinnings that might lead to a terrorist attack, whereas A_2 focuses on indicators that an individual or group is committed to involvement in a cause, and, as such, this activity generally represents a more advanced developmental state. Consequently, we

would conclude that the proposition has no support and is therefore not a focal element.

Belief Functions

The discussion so far has focused on a method for assessing the belief we are willing to assign to all subsets of a frame of discernment. The levels of belief are derived from the evidence provided by a single observation of a single indicator. The next step is to examine a method for combining the evidence from multiple indicators, both similar and disparate or from repeated observations from a single indicator. As we have shown, in Bayesian analysis, we conditionally update the probabilities on the hypotheses based on the collected evidence. For belief functions, we apply Dempster's rule of combination (Dempster, 1967).

Dempster's rule allows us to focus on the focal elements of the belief functions developed from the evidence produced from two sources and compute an *orthogonal sum* of the two that results in a third combined belief function. The combined belief function can then be combined with yet another belief function, and so forth. Although the combining algorithm is rather simple, its mathematical development is complex, and so we omit it here and explain the process with an example.[*]

We start by defining a belief function using the three activities discussed earlier. This produces the set of hypotheses $\mathbf{P} = \{A_1, A_2, A_3\}$, where each of the set elements are activities associated with this phase. Then, by definition, something is a *belief function* over $\mathbf{P}$ if it satisfies the following conditions:

1. $Bel(\mathbf{P}) = 1$,
2. $Bel(\varphi) = 0$, and
3. Given that the cardinality of the power set of $\mathbf{P}$ is 8, then for the collection B_1, B_2, K B_8 subsets of $\mathbf{P}$,

$$Bel(B_1 \cup ... \cup B_8) \geq \sum_i Bel(B_i) - \sum_{i<j} Bel(B_i \cap B_j) + ...(-1)^9 Bel(B_1 \cap ... \cap B_8).$$

---

[*] An excellent discussion of the method and its mathematical development can be found in Chapter 3 of Shafer (1976).

The first two properties are consistent with the axiomatic definition of probability. The third is where the two depart. The third property is better explained by example. First, the eight subsets of **P** is the frame of discernment. Suppose we have that $B_1 = \{A_1, A_2\}$ and $B_2 = \{A_2, A_3\}$; then the total belief committed to the disjunction of the two satisfies the following inequality:

$$Bel(B_1 \cup B_2) \geq Bel(B_1) + Bel(B_2) - Bel(B_1 \cap B_2) =$$
$$Bel(\{A_1, A_2\}) + Bel(\{A_2, A_3\}) - Bel\{A_2\}.$$

To make this a bit more concrete, suppose the evidence from a single indicator source results in the support levels for the following propositions (hypotheses):

- "The individual is exhibiting motivational development reflecting inherent characteristics and experiences":
  $m(\{A_1\}) = 0.2$
- "The individual is experiencing psychological convergence":
  $m(\{A_2\}) = 0$
- "The individual is experiencing psychological convergence and he is exhibiting motivational development":
  $m(\{A_1, A_2\}) = 0.3$
- "The individual is being recruited by a radical group":
  $m(\{A_3\}) = 0.2$.

  All other propositions:
  $m(\{A_1, A_3\}) = m(\{A_2, A_3\}) = 0$ and $m(\mathbf{P}) = 0.3$.

This example illustrates the methodology. For now, we assume that the evidence we have gathered comes from a single indicator source. We deal with evidence from multiple sources later in the discussion of Dempster's rule of combination. The evidence received allows us to make basic probability assignments to some of the subsets of the frame.

Note that it is possible to support disjunctions independently. That is, if the assignments were probabilities, then we would have that $m(\{A_1, A_2\}) = m(\{A_1\} + m(\{A_2\} = 0.2$ and not 0.3. The restriction that

the assignments sum to 1.0 forces us to assign 0.3 to the frame. Next, we observe that the belief functions arising from these assignments produce some interesting results:

- $Bel(\{A_1\}) = m(\{A_1\}) = 0.2$
- $Bel(\{A_2\}) = m(\{A_2\}) = 0$
- $Bel(\{A_3\}) = m(\{A_3\}) = 0.2$
- $Bel(\{A_1, A_2\}) = m(\{A_1, A_2\}) + m(\{A_1\}) + m(\{A_2\}) = 0.5$
- $Bel(\{A_1, A_3\}) = m(\{A_1, A_3\}) + m(\{A_1\}) + m(\{A_3\}) = 0.4$, and
- $Bel(\{A_2, A_3\}) = m(\{A_2, A_3\}) + m(\{A_2\}) + m(\{A_3\}) = 0.2$.

Note that the sum of the total beliefs of all subsets of the frame is considerably greater than 1. We can also apply the third condition for belief functions to this example. We have for example that $Bel(\{A_1, A_2\} \cup \{A_1, A_3\} \cup \{A_2, A_3\}) = Bel(\mathbf{P}) = 1$, and that $Bel(\{A_1\} \cup \{A_1, A_2\}) = Bel(\{A_1, A_2\}) = 0.5$. Evaluating the right side of the inequality, we get:

$$Bel(\{A_1\} \cup \{A_1, A_2\}) = Bel(\{A_1\}) + Bel(\{A_1, A_2\}) - Bel(\{A_1\} \cap \{A_1, A_2\})$$
$$= Bel(\{A_1\}) + Bel(\{A_1, A_2\}) - Bel(\{A_1\}) = 0.5,$$

and therefore the equality condition holds. A more interesting case is to evaluate the right side of the inequality for $Bel(\{h_1, h_2\} \cup \{h_1, h_3\} \cup \{h_2, h_3\}) = Bel(\mathbf{H}) = 1$. This gives us:

$$\begin{aligned} Bel(\{h_1, h_2\} \cup \{h_1, h_3\} \cup \{h_2, h_3\}) &= Bel(\{h_1, h_2\}) + Bel(\{h_1, h_3\}) \\ &+ Bel(\{h_2, h_3\}) - Bel(\{h_1, h_2\} \cap \{h_1, h_3\}) \\ &- Bel(\{h_1, h_2\} \cap \{h_2, h_3\}) \\ &- Bel(\{h_1, h_3\} \cap \{h_2, h_3\}) \\ &+ Bel(\{h_1, h_2\} \cap \{h_1, h_3\} \cap \{h_2, h_3\}) \\ &= Bel(\{h_1, h_2\} + Bel(\{h_1, h_3\}) \\ &+ Bel(\{h_2, h_3\}) - Bel(\{h_1\}) - Bel(\{h_2\}) \\ &- Bel(\{h_3\}) + Bel(\varphi) \\ &= 0.5 + 0.4 + 0.2 - 0.2 - 0 - 0.2 + 0 \\ &= 0.7. \end{aligned}$$

Therefore, in this case the inequality holds.

Dempster's Rule of Combination

For belief functions, we apply Dempster's rule of combination (Dempster, 1967).

   We illustrate the process with the example developed so far. We assume we have two belief functions defined on the same frame of discernment, $\mathbf{P} = \{A_1, A_2, A_3\}$. The evidence to support these beliefs would have come from fused indicator reports. We also assume that the two assessments are independent. Regardless, we denote the first Bel_1 and the second Bel_2 and depict the orthogonal sum as $Bel_1 \oplus Bel_2$.

   The basic probability support levels for the focal elements from the two sources are as follows. Note that we are taking Bel_1 to be the belief function already developed. Also note that only propositions with support are listed as focal elements, consistent with its definition.

- Bel_1: $m_1(\{A_1\}) = 0.2$, $m_1(\{A_3\}) = 0.2$, $m_1(\{A_1, A_2\}) = 0.3$, $m_1(\mathbf{P}) = 0.3$
- Bel_2: $m_2(\{A_3\}) = 0.2$, $m_2(\{A_1\ A_2\}) = 0.3$, $m_2(\mathbf{P}) = 0.5$.

   The combined support level for the general focal element A, $m_{1,2}(A)$ for any two focal elements is the normalized sum of the product of focal elements from both belief functions whose intersection is A. The normalizing divisor is the complement of the product of the support levels for all disjoint focal elements. Rather than discuss the mathematics of this definition, we resort to a simple algorithmic process. The matrix depicted in Table D.1 represents the products of all combinations of support levels between Bel_1 and Bel_2. Note that the matrix is not necessarily square. We also include the frame because of the requirement that the basic probability assignments sum to 1. Each entry in the table is the product of the support levels for the row and column entries. For example, $m_{1,2}(\{A_3\}) = m_1(\{A_3\})m_2(\{A_3\}) = 0.2 \times 0.2 = 0.04$. The entries in parentheses represent the product of two support levels for disjoint focal elements.

   Next we calculate the normalizing divisor by summing the "disjoint" entries and subtracting from 1. This is then used to divide each support level entry in Table D.1 to arrive at the normalized support level entries in Table D.2.

Table D.1
Support-Level Products

Bel_1 \ Bel_2	$m_1(\{A_1\})$	$m_1(\{A_3\})$	$m_1(\{A_1,A_2\})$	$m_1(\{\mathbf{P}\})$
$m_2(\{A_3\})$	(0.04)	0.04	(0.06)	0.06
$m_2(\{A_1,A_2\})$	0.06	(0.06)	0.09	0.09
$m_2(\{\mathbf{P}\})$	0.10	0.10	0.15	0.15

Table D.2
Normalized Support Levels

Bel_1 \ Bel_2	$m_1(\{A_1\})$	$m_1(\{A_3\})$	$m_1(\{A_1,A_2\})$	$m_1(\{\mathbf{P}\})$
$m_2(\{A_3\})$	0	0.0476	0	0.0714
$m_2(\{A_1,A_2\})$	0.0714	0	0.1071	0.1071
$m_2(\{\mathbf{P}\})$	0.1190	0.1190	0.1785	0.1785

$$N = 1 - (0.04 + 0.06 + 0.06) = 0.84.$$

Table D.2 represents the normalized entries with support for disjoint focal elements set to 0. This table is used to calculate the combined belief function, $Bel_1 \oplus Bel_2$.

The combined basic probability assignments for each focal element are calculated as the sum of the entries in Table D.2 for which the intersection of the row focal elements and column focal elements are the focal element being evaluated. For example, for the focal element $B = \{h_3\}$ we get the following (from Table D.2):

$$m_{1,2}(\{A_3\}) = 0.0476 + 0.0714 + 0.1190 = 0.2380.$$

Using this same method, the resultant combined support levels for each focal element then is

$$Bel_1 \oplus Bel_2: m_{1,2}(\{A_1\}) = 0.1904, \ m_{1,2}(\{A_3\}) = 0.2380,$$
$$m_{1,2}(\{A_1, A_2\}) = 0.3927, \ m_{1,2}(\mathbf{P}) = 0.1785.$$

It is interesting to compare the combined support levels to the constituent levels before combining. They are repeated here for convenience:

- Bel_1: $m_1(\{A_1\}) = 0.2$, $m_1(\{A_3\}) = 0.2$, $m_1(\{A_1, A_2\}) = 0.3$, $m_1(\mathbf{P}) = 0.3$
- Bel_2: $m_2(\{A_3\}) = 0.2$, $m_2(\{A_1 A_2\}) = 0.3$, $m_2(\mathbf{P}) = 0.5$.

For example the focal element $B = \{A_3\}$ has the following constituent support levels: $m_1(\{A_3\}) = m_2(\{A_3\}) = 0.02$. Compare this with the combined support level: $m_{1,2}(\{A_3\}) = 0.2380$. In other words, two separate (and independent) set of observed indicators indicate that there is a 20 percent likelihood that our suspect is about to join a radical group. However, the fusion of these two assessments results in an increase in the likelihood to almost 24 percent. In addition, note that the support level for the frame has decreased considerably (from 0.5 and 0.3 down to 0.1785). The combination of two independent indicator reports has reduced the level of uncertainty as more belief was assigned to focal elements. We observe the same phenomenon for the proposition $\{A_1, A_2\}$; that is, the proposition that our suspect individual is developing the cognitive and emotional underpinnings that could later support involvement in a hostile act or that such underpinnings have matured. In contrast, the support for $\{A_1\}$ decreased due to lack of direct Bel_2 support. In fact, the only reason we have support at all is because of the support derived from the support for the proposition $\{A_1, A_2\}$.

Ultimately, there are three shortcomings in using Dempster's rule of combination to fuse indicator reports and activities. The first is that there is simply no good way to deal with total conflict, that is, when two indicator reports support opposing propositions or when two activities cannot exist at the same time. One of many ways this can happen is when two sources of information have different subjective interpretations of the same thing. Even the existence of partial conflict means that one or more reports must be disregarded. Nevertheless, Dempster's rule of combination at least allows us to measure the degree of conflict that exists between reports. In Bayesian updating,

the erroneous probability is simply combined with the a priori probability using Bayes' rule.

The second is the condition that the frame of discernment consist of a finite set of *mutually exclusive and collectively exhaustive* set of hypotheses. Because the nature of the problem we are addressing is likely to result in imprecise, conflicting, and perhaps fuzzy reports, the "exclusion of the middle" principle (i.e., excluding the complement of the hypotheses) limits the range of possible states because support for the negation of a proposition is not allowed. That is a serious problem, because evidence is often conflicting and people reporting "the same thing" subjectively often disagree. As another example, it is possible that some indicators suggest that an individual is attempting to join a radical group, while others suggest that is *not* the individual's intention.

The third issue has to do with the fact that Dempster's rule of combination focuses on combining evidence from two possibly disparate sensors or sources—in our case, indicator reports or activities. The application of the rule involves a particular normalization that has the effect of ignoring conflict and attributing any belief associated with conflict to the null set, that is, in effect, has the effect of zeroing out evidence conflicting with the proposition having the highest belief score. As a result, the application of the rule may have the effect of producing inconsistent results (Sentz and Ferson, 2002). Sentz and Ferson attribute the identification of this problem to Lotfi Zadeh in his review of Shafer's book, *A Mathematical Theory of Evidence* (Zadeh, 1984). Zadeh provides a compelling example of erroneous results. Suppose that a patient is seen by two physicians regarding the patient's neurological symptoms. The first doctor believes that the patient has either meningitis, with a probability of 0.99, or a brain tumor, with a probability of 0.01. The second physician believes the patient actually suffers from a concussion, with a probability of 0.99, but admits the possibility of a brain tumor, with a probability of 0.01. Using the values to calculate the m (brain tumor) with Dempster's rule, we find that m(brain tumor) = 1. This rule of combination, then, yields a result that implies complete support for a diagnosis that both physicians considered to be very unlikely (discussed also in Sentz and Ferson [2002]).

This can be easily demonstrated using the algorithm described above. Table D.3 records support-level products for the three propositions.

The entries in parentheses are disjoint. That is, we rule out the possibility of the patient having both a brain tumor and a concussion (treating them like the mutually exclusive propositions that Dempster-Schafer deals with), likewise meningitis and a brain tumor and finally meningitis and a concussion. The normalizing term is then $1 - (0.0099 + 0.0099 + 0.9801) = 0.0001$. This results in a normalized table with a 1.0 in the upper left-hand cell and zeros elsewhere. Consequently, the fused support for brain tumor is 1.0—clearly not consistent with either physician's best guess.

Yet another shortcoming is more problematic but for a very practical reason. The combinatorial complexity of the rule of combination grows in proportion to $O(2^{2|\mathbf{P}|}k)$, where $|\mathbf{P}|$ is the cardinality of the frame (number of cases) and k is the number of activities to combine.[*] For small frames, i.e., when the number of activities is small, complexity is no problem. However, as the number increases, the problem magnifies considerably. For $|\mathbf{P}| = 4$, for example, the combinatorial com-

Table D.3
Support Levels for Neurological Symptoms

		Doctor 1	
		Brain Tumor	Meningitis
Doctor 2	Brain Tumor	.0001	(.0099)
	Concussion	(.0099)	(.9801)

---

[*]   This is an upper bound based on the assumption that all propositions in the frame are focal elements (i.e., they all have support). The complexity is greatly reduced if the number of focal elements is small.

plexity is order 256 for each combination. However, for 5, it increases to 1024 and for 10, it reaches 1,048,576! For large frames, other methods are required. One such method is described in Fixsen (1977) and Perry and Stephanou (1993).

The Dempster-Shafer theory as usually presented is not likely to be valuable in our problem, but spinoff methods may be—especially if care is taken in defining the propositions. One version is Dezert-Smarandache theory, presented next.[*]

The Dezert-Smarandache Theory of Plausible and Paradoxical Reasoning

The indicator reports that point to various threatening activities are likely to be imprecise, fuzzy, paradoxical, and highly conflicting. Combining the information from such reports can therefore be just as imprecise, and dealing with likely conflicting information is certainly problematic. We noted earlier that although resolving conflicting information may not be possible in all cases, Dempster's rule of combination at least allows us to measure the severity of the conflict. Several other methods of combination have been advanced in efforts to improve the ability to resolve conflicts arising from disconfirming and conflicting evidence. The Dezert-Smarandache theory (DSmT) is one of these (Smarandache and Dezert, 2009a, 2009b).

DSmT was developed specifically to deal with combining evidence from sources and sensors that produce, imprecise, fuzzy, paradoxical, and highly conflicting reports—precisely the type of reports we might expect from the indicators we have identified. It also purports to remove the three fundamental conditions imposed by the Dempster-Shafer theory discussed above (Smarandache and Dezert, 2009a).

---

[*]   We are aware that controversy exists, not so much about the theory, but about what has been referred to as self-promoting and mischievous behavior by Smarandache. Further, we are aware that the vast percentage of literature references to the work originates, directly or indirectly, with the authors. Nonetheless, we believe that there are important and practical aspects to the theory that need to be pursued.

The first is the assumption that the Shafer frame of discernment is too restrictive in that it excludes negation. If a proposition is a member of the power set $2^{\mathbf{P}}$, then its negation is not, another reminder that Dempster-Shafer theory comes from a tradition of propositions with binary values, yes or no. For example, it is not possible for both $\{A\}$ and $\{\overline{A_1}\}$ to be members of the power set. However, the Shafer frame excludes more than that. The frame of discernment is defined as the power set of the fundamental hypotheses. Each subset in the power set is taken to be a proposition. At the first fusion level, the hypothesis set for "developing intent" is $\mathbf{P} = \{A_1, A_2, A_3\}$. Thus, $B = \{A_1, A_3\}$ is a logical proposition. The power set includes all unions of the hypotheses and nothing else. Formally, the power set $2^{\mathbf{P}}$ is defined as the set of all composite propositions/subsets built from the elements of $\mathbf{P}$ with the operator (Smarandache and Dezert, 2009a, 2009b). Hence, more than just negation is excluded: Conjunction is excluded as well.

To overcome this deficiency, DSmT includes two more inclusive sets: the *hyper power set* and the *super power set*. Both of these sets allows for an expansion of the Shafer power set to include a richer representation of the propositions derived from the frame of discernment. This allows us to assess belief in more varied ways consistent with a more realistic representation of the likely fuzziness of estimates. For example, if $A_1 =$ "an individual attended a hostile group meeting," and $A_2 =$ "the same individual accessed the group's website," then an observer might report that he is 70 percent certain that he did both. Another may report that he is 50 percent certain that he did one or the other—but not both and a third may report that he is 80 percent certain that he *did not* attend a meeting, but has nothing to say about his accessing the group's website. This richer representation allows for a more realistic set of propositions that reflect real-world fuzziness inherent in proposition estimates. The two sets below allow for this richer representation:

- **The hyper-power set:** The hyper-power set, denoted $D^{\mathbf{P}}$, is defined to be the set of all composite propositions/subsets built from the element of the frame, $\mathbf{P}$, with both $\cup$ and $\cap$ operators. The cardinality of this set is such that $|D^{\mathbf{P}}| \geq |2^{\mathbf{P}}|$, thus admitting a greatly expanded set from the same frame of discernment. For example,

two of the indicators of the activity of "joining" are $I_1 =$ "Explore different organizations and strategies" and $I_2 =$ "Pure adherence to organizational standards." The Shafer model would allow for the following propositions: $\{\varphi, \{I_1 \cup I_2\}, \{I_1\}, \{I_2\}\}$. The hyper power set expands the Shafer model to $\{\varphi, \{I_1 \cup I_2\}, \{I_1 \cap I_2\}, \{I_1\}, \{I_2\}\}$. The set now allows us to assign belief to the proposition that one or the other reports and to the proposition that they both report. The actual cardinality relating to sets with increasing numbers of propositions follows the sequence of Dedekind's numbers: 1, 2, 5, 19, 167 . . . (Tombak, Isotamm, and Tamme, 2001). For example, if P is the degenerate case, then $\mathbf{P} = \{\varphi\}$ and $|D^{\mathbf{P}}| = 1$. Similarly, if $\mathbf{P} = \{A, B\}$, $D^{\mathbf{P}} = \{A, B, A \cap B, A \cup B, \varphi\}$ and then $|D^{\mathbf{P}}| = 5$. In our case, the cardinality for hyper-power set is 19. In addition to admitting all of the disjunctions and conjunctions of the frame, $\mathbf{P}$, it also admits combinations of both.

- **The super-power set:** The super-power set, denoted $S^{\mathbf{P}}$, adds complementation or logical negation to the hyper-power set. Hence, $S^{\mathbf{P}}$ consist of all composite propositions/subsets built from the elements of the frame, $\mathbf{P}$. In this case, we must have that $\{I_1, \cap I_2\} \neq \varphi$, and I_1 and I_2 are not disjoint (see Figure D.2). This admits an even greater number of elements. Expanding the set above to include the additional propositions gets us $\{\varphi, \hat{\varphi}, \{I_1 \cup I_2\}, \{I_1 \cap I_2\}, \{I_1\}, \{I_2\}, \{\overline{I_1}\}, \{\overline{I_2}\}, \{\overline{I_1 \cup I_2}\}, \{\overline{I_1 \cap I_2}\}\}$. However, since $\overline{\varphi} = \{I_1 \cup I_2\}$ and $\varphi = \{\overline{I_1 \cup I_2}\}$, we get the super power set $\{\varphi, \{I_1 \cup I_2\}, \{I_1 \cap I_2\}, \{I_1\}, \{I_2\}, \{\overline{I_1}\}, \{\overline{I_2}\}, \{\overline{I_1 \cap I_2}\}\}$. In the three activities example, Dezert and Smarandache show that the cardinality of $S^{\mathbf{P}}$ is $2^{2^{|\mathbf{P}|-1}}$. By comparison, the cardinality of the power set is 8, and the cardinality of the hyper-power set is 19.

Using these larger power sets, Dezert and Smarandache propose combining rules that overcome the deficiencies in Dempster's rule of combination discussed earlier. Then they compare their rules to several other proposed combining rules. Before describing the new rules, the belief function definitions first need to be generalized to operate on any of the power sets produced from the frame of discernment. Recall that

Figure D.1
Venn Diagram for the Super Power Set

P
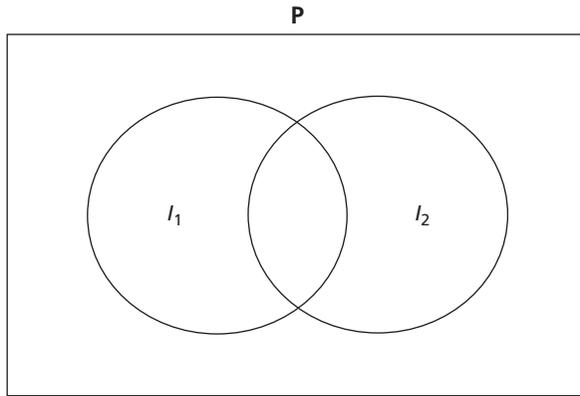
Shafer defined a *basic probability assignment* number, *m*, as $m(\varphi) = 0$ and

$$\sum_{C \subset P} m(C) = 1.$$

The generalized assignment takes into account the fact that any of the three power sets can be generated from the frame **P**. So we get instead $m(\varphi) = 0$ and

$$\sum_{C \in G^P} m(C) = 1,$$

where G^P can be any of the three power sets. Using this definition and the corresponding generalized belief function definition, two DmST combining rules are presented. As in Dempster's rule of combination, we have two independent indicator reports resulting in two belief functions Bel_1 and Bel_2 with associated belief assignments $m_1(\bullet)$ and $m_2(\bullet)$, except this time all four are generalized. Both combining

rules described below assume the hyper-power set is generated from the frame and this is meaningful to the fusion process:

- **The classic DSmT rule of combination:** This model is used when the free DSm model holds for the fusion process.[*] If there are no restrictions on the elements of the frame **P**, that is, for any pair of propositions, their intersection is not necessarily empty. So if $A_1 A_2 \in D^\mathbf{P}$, we have that $A_1 \cap A_2 \neq \varphi$ for some hypotheses in **P**. The combining rule then is

$$m_{1,2}(C) = \sum_{\substack{A,B \in D^\mathbf{P} \\ A \cap B = C}} m_1(A) m_2(B)$$

  for every $C \in D^\mathbf{P}$. Note that this is significantly different from Dempster's rule. Dezert and Smarandache use a simple frame of four elements to illustrate a case where Dempster's rule fails to produce a logical result whereas the classic DSmT rule does. More dramatic, however, is the demonstration that Zadeh's problem is resolved more logically using this method. The difference in this latter example is that the normalizing function is not used and the intersection of the elements is allowed so that the fused opinions are the numbers in Table D.3 above. Of course, one might question the logic of assuming that the patient has both a concussion and meningitis.
- **The hybrid DSmT rule of combination:** This second rule is used when the free DSmT model does not hold—i.e., when there is no guarantee that the intersections of the elements of **P** are not empty. This is close to the Shafer model where disjoint combinations of the elements of $2^\mathbf{P}$ are assumed to be empty. The general

---

[*] Many problems, such as assessing hostile intent, involve fuzzy continuous and relative concepts that have no absolute interpretation like "he accessed a hostile group's website." The *free DSm model* considers **P** as a frame of exhaustive elements which can potentially overlap: "he accessed a hostile group's website as part of a research project for school," or "he accessed a hostile group's website and ordered material on the group." Both of these activities overlap with the original hypothesis. This model is *free* because no other assumptions are made about the hypotheses.

form of the hybrid DSm combining rule is mathematically complex, and we omit it here. Essentially, it is referred to as hybrid because it is a combination of the classic DSm rule (for cases where the intersections are not empty) and the logical disjunction (union) of elements whose intersections are indeed empty. This method is also used on the two examples cited above, and both achieve the same results. In trying to discern hostile intent, however, we would expect that the intersection of the elements of **P** to be non-empty—consistent with the fuzzy nature of the activities being observed.

These combining rules along with the Dempster-Shafer model should be used where they fit the fusion model. If in the simple example of the doctors' opinions, we must rule out cases in which the patient has two causes for his or her illness, then the hybrid DSm model can be used. We rule out Dempster's rule of combination because it produces a nonsense result—even though excluding non-empty intersections makes sense.

Other Combining Methods

Several other combining methods may be of use in information fusion. In this section, we briefly introduce four of them: possibility theory, multi-attribute assessment, mutual information, and filtering. We do not provide much detail, but include them to illustrate the variety of possible methods and to suggest that each be examined more closely.

Possibility Theory

The phrase "theory of possibility" was coined by Lotvi Zadeh in a paper titled "Fuzzy Sets as a Basis for a Theory of Possibility" (Zadeh, 1978). Possibility theory is an uncertainty theory that deals with incomplete information, and is therefore well suited to the problem of discerning individual or group activity that may indicate hostile intent. Possibility theory states that any hypothesis not known to be impossible cannot be ruled out. In his paper on possibility theory, Zadeh states that ". . .

when our main concern is with the meaning of information—rather than with its measure—the proper framework for information analysis is possibilistic."[*] Our main concern in this work is the meaning of the indicator reports—not on how they are measured. However, we will discuss measurement as a prerequisite to combining.

A possibility distribution, denoted π_x, is taken to be a membership function of a fuzzy set of possible values of the quantity x (Dubois and Prade, 1994). All the values of x are assumed to be mutually exclusive, since it only takes on its true value. We further assume that all the values of x are contained in a closed and bounded set, $\mathbf{S}$. In the level 2 example, then, we might have the set $\mathbf{S} = \mathbf{A} = \{a_1, a_2, \text{L}, a_n\}$, where the elements of the set $\mathbf{A}$ represents an individual's level of engagement in activity A. This set is bounded and the elements are mutually exclusive. The actual value of x is unknown, but it must be true that one of the elements of the set A is the true value of x. If the true value of $x = a^*$, then we set $\pi_x(a^*) = 1$. Likewise, if $x \neq a$, i.e., if x cannot possibly be a, then $\pi_x(a) = 0$. In general, if $\pi_x(a_i) > \pi_x(a_j)$, then we say that a_i is considered more *plausible* than a_j. Furthermore, in complete ignorance (where every a_i is equally possible), we have that $\pi_x(a_i) = 1$ for all i. A possibility distribution, then, is a mapping of the set $\mathbf{S}$ to the unit interval [0,1]. As such, possibility complements probability theory.

As with all the combining methods discussed, we first start with some measure of how likely it is that an individual or group is engaged in some hostile activity based on some indicator report. With the possibility distribution defined, Dubois and Prade introduce two set-functions: the *possibility measure* and the *necessity measure* (Dubois and Prade, 1988). The two concepts are duals. Although these measures are interesting, they are generally not used to combine evidence. Consequently, we omit their mathematical development. The interested reader can find a development of the concept in Didier Dubois's paper *Possibility Theory and Statistical Reasoning* (Dubois, 2006).

In their paper on possibility theory and data fusion, Dubois and Prade focus on using possibility theory to fuse information when the

---

[*]   Zadeh credits the use of the term "possibilistic" to Gaines and Kohout in their paper on possible automata (Gaines and Kohout, 1975).

information available from sources and sensors may be of poor quality. They refer to this form of fusion as "pooling" or "aggregation" of information from disparate sources. We consider these terms to be synonymous with information fusion.

There is no unique combination mode using the possibilistic approach as there is in the previous methods. The method chosen will depend on what assumptions we make about the reliability of our sources of information. This is the first time the reliability of our sources is taken into account. In earlier chapters, we discuss indicators and the likely means of observing subjects looking for these indicators. In some cases, we rely on humans to provide indicator reports from direct or indirect observation and we also discuss technical means. In combining reports using possibility theory, the reliability of these sources will dictate the combining method to be used. There are three possible source deficiencies, and these may apply to human sources of technical sensors and sources (Dubois and Prade, 1988):

- **Inaccuracy:** The reports rendered are inconsistent with actual information about the individual or group. For example, the source always underestimates the true likelihood that the individual or group is engaged in a possibly hostile activity.
- **Overcautiousness:** The reports rendered by the source are too broad. The source always hedges by providing too great a range of possibilities. For example, after observing our suspect attend meetings with a radical group for three months, he reports that the likelihood that the individual may be joining the group is the same as the likelihood that he is just supporting the group with no intention of joining.
- **Overconfidence:** This is the dual of overcautiousness. In this case, the source reports that it is highly likely that the individual or group is engaged in a hostile activity based on very little evidence: for example, a report from an observer that an individual is definitely on the verge of joining a radical group based on his single access to the group's website.

Dubois and Prade develop an index that allows them to account for these deficiencies in report sources. In general, the index measures how accurate and informative the source is. We omit the development of this index. However, it is used in possibilistic combining.

There are basically two modes of combining reports from two or more disparate or similar sources: the conjunctive mode and the disjunctive mode. The former is used when all the sources agree somewhat and are reliable (based on the index discussed above). The latter is used when the sources disagree so that at least one of them is wrong.

For example, suppose two sources report on our suspect. One reports that it is likely that he will join the radical group but it is also possible that he intends only to support it financially. So he reports that the truth is in the set $A_1 = \{a_1, a_2\} \subseteq A = \{a_1, a_2, a_3, a_4\}$.[*] The second source reports that it is highly likely that the individual will not join the group, but there is some likelihood that he will support it financially. So the second source reports that truth is in the set $A_2 = \{a_2, a_4\} \subseteq A = \{a_1, a_2, a_3, a_4\}$. In this case, both reports agree somewhat. If they are also reliable, then the conjunction of the two sets represents the "pooled" consensus, giving us the fused assessment $A_3 = \{a_2\}$. Note that both sources considered financial support to be minimally likely, but because two reports included this possibility, it was considered to be the truth. This has the advantage of dealing directly with disconfirming and contradicting evidence. However, the next question is the likelihood estimate to be assigned. If source one assigned a likelihood of 40 percent and source two of 60 percent, what is the combine likelihood? Possibility theory is silent on this subject.

Multi-Attribute Assessment

The simplest (but perhaps not the most accurate) way to deal with the problem of fusing information from multiple and possibly disparate sources is to create a weighted sum of the activity likelihoods included

---

[*]   The four elements of the set A are the levels of participation in the radical group. They are, respectively, $a_1 =$ "the individual is on the verge of joining the radical group"; $a_2 =$ "the individual is about to support the group financially but will not join"; $a_3 =$ "the individual will write support the group by favorable writing on his blog"; and $a_4 =$ "the individual will not join the group."

in the indicator reports. Weights generally imply some notion of relative importance, and, in this case, the weights would be assigned to the reports—and ultimately to the sources. This, then, is another way to account for the reliability of the sources. However, as we will discuss below, it is better to consider the weights as reflecting the relative reliability of the reports. Although this is indeed desirable, it is not enough in all cases. What is needed is some way to represent the inherent dependencies among the reports and/or sources. Regardless of how well we are able to assign weights that truly reflect the relative reliability of the various reports and report sources, a weighted sum is inherently flawed because the likelihood estimates need not be additive. Nevertheless, as a means of comparison, the method sometimes is useful.[*]

The objective of multi-attribute assessment is to derive a single assessed likelihood that an individual or group is indeed engaged in some hostile activity. In this formulation, we assume that several indicator reports consisting of the likelihood that an individual or group is engaged in one or more of the hostile activities we have identified. The methods we discuss to develop this single assessment derive from Multiple Attribute Decision Making (MADM) theory (Hwang and Yoon, 1981).

Generally, MADM methods are used when a decision must be made between two or more alternatives based on multiple attributes that have incommensurable units, for example, speed and direction. We are suggesting its use here as a way to assess the likelihood that an individual or group is about to engage in each of the threatening activities we have identified. In this case, the "attribute" is the source, and the value of the attribute is the reported likelihood estimate based on the observed indicator. The choice of one technique over another depends on the nature of the sources whose likelihoods are being combined and their relation to one another. Here we discuss three methods: Simple Additive Weights (SAW), Weighted Product, and the Keeney-Raiffa multi-attribute utility method.

---

[*] This section is adapted from Perry and Moffatt (2004).

- **Simple Additive Weights (SAW) Method:** The SAW method is perhaps the simplest method of aggregation or fusion (Fishburn, 1967). A relatively old technique, it is cited in Article I, Section 2 of the U.S. Constitution as a method to determine the degree of a state's representation in the union (Yoon and Hwang, 1995). It is best used when the reports from the sources and sensors are independent of each other. For a case where there are n reports from several disparate sources or n reports from a single source concerning the activity, A_j, we get

$$L(A_j) = \sum_{i=1}^{n} \omega_i L_i(A_j),$$

where $L(A_j)$ is the combined likelihood that an observed individual or group is engaged in threatening activity A_j and $\mathbf{A} = \{A_1, A_2, L, A_m\}$ is the set of all activities identified in Figure 1.3 as threatening regardless of phase. The term ω_i, with $\sum_{i=1}^{n} \omega_i = 1$, is the weight (reliability) of the ith likelihood-estimate report $L_i(A_j)$. This method is problematic for two reasons: (1) The process must be repeated for each activity, and (2) assigning the reliability weights to each set of reports can be onerous. Nevertheless, if these two problems can be overcome, this method is attractive because of its simplicity. The result is a set of likelihoods, $L(\mathbf{A}) = \{L(A_1), L(A_2), L, L(A_m)\}$. Since each likelihood estimator is a fraction, threshold criteria can be set so that when the likelihood of any activity exceeds that value, we conclude that the individual or group is about to engage in that threatening activity.
- **Weighted Product Method:** The weighted product method is similar to the simple additive weights technique except in this case the likelihood values of the different reports are multiplied (Bridgman, 1922). The general form of this method is

$$L(A_j) = \prod_{i=1}^{n} [L_i(A_j)]^{\omega_i},$$

where $L(A_j)$, A_j, and ω_i are as above.

Although $L(A_j)$ might be used directly as a measure of the likelihood of the activity A_j based on n indicator reports, the A_j's fused likelihood in relation to a some threshold may also be used instead, so we obtain:

$$\overline{L}(A_j) = \frac{L(A_j)}{T_j},$$

where T_j is the threshold above which we conclude that the individual or group is engaged in activity A_j so we have that $T_j \geq L(A_j)$.

- **A Multiplicative Method from Keeney and Raiffa:** As discussed in Keeney and Raiffa (1976), on multi-attribute utility theory, it is sometimes important to use nonlinear utility functions. Adapting this method, the fusion algorithm takes the form

$$\Omega L(A_j) + 1 = \prod_{i=1}^{n} [\Omega \omega_i L_i(A_j) + 1],$$

where Ω is a normalizing factor used to insure consistency between the definition of $L(A_j)$ and the $L(A_j)$'s (de Neufville, 1990). The value of Ω is given by

$$\Omega = \prod_{i=1}^{n} [\Omega \omega_i + 1] - 1.$$

This technique is advantageous in that it allows for the consideration of possible interactions between the reports, which could be important.

As an example, if $n = 2$, and $\{L_2(A_j), L_1(A_j)\}$ is the set of indicator-likelihood reports for activity A_j, we get

$$L(A_j) = \omega_1 L_1(A_j) + \omega_2 L_2(A_j) + \Omega \omega_1 \omega_2 L_1(A_j) L_2(A_j),$$

with $\Omega = \dfrac{1 - \omega_1 - \omega_2}{\omega_1 \omega_2}.$

Mutual Information

This next section examines the relationship among the threatening activities. It is less concerned with fusing indicator reports than were the previous sections. The question here is what can we learn about other threatening activities given what we know about one or more particular threatening activities? For example, suppose that, based in several indicator reports, we conclude that it is highly likely that our suspect will join a radical group. Does that tell us anything about the likelihood that he will participate in target-identification activities? We refer to such questions in terms of asking about *mutual information*.

Mutual information is derived from information entropy (Cover and Thomas, 1991; Kullback, 1978; Shannon, 1948); it deals directly with independence among the activities. What we desire is a mathematical construct that will allow us to modify our estimates of the likelihood that an individual or group is about to engage in activity A_j through our knowledge that the individual or group is likely engaged in threatening activity A_k. We begin by treating $L(A_j) = L_j$ and $L(A_k) = L_k$ as continuous random variables defined on the interval [0,1]. We assume that they have distributions of values, even if empirical, which we denote for L_j as $f(L_j = l_j) = f(l_j)$. Similarly, we have that the distribution on L_k denoted as $g(L_k = l_k) = g(l_k)$. We assume that L_j and L_k are not independent. Because one random variable informs another, we refer to this construct as *mutual information*. Mutual information is based on the concept of *relative entropy*, which we take up next:

- **Relative Entropy:** Relative entropy measures the difference in entropies as calculated with two probability distributions denoted $D[f(l)\|q(l)]$. It is essentially the error incurred by assuming the true distribution for L is $f(l)$ when it is really $q(l)$. Relative entropy as defined by Cover and Thomas:

$$D[f(l)\,\|\,q(l)] = \int_0^1 f(l)\log\frac{f(l)}{q(l)}\,dl.$$

Note that if $f(l) = q(l)$, $D[f(l)\|q(lo)] = 0$. Relative entropy is not commutative. That is, $D[f(l)\|q(l)] = D[q(l)\|f(l)]$ is not always

true.[*] Kullback refers to the quantity $D[f(l)\|q(l)] + D[q(l)\|f(l)]$ as a measure of *divergence* between $f(l)$ and $q(l)$, and therefore as a measure of the difficulty of *discriminating* between them (Kullback, 1978).

- **Mutual Information:** We use the concept of relative entropy to measure of mutual information. First, we need to define $h(l_j, l_k)$, the joint probability function for our two random variables, L_j and L_k. We then define the mutual information to be the relative entropy between the joint probability distribution and the product of the probability distributions defined above:

$$I(L_j : L_k) = D[h(l_j, l_k) \| f(l_j) g(l_k)] =$$

$$\iint_{l_j l_k \in [01]} h(l_j, l_k) \log \frac{h(l_j, l_k)}{f(l_j) g(l_k)} \, dl_j \, dl_k.$$

Hence, $I(L_j{:}L_k)$ defined in this way is the amount of information about L_j gained from L_k.

The next issue is how to incorporate this knowledge into the estimate for L_j. We note first that $I(L_j{:}L_k) \in [0,1]$. We may employ some heuristic that increases L_j a fractional amount equivalent to $I(L_j{:}L_k)$; i.e., if $L_{j,0}$ is the likelihood that an individual or group is engaged in activity A_j prior to assessing mutual information, then we might calculate the contribution of L_k to be $L_k = L_{j,0}[1 + I(L_j{:}L_k)]$.

The difficulties associated with implementing this method are obvious. It requires that we know the probability distributions on the activity likelihood variables to start. But more problematic is the assumption that the joint probability be known, and since we must

---

[*]   A true metric satisfies the following properties: A *metric space* is a pair $(\mathbf{X}, d)$, where $\mathbf{X}$ is a set and d is a *metric* on $\mathbf{X}$ (or a distance function on $\mathbf{X}$), such that for all we have:

- d is real-valued, finite, and nonnegative.
- $d(x, y) = 0$ if and only if $x = y$.
- $d(x, y) = d(y, x)$.
- $d(x, y) \leq d(x, z) + d(z, y)$ (Kreyzig, 1978).

assume that the two random likelihood variables be dependent, the joint distribution is not merely the product of the separate distributions. Nevertheless, if these difficulties can be overcome, mutual information could be a useful way to assess the effect of what we know about the likelihood of one activity on what we know about another.

Filtering

Filtering is a process that removes noise from a signal. Applied to information fusion, the signal is the indicator report, and the noise is the inaccuracy associated with the uncertainties in the report and the errors introduced by the process itself. Filtering can be used when several sources and sensors issue indicator reports that need to be fused in near real time. Of the various filtering methods, the most commonly used is the Kalman filter, described below. The reader interested in implementation details may wish to consult *Optimal Estimation* by Frank Lewis (Lewis, 1986).

The combining process in a Kalman filter is essentially a sequential update of a state vector based on a prediction-correction process. In our application, the state vector is the vector of all likelihoods for the activities, i.e., $\mathbf{L}(t) = [L_1(t), L_2(t), \ldots L_m(t)]$ with time as one dimension. This is because of the time-sequential update mechanism associated with filtering. The prediction-correction mechanism assumes that in the absence of any indicator report, the state of the system is as of the last update. Hence we "predict" that this state persists into the future until we receive an update in the form of one or more indicator reports. These reports then "correct" the estimate, which then becomes the next prediction and remains so until subsequent reports are rendered.

Certain conditions must be satisfied to use Kalman filtering. The first is that the *dynamical system* be *linear*:

- **Dynamical System:** A system is dynamical if it changes over time (as new additional indicator reports arrive), with the time-dependence dictated by a fixed rule. The next state of the system is a function only of the state of the current system.

- **Linearity:** A function, $f(x)$, is said to be linear if it satisfies the following two properties: (1) $f(x + y) = f(x) + f(y)$ (additivity); and (2) $f(\alpha x) = \alpha f(x)$ for any α (homogeneity) (Grossman, 1980).

Secondly, we assume the reports that are received contain some error. If we are receiving human assessments of activity likelihood based on the indicators, then these assessments are certain to contain errors—including random errors. We refer to this as "noisy input data." The indicator reports are current estimates of the activity likelihoods and as such are treated as means of the probability distributions, $L_j(t)$. The error consists of two components: the distribution variance and random error. If the random error is "white noise," then the Kalman filter produces optimal estimates of the activity likelihoods.[*]

Once the next indicator report arrives, the estimates of the activity likelihoods, $\mathbf{L}(t)$, are updated using a weighted average, with more weight being given to reports with higher certainty. Because the Kalman filter's algorithm is recursive nature, it can run in real time using only the present indicator report(s) and the previously calculated state; no additional past information is required.

An example may help. Suppose we have just two threatening activities and want to assess the likelihood that an individual is engaged in either or both of these activities. The first of these is whether our suspect is indeed about to join a radical group and the second is whether that radical group is planning an attack. So, at time t, we have the state vector:

$$\mathbf{L}(t) = \begin{bmatrix} L_1(t) \\ L_2(t) \end{bmatrix}.$$

Let us further assume that indicator reports are arriving periodically and with some regularity. We are interested in the likelihood

---

[*] White noise is a term used for a random process with zero mean and finite standard deviation. If the noise is normally distributed with zero mean and finite standard deviation, it is referred to as Gaussian noise (Lewis, 1986).

estimates after a new report has arrived and has been processed, i.e., we are interested in $\mathbf{L}(t + 1)$. The Kalman filter representation of this system consists of several components. We discuss these here, not to completely define the process, but rather to illustrate the information needed to implement a Kalman filter fusion process for this problem. The first component is the *discrete time system* equations:

- **State transition equation:** The state transition equation describes how we expect $\mathbf{L}(t)$ to transition to $\mathbf{L}(t + 1)$ in the absence of a new indictor report. This equation models what is referred to as the "plant," i.e., where processing takes place. It consists of information on how the transition occurs, process noise (random inaccuracies in applying the transition process), and controls on inputs. All of these entities require subjective assessments.
- **Measurement system equation:** This equation measures the effect of the indicator report on the current estimate, $\mathbf{L}(t)$. This equation also requires subjective assessments. The first describes just how the observation affects the current state vector, and the second deals with report noise. This latter quantity can be interpreted as assessing the reliability of the report.
- **Predictor equation:** This equation produces an update at time $t + 1$ based on the information available at that time. It uses the same information included in the state transition equation.
- **Corrector equation:** This equation is used to update the state vector based on one or more indicator reports. It uses information from the measurement system equation but requires an additional assessment of the gain achieved by the Kalman filter. This is in the form of a matrix that is calculated from known quantities.

As we mentioned at the outset, the Kalman filter process is a good way to fuse indicator reports from disparate sources when the reports are periodic and are somewhat regular. The problem is that it requires the subjective assessment of several coefficients, many of which are matrices.

Summing Up

Information fusion is a critical component in detecting threatening behavior on the part of individuals or groups. Indicator reports are likely to originate in a wide variety of sources and sensors—some human and some technical, as discussed in earlier chapters. The common denominator among them all is an assessment of the likelihood that the individual or group observed is engaging in some threatening activity. This allows us to fuse a report from a remote heartbeat sensor with a human observation of some kind.

That said, the method best suited to fuse such information is far from settled. We have suggested several in this appendix, but each has problems. However, since, to our knowledge, none of these has been implemented for the context of our study, it is difficult to gauge accurately whether these problems can be overcome. What is needed is further research aimed at implementing one or more of the techniques using historical and experimental data.

Bibliography

Al Jazeera (2004, November 1), "Full Transcript of bin Laden's Speech." As of April 24, 2013:
http://www.aljazeera.com/archive/2004/11/200849163336457223.html

Alpers, Georg W., Andrew J. Winzelberg, and Catherine Classen (2005), "Evaluation of Computerized Text Analysis in an Internet Breast Cancer Support Group," *Computers in Human Behavior*, Vol. 21, pp. 361–376.

Aikins, Deane E., Daniel J. Martin, and Charles A. Morgan III (2010), "Decreased Respiratory Sinus Arrhythmia in Individuals with Deceptive Intent," *Psychophysiology*, Vol. 47, No. 4, pp. 633–636.

Aikins, Matthieu (2012, March), "The Siege of September 13," *GQ.*

Alaerts, Kaat, Evelien Nackaerts, Pieter Meyns, Stephan P. Swinnen, and Nicole Wenderoth (2011), "Action and Emotion Recognition from Point Light Displays: An Investigation of Gender Differences," *PLOS One*, Vol. 6, No. 6. As of April 20, 2013:
http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0020989

Albanese, David J., Michael J. Wiacek, Christopher M. Salter, and Jeffrey A. Six (2004), *The Case for Using Layered Defenses to Stop Worms*, Fort Meade, Md.: National Security Agency.

Alpers, Georg W., Andrew J. Winzelberg, and Catherine Classen (2005), "Evaluation of Computerized Text Analysis in an Internet Breast Cancer Support Group," *Computers in Human Behavior*, Vol. 21, pp. 361–376.

Arana, Ashley, Tessa Baker, and Sarah Canna (preparers) (2010), *Defining a Strategic Campaign for Working with Partners to Delegitimize Violent Extremism: Proceedings of a Workshop for the Strategic Multilayer Program of the Office of the Secretary of Defense*, Washington, D.C.: NSI Inc. As of April 22, 2013:
http://www.icst.psu.edu/docs/U_SMA_Delegitimizing%20Violent%20Extremism_Public%20Release_Final.pdf

Ascione, Fran (2001), "Animal Abuse and Youth Violence," *Juvenile Justice Bulletin*, September.

Ashworth, Scott, Joshua D. Clinton, Meirowitz, Ada, and Kristopher W. Ramsay (2008), "Design, Inference, and the Strategic Logic of Suicide Terrorism," *American Political Science Review*, Vol. 102, pp. 269–273.

Associated Press (2010, December 25), "Female Bomber Kills 43 at Pakistan Aid Center," CBS News.

——— (2012a, April 16), "Breivik Admits Massacre but Pleads Not Guilty," *CNS News*.

——— (2012b, July 22), "James Holmes Gun Club Membership Rejected Due to 'Bizarre' Behavior," *Huffington Post*. As of April 19, 2013:
http://www.huffingtonpost.com/2012/07/22/james-holmes-gun-club-membership_n_1693376.html

Astorino-Courtois, Allison, Hriar Cabayan, Bill Casebeer, Abigail Chapman, Diane DiEuliis, Charles Ehschlaeger, Dave Lyle, and Christopher Rice, eds. (2012), *National Security Challenges: Insights from Social, Neurobiological, and Complexity Sciences for the Strategic Multilayer Assessment Program, Office of the Secretary of Defense*, Washington, D.C.: NSI Inc. As of April 19, 2013:
http://nsiteam.com/pubs/National%20Security%20Challenges%20White%20Volume%20July%202012%20FINAL.PDF

Bakshy, Etyan (2012, January 17), "Rethinking Information Diversity in Networks," author's Facebook page. As of April 23, 2103:
http://www.facebook.com/notes/facebook-data-team/rethinking-information-diversity-in-networks/10150503499618859

BBC Forum (2001, June 11), "McVeigh Author Dan Herbeck Quizzed." As of April 20, 2013:
http://news.bbc.co.uk/1/hi/talking_point/forum/1378651.stm

BBC News Special Reports (2008, September 30), "London Attacks." As of April 20, 2013:
http://news.bbc.co.uk/2/hi/in_depth/uk/2005/london_explosions/default.stm

Berrebi, Claude (2009), "The Economics of Terrorism and Counterterrorism: What Matters and Is Rational-Choice Theory Helpful?" in Paul K. Davis and Kim Cragin, eds., *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corporation, pp. 151–208. As of April 19, 2013:
http://www.rand.org/pubs/monographs/MG849.html

Bjorgo, Tore, and John Horgan (2009), *Leaving Terrorism Behind: Disengagement from Political Violence*, New York: Routledge.

Blascovich, Jim, Wendy Berry Mendes, Sarah B. Hunter, and Brian Lickel (2000) "Stigma, Threat, and Social Interactions," in Todd F. Heatherton, Robert E. Kleck, Michelle R. Hebl, and Jay G. Hull, eds., *The Social Psychology of Stigma*, New York: Guilford Press, pp. 307–333.

Bloom, Mia (2005), *Dying to Kill: The Allure of Suicide Terror*, New York: Columbia University Press.

Blue Ribbon Panel (2011), *Report of the Mayor's Blue Ribbon Panel on Airport Security*, City of Los Angeles, Calif.

Bond, Charles F., and Bella M. DePaulo (2006), "Accuracy of Deception Judgments," *Personality and Social Psychology Review*, Vol. 10, No. 3, pp. 214–234.

Bond, Charles F., and Ahmet Uysal (2007), "On Lie Detection Wizards," *Law and Human Behavior*, Vol. 31, No. 1, pp. 109–115. As of April 20, 2013: http://www.communicationcache.com/uploads/1/0/8/8/10887248/on_lie_detection_wizards..pdf

Bornstein, Ann, Thyagaraju Damarla, John Lavery, Frank Morelli, and Elmar Schmeisser (2010), *Remote Detection of Covert Tactical Adversarial Intent of Individuals in Asymmetric Operations*, ARL-SR-197, Abderdeen Proving Ground, Md.: Army Research Laboratory. As of April 23, 2013: http://www.arl.army.mil/www/pages/185/ARL-SR-197_2010_04_15_final.pdf

Boukis, Christos, Aristodemos Pnevmatikakis, and Lazaros Polymenkakos, eds. (2007), *Artificial Intelligence and Innovations 2007: From Theory to Applications*: *Proceedings of the 4th IIIFIP International Federation for Information Processing, Vol. 247*, Boston: Springer.

Bridgman, Percy W. (1922), *Dimensional Analysis*, New Haven, Conn.: Yale University Press.

Brown, Angela K., and Michael Graczyk (2010), "Hasan Repeatedly Visited Firing Range Before Fort Hood Rampage," *Washington Post*, p. A1. As of April 23, 2103: http://www.washingtonpost.com/wp-dyn/content/article/2010/10/21/AR2010102106266.html

Brownley, Kimberly A., Barry Hurwitz, and Neil Schneiderman (2000), "Cardiovascular Psychophysiology," in John T. Cacioppo and Louis G. Tassinary, eds., *Handbook of Psychophysiology, 2nd ed.*, New York: Cambridge University Press, pp. 224–264.

Brulliard, Karin, and Shaiq Hussain (2011, May 12), "Pakistani Spy Chief Offers to Resign," *Washington Post*. As of April 23, 2103: http://www.washingtonpost.com/world/2011/05/12/AFdoRh1G_story.html?nav=emailpage

Bueno de Mesquita, Bruce (2011), "Applications of Game Theory in Support of Intelligence Analysis," in Baruch Fischoff and Cherie Chauvin, eds., *Intelligence Analysis: Behavioral and Social Scientific Foundations*, Washington, D.C.: National Academies Press, pp. 57–82.

Burbey, Ingrid, and Thomas L. Martin (2008), "Predicting Future Locations Using Prediction-by-partial-match." As of April 19, 2013: http://dl.acm.org/citation.cfm?id=1410014

Burgoon, Judee K., Douglas P. Twitchell, Matthew L. Jensen, Thomas O. Meservy, M. Adkins, J. Kruse, A. V. Deokar, G. Tschpenakis, Shan Lu, Dimitis N. Metaxas, Jay F. Nunamaker, and Robert E. Younger (2009), "Detecting Concealment of Intent in Transportation Screening: A Proof of Concept," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 10, No. 1, pp. 103–112.

Burns, Robert P. (2008), *Privacy Impact Assessment for the Future Attribute Screening Technology (FAST) Project*, Washington, D.C.: Department of Homeland Security. As of April 22, 2013:
http://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_st_fast.pdf

Butterworth, Bruce R., Shalom Dolev, and Brian Michael Jenkins (2012), *Security Awareness for Public Bus Transportation: Case Studies of Attacks Against the Israeli Public Bus System*, MTI Report 11-07, San Jose, Calif.: Mineta Transportation Institute.

Cacioppo, John T., David J. Klein, Gary G. Berntson, and Elaine Hatfield (1993), "The Psychophysiology of Emotion," in Michael Haviland and Jeannette M. Lewis, eds., *Handbook of Emotions*, New York: Guilford Press, pp. 119–142.

Cacioppo, John T., Bert N. Uchino, and Gary G. Berntson (1994), "Individual Differences in the Autonomic Origins of Heart Rate Reactivity: The Psychometrics of Respiratory Sinus Arrhythmia and Preejection Period," *Psychophysiology*, Vol. 31, No. 4, pp. 412–419.

Cacioppo, John T., Bert N. Uchino, Stephen L. Crites, Mary A. Snydersmith, Gregory Smith, Gary G. Berntson, and Peter Lang (1992), "Relationship Between Facial Expressiveness and Sympathetic Activation in Emotion: A Critical Review, with Emphasis on Modeling Underlying Mechanisms and Individual Differences," *Journal of Personality and Social Psychology*, Vol. 62, No. 1, pp. 110–128.

Cañal-Bruland, R., and M. Schmidt (2009), "Response Bias in Judging Deceptive Movements," *Acta Psychologica*, Vol. 130, No. 3.

Caridakis, George, Ginevra Castellano, Loic Kessous, Amaryllis Raouzaiou, Lori Malatesta, Stelios Asteriadis, and Kostas Karpouzis (2007), "Multimodal Emotion Recognition from Expressive Faces, Body Gestures and Speech," in Christos Boukis, Aristodemos Pnevmatikakis, and Lazaros Polymenakos, eds., *Artificial Intelligence and Innovations 2007: From Theory to Applications: Proceedings of the 4th IFIP International Federation for Information Processing, Vol. 247*, Boston: Springer, pp. 375–388.

Castellano, Ginevra, Loic Kessous, and George Caridakis (2008), "Emotion Recognition Through Multiple Modalities: Face, Body Gesture, Speech Affect and Emotion in Human-Computer Interaction," in Christian Peter and Russell Beale, eds., *Artificial Intelligence and Innovations 2007: From Theory to Applications: Proceedings of the 4th IFIP International Federation for Information Processing, Vol. 247*, Berlin: Springer, pp. 92–103

Cavanagh, J. T. O., A. J. Carson, A. J. Sharpe, and S. M. Lawrie (2003), "Psychological Autopsy Studies of Suicide: A Systematic Review," *Psychological Medicine*, Vol. 33, No. 3, pp. 395–405.

CBS News (2010), "Man Arrested in D.C. Subway Bomb Plot." As of April 8, 2013:
http://www.cbsnews.com/2100-201_162-6996775.html

Chauvin, Cherie (for Board on Behavioral, Cognitive, and Sensory Sciences), ed. (2011), *Threatening Communications and Behavior: Perspectives on the Pursuit of Public Figures*, Washington, D.C.: National Academy Press.

Chen, Hsinchun, Edna Reid, Joshua Sinai, Andrew Silke, and Boaz Ganor, eds. (2009), *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security*, New York: Springer.

Chittaranjan, Gokul, Jan Blom, and Daniel Gatica-Perez (2012), "Mining Large-Scale Smartphone Data for Personality Studies," *Personal and Ubiquitous Computing*, Vol. 17, No. 3, pp. 433–450. As of April 20, 2013:
http://publications.idiap.ch/downloads/papers/2011/Chittaranjan_PUC_2012.pdf

Chouchourelou, Arieta, Toshikhiko Matsuka, Kent Harber, and Maggie Shiffrar (2006), "The Visual Analysis of Emotional Actions," *Social Neuroscience*, Vol. 1, No. 1, pp. 63–74. As of April 20, 2013:
http://nwkpsych.rutgers.edu/roar/reprint%20pdfs/ChouchourelouEtAl06.pdf

Chow, Brian G., Richard Silberglitt, Caroline Reilly, Scott Hiromoto, and Christina Panis (2012), *Toward Affordable Systems III: Portfolio Management for Army Engineering and Manufacturing Development Programs*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG1187.html

Chung, Cindy K., and James W. Pennebaker (2011), "Using Computerized Text Analysis to Assess Threatening Communications and Actual Behavior," in Claire Chauvin (for Board on Behavioral, Cognitive, and Sensory Sciences), ed., *Threatening Communications and Behavior: Perspectives on the Pursuit of Public Figures*, Washington, D.C.: National Academic Press, pp. 3–32.

Clarke, Ronald V., and Marcus Felson, eds. (1993), *Routine Activity and Rational Choice*, New Brunswick and London: Transaction Publishers.

Cohen, Charles J., Frank Morelli, and Katherine A. Scott (2008), "A Surveillance System for the Recognition of Intent Within Individuals and Crowds," in *Proceedings of 2008 IEEE Conference on Technologies for Homeland Security*. As of April 22, 2013:
http://www.openskies.net/papers/Intent%20Recognition.pdf

Cohen, Charles J., Katherine A. Scott, Marcus J. Huber, Steve C. Rowe, and Frank Morelli (2008), "Behavior Recognition Architecture for Surveillance Applications," in *Proceedings of Applied Imagery Pattern Recognition Workshop, 2008*, pp. 1–8.

Cohn, M. A., M. R. Mehl, and J. W. Pennebaker (2004), "Linguistic Markers of Psychological Change Surrounding September 11, 2001," *Psychological Science*, Vol. 15, pp. 687–693.

Cornish, Derek B., and Ronald V. Clarke, eds. (1986), *The Reasoning Criminal: Rational Choice Perspectives on Offending*, New York: Springer-Verlag.

Costa, Paul T., and Robert R. McCrae (1990), "Personality Disorders and the Five-Factor Model of Personality," *Journal of Personality Disorders*, Vol. 4, No. 4, pp. 362–371.

Cover, Thomas M., and Joy A. Thomas (1991), *Elements of Information Theory*, New York: Wiley.

Crenshaw, Martha (2007), "Explaining Suicide Terrorism: A Review Essay," *Security Studies*, Vol. 16, No. 1, pp. 133–162.

Crossett, Chuck, and Jason A. Spitaletta (2010), *Radicalization: Relevant Psychological and Sociological Concepts*, Laurel, Md.: Johns Hopkins University Applied Physics Laboratory.

Cullen, Dave (2009), *Columbine*, New York: Hatchette Bood Group.

Dael, N., M. Mortillaro, and K. R. Scherer (2012), "Emotion Expression in Body Action and Posture," *Emotion*, Vol. 12, No. 5, pp. 1085–1101.

Damphousse, Kelly R. (2008), "Voice Stress Analysis: Only 15 Percent of Lies About Drug Use Detected in Field Test," *National Institute of Justice Journal (NIJ Journal)*, Vol. 259. As of April 22, 2013:
http://www.nij.gov/journals/259/voice-stress-analysis.htm

Damphousse, Kelly R., Laura Pointon, Deidra Upchurch, and Rebecca K. Moore (2007, June), "Assessing the Validity of Voice Stress Analysis Tools in a Jail Setting," *NIJ Journal*, Vol. 259. As of April 8, 2013:
http://www.ncjrs.gov/pdffiles1/nij/grants/219031.pdf

Daraghmeh, Ali (2004), "Boy Bomber Angers Some Palestinians," *Chicago Sun-Times.*

Darilek, Richard E., Walter L. Perry, Jerome Bracken, Brian Nichiporuk, and John Gordon IV (2001), *Measures of Effectiveness for the Information-Age Army*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monograph_reports/MR1155.html

Davis, Paul K. (2002), *Analytic Architecture for Capabilities-Based Planning, Mission-System Analysis, and Transformation*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monograph_reports/MR1513.html

———, ed. (2011), *Dilemmas of Intervention: Social Science for Stabilization and Reconstruction*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG1119.html

Davis, Paul K., and Kim Cragin, eds. (2009), *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG849.html

Davis, Paul K., and Paul Dreyer (2009), *RAND's Portfolio Analysis Tool (PAT): Theory, Methods, and Reference Manual*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/technical_reports/TR756.html

Davis, Paul K., David C. Gompert, Stuart Johnson, and Duncan Long (2008a), *Developing Resource-Informed Strategic Assessments and Recommendations*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG703.html

Davis, Paul K., Eric Larson, Zachary Haldeman, Mustafa Oguz, and Yashodhara Rana (2012), *Understanding and Influencing Public Support for Insurgency and Terrorism*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG1122.html

Davis, Paul K., Russell D. Shaver, and Justin Beck (2008), *Portfolio-Analysis Methods for Assessing Capability Options*, Santa Monica, Calif.: RAND Corporation. As of April 22, 2013:
http://www.rand.org/pubs/monographs/MG662.html

Davis, Paul K., Russell D. Shaver, Gaga Gvineria, and Justin Beck (2008b), *Finding Candidate Options for Investment Analysis: From Building Blocks to Composite Options and Preliminary Screening*, Santa Monica, Calif.: RAND Corporation. As of April 19, 2013:
http://www.rand.org/pubs/technical_reports/TR501.html

Dawes, Robyn M. (1979), "The Robust Beauty of Improper Linear Models," *American Psychologist*, Vol. 34, pp. 571–582.

de Neufville, Richard (1990), *Applied Systems Analysis: Engineering Planning and Technology Management*, New York: McGraw-Hill.

Defense Advanced Research Projects Agency (2012), "Tag Team Threat-Recognition Technology Incorporates Mind, Machine." As of April 8, 2013:
http://www.darpa.mil/NewsEvents/Releases/2012/09/18.aspx

Defense Science Board (2012), *Predicting Violent Behavior*, Washington, D.C.: Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics.

Dempster, Arthur P. (1967), "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals of Mathematical Statistics*, Vol. No. 38, pp. 325–339.

Department of the Army (1988), *Suicide Prevention and Psychological Autopsy*, Pamphlet 600-24, Washington, D.C.: Department of the Army.

Department of Homeland Security (2011), *Validation of Behavior-Based Screening Techniques*, Washington, D.C.: Department of Homeland Security, Science and Technology Directorate. (Not publicly released, but cited in U.S. Government Accountability Office, *Aviation Security: Efforts to Validate TSA's Passenger Screening Behavior Detection Program Are Underway, But Opportunities Exist to Strengthen Validation and Address Operational Challenges*, Washington, D.C., 2010.)

DePaulo, Bella M., James J. Lindsay, Brian E. Malone, Laura Muhlenbruck, Kelly Charlton, and Harris Cooper (2003), "Cues to Deception," *Psychological Bulletin*, Vol. 129, No. 1, pp. 74–118.

Deresiewicz, William (2011), "Faux Friendship," in James S. Miller, ed., *Acting Out Culture*, Boston: St. Martin's Press, pp. 470–480.

Derrick, Douglas C., Aaron C. Elkins, Judee K. Burgoon, Jay F. Nunamaker, and Daniel D. Zeng (2010, May–June), "Border Security Credibility Assessments via Heterogeneou Sensor Fusion," *Intelligent Systems, IEEE*, pp. 41–49.

Dietz, Park Elliott, Daryl B. Matthews, Cindy Van Duyne, and Daniel A. Martell (1991a), "Threatening and Otherwise Inappropriate Letters to Hollywood Celebrities," *Journal of Forensic Sciences, JFSCA*, Vol. 36, No. 1, pp. 185–209.

——— (1991b), "Threatening and Otherwise Inappropriate Letters to Members of the United States Congress," *Journal of Forensic Sciences, JFSCA*, Vol. 36, No. 5, pp. 1445–1468.

Dittrich, Winand H., and Anthony P. Atkinson (2008) "The Perception of Bodily Expressions of Emotion and the Implications for Computing," in Jimmy Or, ed., *Affective Computing, Focus on Emotion Expression, Synthesis, and Recognition*, Vienna: InTech, pp. 157–184.

Dubois, Didier (2006), "Possibility Theory and Statistical Reasoning," *Elsevier Journal of Computational Statistics and Data Analysis*, Vol. 51, pp. 47–69.

Dubois, Didier, and Henri Prade (1988), *Possibility Theory—An Approach to Computerized Processing of Uncertainty*, New York: Plenum Press.

——— (1994), "Possibility Theory and Data Fusion in Poorly Informed Environments," *Control Engineering Practice*, Vol. 25, pp. 811–823.

Dutton, Donald (2003), "Theoretical Approaches to the Treatment of Intimate Violence Perpetrators," *Journal of Aggression, Maltreatment & Trauma*, Vol. 7, No. 1–2, pp. 7–23.

Egerton, Brooks (2009, November 28), "Immam's E-Mails to Fort Hood Suspect Hasan Tame Compared to Online Rhetoric," *Dallas Morning News*. As of April 20, 2013:
http://www.dallasnews.com/news/state/headlines/20091128-Imam-s-e-mails-to-Fort-3556.ece

Eisenberg, Daniel (2003, June 22), "The Triple Life of a Qaeda Man," *Time*. As of April 22, 2013:
http://www.time.com/time/magazine/article/0,9171,460158,00.html

Ekman, Paul (1970), "Universal Facial Expressions of Emotion," *California Mental Health Research Digest*, Vol. 8, No. 4, pp. 151–158.

——— (1981), "Mistakes When Deceiving," *Annals of the New York Academy of Sciences*, Vol. 364, No. 1, pp. 269–278.

——— (1992a), "Are There Basic Emotions?" *Psychological Review*, Vol. 99, No. 3, pp. 550–553.

——— (1992b), "Facial Expressions of Emotion: New Findings, New Questions," *Psychological Science*, Vol. 3, No. 1, pp. 34–38.

——— (1993), "Facial Expression and Emotion," *American Psychologist*, Vol. 48, No. 4, pp. 384–392.

——— (2003), "Darwin, Deception, and Facial Expression," *Annals of the New York Academy of Sciences*, Vol. 1000, No. 1, pp. 205–221.

——— (2011), "Testimony to Hearing on Behavioral Science and Security: Evaluating TSA's SPOT Program."

Ekman, Paul, and Erika L. Rosenberg (2005), "What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (Facs), Second Edition," Vol. 639.

Elias, Bart (2009), *Airport Passenger Screening: Background and Issues for Congress*, Washington, D.C.: Congressional Research Service. As of April 22, 2013:
http://www.hsdl.org/?view&did=37127

——— (2011), *Changes in Airport Passenger Screening Technologies and Procedures: Frequently Asked Questions*, Washington, D.C.: Congressional Research Service. As of April 22, 2013:
http://www.fas.org/sgp/crs/homesec/R41502.pdf

Elkins, Aaron C., Judee K. Burgoon, and Jay F. Nunamaker (April 2012), "Vocal Analysis Software for Security Screening: Validity and Deception Detection Potential," *Homeland Security Affairs*, Supplement 4, Article 1. As of April 20, 2013:
http://www.hsaj.org/?special:fullarticle=0.4.1

Elliott, Andrea (2010, January 27), "The Jihadist Next Door," *New York Times*, *The Times Magazine*. As of April 22, 2013:
http://www.nytimes.com/2010/01/31/magazine/31Jihadist-t.html?pagewanted=all&_moc.semityn.www

Elson, Sara Beth, Doug Yeung, Parisa Roshan, Sue Bohandy, and Alireza Nader (2012), *Using Social Media to Gauge Iranian Public Opinion and Mood After the 2009 Election*, Santa Monica, Calif.: RAND Corporation. As of April 22, 2013: http://www.rand.org/pubs/technical_reports/TR1161.html

Erickson, William H. (2001, May), *The Report of Governor Bill Owen's Columbine Review Commission*, Denver, Colo.: State of Colorado. As of April 22, 2013: http://www.state.co.us/columbine/Columbine_20Report_WEB.pdf

Eriksson, Anders, and Francisco Lacerda (2007), "Charlatantry in Forensic Speech Science: A Problem to Be Taken Seriously," *International Journal of Speech, Language, and the Law*, Vol. 2, pp. 169–193.

Esposito, Richard, Mary-Rose Abraham, and Rhonda Schwartz (2009, November 12), "Major Hasan: Soldier of Allah; Many Ties to Jihad Web Sites," *ABC News: The Blotter*. As of April 8, 2013: http://abcnews.go.com/Blotter/hasan-multiple-mail-accounts-officials/story?id=9065692

Fabozzi, Frank J., and Harry M. Markowitz (2002), *The Theory and Practice of Investment Management*, New York: Wiley.

Farrahi, Katayoun, and Daniel Gatica-Perez (2012), "Extracting Mobile Behavioral Patterns with the Distant N-Gram Topic Model," in *Proceedings of IEEE Int. Symposium on Wearable Computers (ISWC)*, p. unknown. As of April 8, 2013: http://publications.idiap.ch/downloads/papers/2012/Farrahi_ISWC_2012.pdf

Fawcett, Tom (2006), "An Introduction to ROC Analysis," *Pattern Recognition Letters*, Vol. 27, pp. 861–874.

Fein, Robert A., Paul Lehner, and Bryan Vossekuil (2006), *Educing Information—Interrogation: Science and Art: Foundations for the Future: Intelligence Science Board Phase 1 Report*, Washington, D.C.: National Defense Intelligence College.

Feller, William (1950), *An Introduction to Probability Theory and Its Applications*, New York: Wiley.

Fenstermacher, Laurie, Larry Kuznar, Tom Rieger, and Anne Speckhard, eds. (2010), *Protecting the Homeland from International and Domestic Terrorism Threats: Current Multi-Disciplinary Perspectives on Root Causes, the Role of Ideology, and Programs for Counter-Radicalization and Disengagement*, Washington, D.C.: Office of the Director, Defense Research and Engineering.

Ferrero, Mario (2006), "Martyrdom Contracts," *Journal of Conflict Resolution*, Vol. 50, No. 6, pp. 855–877. As of April 22, 2013: http://jcr.sagepub.com/content/50/6/855.abstract

Fessler, Daniel M. T., Colin Holbrook, and Jeffrey K. Snyder (2012), "Weapons Make the Man (Larger): Formidability Is Represented as Size and Strength in Humans," *PLOS ONE*, Vol. 7, No. 4. As of April 22, 2013: http://www.plosone.org/article/info:doi/10.1371/journal.pone.0032751

Fiedler, Klaus (1988), "The Dependence of the Conjunction Fallacy on Subtle Linguistic Factors," *Psychological Research*, Vol. 50, pp. 123–129.

Fischhoff, Baruch, and Cherie Chauvin, eds. (2011), *Intelligence Analysis: Behavioral and Social Scientific Foundations*, Washington, D.C.: National Academies Press.

Fishburn, Peter C. (1967), "Additive Utilities with Incomplete Product Set: Applications to Priorities and Assignments," Baltimore, Md.: ORSA.

Fixsen, Dale (1977), "The Modified Dempster-shafer Approach to Classification," *IEEE Transactions on Systems, Man and Cybernetics.*

Frank, Mark G., and Elena Svetieva (2012), "Lies Worth Catching Involve Both Emotion and Cognition," *Journal of Applied Research in Memory and Cognition*, Vol. 1, pp. 131–133.

Freeman, A. M. (2012), *Dismount Threat Recognition Through Automatic Pose Identification*. As of April 22, 2013: http://www.dtic.mil/dtic/tr/fulltext/u2/a558005.pdf

Gaines, B., and L. Kohout (1975), "Possible Automata," in *Proceedings of Proceedings of the International Symposium on Multi-Valued Logic*, pp. 183–196.

Gambetta, Diego, and Steffen Hertog (2009), "Why Are There So Many Engineers Among Islamic Radicals?" *European Journal of Sociology*, Vol. 50, No. 2, pp. 201–230. As of April 20, 2013: http://eprints.lse.ac.uk/29836/1/Why_are_there_so_many_Engineers_among_Islamic_radicals_(publisher).pdf

GAO—*See* Government Accountability Office.

Garfinkel, Simson L. (2012, August 13), "The iPhone Has Passed a Key Security Threshold," *Technology Review*. As of April 22, 2013: http://www.technologyreview.com/news/428477/the-iphone-has-passed-a-key-security-threshold/

Garvey, Paul, Richard A. Moynihan, and Les Servi (2012, December), "A Macro Method for Measuring Economic-Benefit Returns on Cybersecurity Investments," *Journal of Systems Engineering.*

Goldstein, Scott (2009), "Alleged Terrorist Smadi's Neighbors in Small Town of Italy Recall a Fun, Easygoing Friend," *Dallas News*. As of April 25, 2013: http://crimeblog.dallasnews.com/2009/09/alleged-terrorist-smadis-neigh.html/

Gortner, Eva-Maria, and James W. Pennebaker (2003), "The Archival Anatomy of a Disaster: Media Coverage and Community-Wide Health Effects of the Texas A&M Bonfire Tragedy," *Journal of Social and Clinical Psychology*, Vol. 22, No. 5, pp. 580–603.

Government Accountability Office (2010, May 20), *Aviation Security: Efforts to Validate TSA's Passenger Screening Behavior Detection Program Underway, but Opportunities Exist to Strengthen Validation and Address Operational Challenges*, Washington, D.C., GAO-10-763. As of April 23, 2013: http://www.gao.gov/products/GAO-10-763

——— (2012a), *Transportation Security Administration: Progress and Challenges Faced in Strengthening Three Key Security Programs, Including Testimony of Stephen M. Lord, Director: Homeland Security and Justice Issues*, GAO-12-541T, Washington, D.C. As of April 22, 2013: http://www.gao.gov/assets/590/589588.txt

——— (2012b), *Aviation Security: 9/11 Anniversary Observations on TSA's Progress and Challenges in Strengthening Aviation Security, Statement of Stephen M. Lord, Director Homeland Security and Justice Issues*, GAO-12-1024T, Washington, D.C. As of April 20, 2013: http://homeland.house.gov/sites/homeland.house.gov/files/Testimony%20Lord_2.pdf

Greenemeier, Larry (2011), "Something in the Way You Move: Cameras May Soon Recognize Criminals by Their Gait [Video]," As of April 8, 2013: http://www.scientificamerican.com/article.cfm?id=motion-capture-surveillance

Grubin, Don (2010), "The Polygraph and Forensic Psychiatry," *American Academy of Psychiatry and the Law*, Vol. 38, No. 4, pp. 446–451.

Grubin, Don, and Lars Madsen (2005), "Lie Detection and the Polygraph: A Historical Review," *Journal of Forensic Psychiatry & Psychology*, Vol. 16, No. 2, pp. 357–369.

Gruenewald, Tara L., Teresa E. Seeman, Carol D. Ryff, Arun S. Karlamangla, and Burton H. Singer (2006), "Combinations of Biomarkers Predictive of Later Life Mortality," *Proceedings of the National Academy of Sciences*, Vol. 103, No. 38, pp. 14158–14163.

Guardian (2012), "Terrorism Charges Against Five After London Arrests," July 19.

Gupta, Dipak K. (2008), *Understanding Terrorism and Political Violence: The Life Cycle of Birth, Growth, Transformation, and Demise*, New York: Routledge.

Guss, C. Dominik, Ma Teresa Tuason, and Vanessa B. Teixeira (2007), "A Cultural-Psychological Theory of Contemporary Islamic Martyrdom," *Journal for the Theory of Social Behaviour*, Vol. 37, No. 4, pp. 415–445.

Haddad, Darren, Sharon Walter, Roy Ratley, and Megan Smith (2002), *Investigation and Evaluation of Voice Stress Analysis Technology*, Rome, N.Y.: Air Force Research Laboratory

Hafez, Mohammed M. (2007), "Martyrdom Mythology in Iraq: How Jihadists Frame Suicide Terrorism in Videos and Biographies," *Terrorism and Political Violence*, Vol. 19, No. 1, pp. 95–115.

Hall, David L., and James Llinas (1997), "An Introduction to Multisensor Data Fusion," *Proceedings of the IEEE*, Vol. 85, No. 1.

Halovic, Shaun, and Christian Kroos (2009, April 6–9), "Facilitating the Perception of Anger and Fear in Male and Female Walkers," in *Proceedings of Symposium on Mental States*, held at Heriot-Watt University, Edinburgh, Scotland. As of April 8, 2013:
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.160.4581&rep=
rep1&type=pdf#page=4

Hamburg, M. (1983), *Statistical Analysis for Decision Making*, third edition, Harcourt Brace Jovanovich, 1983.

Harnsberger, J. D., H. Hollien, C. A. Martin, and K. A. Hollien (2009), "Stress and Deception in Speech: Evaluating Layered Voice Analysis," *Journal of Forensic Science*, Vol. 54, No. 3, pp. 642–650. As of April 22, 2013:

Harris, Sam, and Bruce Schneier (2012, May 25), "To Profile or Not to Profile: A Debate," authors' respective blogs. As of July 25, 2012:
http://www.samharris.org/blog/item/to-profile-or-not-to-profile and
http://www.schneier.com/essay-397.html

Heger, Lindsay (2012), *Terrorist Target Choice*, Los Angeles, Calif.: CREATE, University of Southern California. As of April 22, 2013:
http://research.create.usc.edu/current_synopses/27

Helmus, Todd C. (2009), "Why and How Some People Become Terrorists," in Paul K. Davis and Kim Cragin, eds., *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corporation, pp. 71–112. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG849.html

Helmus, Todd C., Christopher Paul, and Russell W. Glenn (2007), *Enlisting Madison Avenue: The Marketing Approach to Earning Popular Support in Theaters of Operations*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG607.html

Hess, Pamela (2009, November 9), "Radical Imam Praises Alleged Fort Hood Shooter," Associated Press.

Hoffman, Bruce (2006), *Inside Terrorism, 2nd Ed*, New York: Columbia University Press.

Hollywood, John, Diane Snyder, Kenneth N. McKay, and John E. Boon, Jr. (2004), *Out of the Ordinary: Finding Hidden Threats by Analyzing Unusual Behavior*, Santa Monica, Calif.: RAND Corporation. As of April 25, 2013: http://www.rand.org/pubs/monographs/MG126.html

Honts, Charles R., David C. Raskin, and John C. Kircher (1994), "Mental and Physical Countermeasures Reduce the Accuracy of Polygraph Tests," *Journal of Applied Psychology*, Vol. 79, No. 2, pp. 252–259.

Honts, Charles R., and William Schweinle (2009), "Information Gain of Psychophysiological Detection of Deception in Forensic and Screening Settings," *Applied Psychophysiology and Feedback*, Vol. 34, pp. 161–172.

Hopkins, Clifford S., Roy J. Ratley, Daniel S. Benincasa, and John J. Grieco (2005), "Evaluation of Voice Stress Analysis Technology," in *Proceedings of 38th Hawaii International Conference on System Sciences*. As of April 22, 2013: http://www.secintel.com/media/pdf/FV-VSASTUDY.pdf

Horgan, John (2009), *Walking Away from Terrorism (Political Violence)*, New York: Routledge.

Hosenball, Mark, Michael Isikoff, and Evan Thomas (2010, January 1), "The Radicalization of Umar Farouk Abdulmutallab," *Newsweek,* pp. 37–41.

Hospedales, Timothy, Shaogang Gong, and Tao Xiang (2009), "A Markov Clustering Topic Model for Mining Behaviour in Video," in *Proceedings of Computer Vision, 2009 IEEE 12th International Conference*, pp. 1165–1172.

Hubbard, Douglas, and Dylan Evans (2010), "Problems with Scoring Methods and Ordinal Scales in Risk Assessment," *IBM Journal of Research and Development*, Vol. 54, No. 3, pp. 2:1–2:10.

Hudson, Rex (1999, September), *The Sociology and Psychology of Terrorism: Who Becomes a Terrorist and Why*, Washington, D.C.: Federal Research Division, Library of Congress. As of April 20, 2013: http://www.loc.gov/rr/frd/pdf-files/Soc_Psych_of_Terrorism.pdf

Hwang, Ching-Lai, and K. Paul Yoon (1981), *Multiple Attributes Decision Making Methods and Applications*, Berlin Heidelberg: Springer.

Intelligence and Security Committee (UK) (2006), *Report into the London Terrorist Attacks on 7 July 2005.* As of April 24, 2013: http://www.scribd.com/doc/80193339/Intelligence-and-Security-Committee-Report-into-the-London-Terrorist-Attacks-on-7-July-2005

Intelligence and Terrorism Information Center at Center for Special Studies (Israel) (2004), "Passover Eve Massacre at Park Hotel in Netanya." As of January 19, 2013: http://www.terrorism-info.org.il/Data/pdf/PDF1/JUNE_6_2_1700435765.pdf

Jack, Rachael E., Oliver G. B. Garrod, Hui Yu, Roberto Caldara, and Philippe G. Schyns (2012), "Facial Expressions of Emotion Are Not Culturally Universal," *Proceedings of the National Academy of Sciences.* As of April 22, 2013:
http://www.pnas.org/content/early/2012/04/10/1200155109.abstract

Jackson, Brian A., John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, and Horacio R. Trujillo (2005), *Aptitude for Destruction, Volume 1: Organizational Learning in Terrorist Groups and Its Implications for Combating Terrorism*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG331.html

Jackson, Brian A., Peter Chalk, Kim Cragin, Bruce Newsome, John V. Parachini, William Rosenau, Erin M. Simpson, Melanie W. Sisson, and Donald Temple (2007), *Breaching the Fortress Wall: Understanding Terrorist Efforts to Overcome Defensive Technologies*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG481.html

Jackson, Brian A., Edward W. Chan, and Tom LaTourrette (2011), "Assessing the Security Benefits of a Trusted Traveler Program in the Presence of Attempted Attacker Exploitation and Compromise," Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/working_papers/WR855.html

Jackson, Brian A., Tom LaTourette, Edward W. Chan, Russell Lundberg, Andrew R. Morral, and David R. Felinger, eds. (2011), *Efficient Aviation Security: Strengthening the Analytic Foundation for Making Air Transportation Security Decisions*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG1220.html

Jacobson, Sheldon H. (2012), "Watching Through the 'I's' of Aviation Security," *Journal of Transportation Security*, Vol. 5, pp. 35–38.

James, David V., Paul E. Muellen, Michele T. Pathé, J. R. Melroy. L. F. Presto, B. Darnley, and F. R. Farnham (2009), "Stalkers and Harassers of Royalty: The Role of Mental Illness and Motivation," *Physiological Medicine*, Vol. 39, No. 9, pp. 1479–1490.

Jenkins, Brian Michael (2011), *Stray Dogs and Virtual Armies: Radicalization and Recruitment to Jihadist Terrorism in the United States Since 9/11*, Santa Monica, Calif.: RAND Corporation. As of April 20, 2013:
http://www.rand.org/pubs/occasional_papers/OP343

Jensen, Bjørn Sand, Jakob Eg Larsen, Kristian Jensen, Larsen Jan, and Lars Kai Hansen (2010), "Estimating Human Predictability from Mobile Sensor Data," *IEEE International Workshop on Machine Learning for Signal Processing.*

Jensen, Matthew L., Thomas O. Meservy, Judee Burgoon, and Jay F. Nunamaker (2010), "Automatic, Multimodal Evaluation of Human Interaction," *Group Decision and Negotiation*, Vol. 19, No. 4, pp. 367–389.

Johnson, Kerri L., Lawrie S. McKay, and Frank E. Pollick (2011), "He Throws Like a Girl (But Only When He's Sad): Emotion Affects Sex-Decoding of Biological Motion Displays," *Cognition*, Vol. 119, No. 2, pp. 265–280.

Jones, Josh, and Mike Ahlers (2004, August 7), "Slowdown in 'Chatter' Worries Officials," CNN.com. As of April 8, 2013:
http://articles.cnn.com/2004-08-06/us/terror.
wrap_1_qaeda-key-al-al-hindi?_s=PM:US

Juvenile Justice Bulletin (2001, September), "Animal Abuse and Violent Offending." As of April 8, 2013:
http://www.ncjrs.gov/html/ojjdp/jjbul2001_9_2/page4.html

Kagan, Jerome (1997), "In the Beginning: The Contribution of Temperament to Personality Development," *Modern Psychoanalysis*, Vol. 22, No. 2, pp. 145–155.

Kahn, Jeffrey, Renée M. Tobin, Audra E. Massey, and Jennifer A. Anderson (2007), "Measuring Emotional Expression with the Linguistic Inquiry and Word Count," *American Journal of Psychology*, Vol. 120, No. 2.

Kallioniatis, Alexander, and Iain MacLeod (2010), "Formalization and Agility in Military Headquarters Planning," *International C2 Journal*, Vol. 4, No. 1.

Kanade, Takeo, Jeffrey F. Cohn, and Yingli Tian (2000), "Comprehensive Database for Facial Epression Analysis," in *Proceedings of Fourth IEEE International Conference on Automatic Face and Gesture Recognition*, pp. 45–53. As of April 21, 2013:
http://www.ri.cmu.edu/pub_files/pub2/kanade_takeo_2000_1/kanade_takeo_2000_1.pdf

Kang, Chaogui, Song Gao, Xing Lin, Yu Xiao, and Yuan Yihong (2010), "Analyzing and Geo-visualizing Individual Human Mobility Patterns Using Mobile Call Records," *Geoinformatics, 2010 18th International Conference*, pp. 1–7. As of April 20, 2013:
http://www.klmp.pku.edu.cn/Paper/UsrFile/1006.pdf

Karg, Michelle, Robert Jenke, Kolja Kühnlenz, and Martin Buss (2009), "A Two-Fold PCA-Approach for Inter-Individual Recognition of Emotions in Natural Walking," in *Proceedings of International Conference on Machine Learning and Data Mining (Poster Proceedings)*. As of April 8, 2013:
http://www.lsr.ei.tum.de/fileadmin/publications/karg_mldm2009.pdf

Karg, Michelle, K. Kühnlenz, and Martin Buss (2010), "Recognition of Affect Based on Gait Patterns," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 40, No. 4, pp. 1050–1061.

Kass, Jeff (2009), *Columbine: A True Crime Story, a Victim, the Killers and the Nation's Search for Answers*, Ghost Road Press.

Kassin, Saul M., Steven A. Drizin, Thomas Grisso, Gisli H. Guðjónsson, Richard A. Leo, and Allison D. Redlich (2010), "Police-Induced Confessions: Risk Factors and Recommendations," *Law and Human Behavior*, Vol. 34, No. 1, pp. 3–38.

Ke-Xue, Dai, Li Guo-Hui, and Can Ya-Li (2006), "A Probabilistic Model for Surveillance Video Mining," in *Proceedings of* the *2006 International Conference on Machine Learning and Cybernetics*, pp. 1144–1148.

Keeney, Ralph L. (1992), *Value-Focused Thinking: A Path to Creative Decisionmaking*, Cambridge, Mass.: Harvard University Press.

Keeney, R., and H. Raiffa (1976), *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, Inc.

Keila, P. S., and D. B. Skillicorn (2005), "Detecting Unusual Email Communication," in *Proceedings of the 2005 Conference of the Centre for Advanced Studies on Collaborative Research*, pp. 117–125.

Kimmage, Daniel, and Kathleen Ridolfo (2007), "Iraqi Insurgent Media: The War of Images and Ideas," *Radio Liberty*, p. 13.

King, Jennifer L. (2007), "Deception Detecton: Creating Realistic Facial Expressions Using the Facs Methodology in Training Simulation," in *Proceedings of BRIMS Conference (Behavior Representation in Modeling and Simulation)*, p. 1. As of April 20, 2013:
http://brimsconference.org/archives/2007/abstract/07-BRIMS-022.htm

Ko, T. (2008), "A Survey on Behavior Analysis in Video Surveillance for Homeland Security Applications," in *Proceedings of Applied Imagery Pattern Recognition Workshop, 2008. AIPR '08. 37th IEEE,* pp. 1–8.

Koehler-Derrick, Gabriel (2012, May 22), "The Abbottabad Documents: Bin Ladin's Cautious Strategy in Yemen," *CTC Sentinel*.

Kreuger, Alan B., and Jitka Maleckova (2003), "Education, Poverty and Terrorism: Is There a Causal Connection?" *Journal of Economic Perspectives*, Vol. 17, pp. 119–144.

Kugler, Richard L. (2006), *Policy Analysis in National Security Affairs: New Methods for a New Era*, Fort McNair, D.C.: National Defense University.

Kull, Steven, Clay Ramsay, Stephen Weber, Evan Lewis, and Ebrahim Mohseni (2009, February 25), *Public Opinion in the Islamic World on Terrorism, Al Qaeda, and U.S. Policies*, College Park, Md.: WorldPublicOpinion.org, University of Maryland. As of April 8, 2013:
http://www.worldpublicopinion.org/pipa/pdf/feb09/STARTII_Feb09_rpt.pdf

Kullback, Solomon (1978), *Information Theory and Statistics*, Magnolia, Mass.: Peter Smith.

Kunde, Wilfried, Stefanie Skirde, and Matthias Weigelt (2011), "Trust My Face: Cognitive Factors of Head Fakes in Sports," *Journal of Experimental Psychology Applied*, pp. 110–127. As of April 22, 2013:
http://www.i3.psychologie.uni-wuerzburg.de/fileadmin/06020300/user_upload/Kunde/Kunde_Skirde_Weigelt_2011_JEP_Applied_.pdf

Kuznar, Larry, Allison Astorino-Courtois, and Sarah Canna, eds. (2011), *From the Mind to the Feet: Assessing the Perception-to-Intent-to-Action Dynamic*, Maxwell Air Force Base, Ala.: Air University Press.

La Fon, Dana (2008), "The Psychological Autopsy," in Brent E. Turvey, ed., *Criminal Profiling: an Introduction to Behavioral Evidence Analysis*, San Diego, Calif.: Elsevier, pp. 419–428.

Lalich, Janja A. (2004), *Bounded Choice: True Believers and Charismatic Cults*, University of California Press.

Lampe, Cliff, Nicole B. Ellison, and Charles Steinfield (2008), "Changes in Use and Perception of Facebook," *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, pp. 721–730.

Landree, Eric, Richard Silberglitt, Brian G. Chow, Michael S. Tseng, Eric Landree, and Lance Sherry (2009), *A Delicate Balance: Portfolio Analysis and Management for Intelligence Information Dissemination Programs*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG939.html

LaTourette, Tom (2012), "The Benefits of Security Depend on How Different Security Measures Work Together," in Brian A. Jackson, Tom LaTourette, Edward W. Chan, Russell Lundberg, Andrew R. Morral, and David R. Felinger, eds. (2011), *Efficient Aviation Security: Strengthening the Analytic Foundation for Making Air Transportation Security Decisions*, Santa Monica, Calif.: RAND Corporation, pp. 67–80. As of April 23, 2013:
http://www.rand.org/pubs/monographs/MG1220.html

Le, Quoc V., Marc'Aurelion Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg S. Corrado, Jeff Dean, and Andrew W. Ng (2012), "Building High-Level Features Using Large Scale Unsupervised Learning," in *Proceedings of 29th International Conference on Machine Learning*, Edinburgh, Scotland. As of April 8, 2013:
http://ai.stanford.edu/~quocle/faces_full.pdf

Leech, Rob (2011), "My Brother the Jihadi," BBC3. As of April 8, 2013:
http://www.bbc.co.uk/programmes/b010758h

Lewis, Frank L. (1986), *Optimal Estimation with an Introduction to Stochastic Control Theory*, Wiley.

Li, Shing-Han, David C. Yen, Wen-Hui Lu, and Chian Wang (2012), "Identifying the Signs of Fraudulent Accounts Using Data Mining Techniques," *Computers in Human Behavior*, Vol. 28, No. 3, pp. 1002–1013. As of April 22, 2013: http://www.sciencedirect.com/science/article/pii/S0747563212000040

Lincoln, Bruce (2006), *Holy Terrors, Second Edition: Thinking About Religion After September 11*, Chicago, Ill.: University of Chicago Press.

Llinas, James, Ann Bisantz, Colin Drury, Younho Seong, and Jiun-Jim Jian (1998), *Studies and Analyses of Aided Adversarial Decision-making: Phase 2: Research on Human Trust in Automation*, Buffalo, N.Y.: State University of New York at Buffalo.

London Regional Resilience Forum (2006, September), *Looking Back, Moving Forward: The Multi-Agency Debrief: Lessons Identified and Progress Since the Events of 7 July 2005*, London: Government Office for London. As of April 20, 2013: http://ia201119.eu.archive.org/tna/20061004085342/londonprepared.gov.uk/

Los Angeles Times Staff (2012, May 10), "Chinese Woman Facing Home's Demolition Blows Up Herself, Two Others," *Los Angeles Times*, p. 1. As of April 20, 2013: http://latimesblogs.latimes.com/world_now/2012/05/chinese-suicide-bomber-kills-herself-two-others.html

Lou, Michel, and Dan Herbeck (2001), *American Terrorist: Timothy McVeigh and the Oklahoma City Bombing*, New York: Regan Books.

Lyon, Elizabeth, and Becky Afergan (2012), "Social Media Conference Summary." As of April 23, 2013: http://www.ms.army.mil/news/HSCB%20Summer%202012%20Newsletter.pdf

Madan, Anmol, Ketayoun Farrahi, Daniel Gatica-Perez, and Alex Pentland (2011), "Pervasive Sensing to Model Political Opinions in Face-to-Face Networks," April 20, 2013, in Kent Lyons, Jeffrey Hightower, and Elaine M. Huang, eds., *Pervasive Computing*, Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 214–231. As of April 20, 2013: http://hdl.handle.net/1721.1/65862

Mairesse, François, Marilyn A. Walker, Matthias R. Mehl, and Roger K. Moore (2007), "Using Linguistoc Cues for the Automatic Recognition of Personality in Conversation and Text," *Journal of Artificial Intelligence Research*, Vol. 30, pp. 457–500. As of April 20, 2013: http://www.aaai.org/Papers/JAIR/Vol30/JAIR-3012.pdf

Marquise, R. (2008), "Terrorist Threat Indicators," *The Counter Terrorist*, Vol. 1, No. 3, pp. 25–32.

Martinez, Thomas, and John Guinther (1988), *Brotherhood of Murder: How One Man's Journey Through Fear Brought The Order—the Most Dangerous Racist Gang in America—to Justice*, New York: McGraw-Hill Companies.

Maschke, George W., and Gino J. Scalabrini (2005), *The Lie Behind the Lie Detector*, AntiPolygraph.org. As of April 22, 2013: http://69.90.109.228/lie-behind-the-lie-detector.pdf

Matsumoto, David (1990), "Cultural Similarities and Differences in Display Rules," *Motivation and Emotion*, Vol. 14, No. 3, pp. 195–214.

Mayew, William J., and Mohan Venkatachalam (2012), "The Power of Voice: Managerial Affective States and Future Firm Performance," *Journal of Finance*, Vol. 67, No. 1, pp. 1–43.

McCorkell, W. J. J., and R. M. E. Griffin (1998), "An Overview of the Scientific Examinations Performed After an Explosion on the Shankhill Road," *Science and Justice*, Vol. 38, No. 2, pp. 75–79.

McDuff, Daniel, Rana el Kaliouby, and Rosalind Picard (2011), "Crowdsourced Data Collection of Facial Responses," *Proceedings of the 13th international Conference on Multimodal Interfaces (ICMI)*, pp. 11–18. As of April 22, 2013: http://affect.media.mit.edu/pdfs/11.McDuff-etal-Crowdsourced-2011.pdf

McGrory, Daniel, and Zahid Hussain (2005, July 22), "Cousin Listened to Boasts About Suicide Mission," *The Times*.

McKinley, James C., and James Dao (2009, November 9), "After Years of Growing Tensions, 7 Minutes of Bloodshed," *New York Times*, pp. A1 and A16.

McLay, Laura A., Sheldon H. Jacobson, and John E. Kozba (2006), "Multilevel Passenger Screening Strategies for Aviation Security Systems," *Naval Research Logistics*, Vol. 53, No. 3, pp. 183–197.

Meaney, Michael J., Seema Bhatnagar, Sylvie Larocque, Cheryl M. McCormick, Nola Shanks, Shakti Sharma, James Smythe, Victor Viau, and Paul M. Plotsky (1996), "Early Environment and the Development of Individual Differences in the Hypothalamic-Pituitary-Adrenal Stress Response," in Cynthia R. Pfeffer, ed., *Severe Stress and Mental Disturbance in Children*, Washington, D.C.: American Psychiatric Press, Inc., pp. 85–127.

Meixner, John B., and J. Peter Rosenfeld (2011), "A Mock Terrorism Application of the P300-Based Concealed Information Test," *Pscyophysiology*, Vol. 48, No. 2, pp. 149–154.

Meloy, J. Reid (2011), "Approaching and Attacking Public Figures: A Contemporary Analysis of Communications and Behavior," in Cherie Chauvin (for Board on Behavioral, Cognitive, and Sensory Sciences), ed., *Threatening Communications and Behavior: Perspectives on the Pursuit of Public Figures*, Washington, D.C.: National Academies Press, pp. 75–102.

Memmott (NPR), Mark (2011, May 3), "Bin Laden's Death Prompts 'Fist-Pumping Joy' for One Ticket Agent," *NPR*. As of April 23, 2103: http://www.npr.org/blogs/thetwo-way/2011/05/03/135948879/bin-ladens-death-prompts-fist-pumping-joy-for-one-ticket-agent

Merari, Ariel (2010), "Driven to Death: Pyschological and Social Aspects of Suicide Terrorism."

Merritt, Rob, and Brooks Brown (2002), *No Easy Answers: The Truth Behind Death At Columbine*, New York: Lantern Books.

Meservy, Thomas O., Matthew L. Jensen, W. John Kruse, Judee K. Burgoon, and Jay F. Nunamaker, Jr. (2005a), "Deception Detection Through Automatic, Obtrusive Analysis of Nonverbal Behavior," *Intelligent Systems, IEEE*, Vol. 20, No. 5, pp. 36–43.

Meservy, Thomas O., Matthew L. Jensen, W. John Kruse, Judee K. Burgoon, Jay F. Nunamaker, Jr., Douglas P. Twitchell, Gabriel Tsechpenakis, and Dmitri N. Metaxas (2005b), "Automatic Extraction of Deceptive Behavioral Cues from Video," in *Proceedings of 2005 IEEE International Conference on Intelligence and Security Informatics*, pp. 198–208.

Middleton, Robert D. (2011), *Privacy Assessment Impact Update: Future Attribute Screening Technology (FAST)/Passive Methods for Precision Behavioral Screening*, DHS/S&T/PIA-012(a), Washington, D.C.: Department of Homeland Security. As of April 22, 2013:
http://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_st_fast-a.pdf

Moffitt, Terrie E. (2006), "Life-Course-Persistent Versus Adolescence-Limited Antisocial Behavior," in D. Cicchetti and D. J. Cohen, eds., *Developmental Psychopathology, Vol. 3: Risk, Disorder, and Adaptation (2nd ed.)*, Hoboken, N.J.: John Wiley & Sons Inc., pp. 570–598.

Moffitt, Terrie E., Avshalom Caspi, Honalee Harrington, and Barry J. Milne (2002), "Males on the Life-Course-Persistent and Adolescence-Limited Antisocial Pathways: Follow-Up at Age 26 Years," *Development & Psychopathology*, Vol. 14, No. 1, pp. 179–207.

Moniquet, Claude (2006), *"Omar Nasiri Book": An Anti-French Manipulation*, ESISC Background Analysis, ESISC (European Strategic Intelligence and Security Center). As of April 15, 2013:
http://archive.wikiwix.com/cache/?url=http://www.esisc.eu/documents/nasiri.pdf&title="Omar%20Nasiri"%20book%3A%20An%20anti-French%20manipulation

Montoliu, Raul, Jan Blom, and Daniel Gatica-Perez (2012), "Discovering Places of Interest in Everyday Life from Smartphone Data," *Multimedia Tools and Applications*.

Mood, Alexander, and Franklin Graybill (1963), *Introduction to the Theory of Statistics*, New York: McGraw-Hill.

Moreau, Ron, and Sudip Mazumdar (2008), "The Pakistan Connection." As of April 8, 2013:
http://www.thedailybeast.com/newsweek/2008/11/26/the-pakistan-connection.html

Morosan, Cristian (2012), "Understanding the Antecedents of Perceived Value of Registered Traveler Biometric Systems," *Journal of Hospital Management*, Vol. 21, No. 8, pp. 872–896.

Mosher, Daniel (2010), *Linguistic Deception Theory: Is the World of E-Discovery Ready?* Deloitte Financial Advisory Services LLP.

Moskowitz, Andrew (2004), "Dissociation and Violence: A Review of the Literature," *Trauma, Violence and Abuse*, Vol. 5, No. 1, pp. 21–46.

Moynihan, Colin (2011), "2nd Relative Says Man Planned to Destroy Evidence of Subway Terror Plot," *New York Times*, p. A25. As of April 23, 2103: http://www.nytimes.com/2011/07/20/nyregion/najibullah-zazis-father-destroyed-evidence-of-bomb-plot-relative-says.html

Moynihan, Richard A. (2005), *Investment Analysis Using the Portfolio Analysis Machine (PALMA) Tool*, McLean, Va.: MITRE Corporation. As of April 22, 2013: http://www.mitre.org/work/tech_papers/tech_papers_05/05_0848/05_0848.pdf

Mullaney, Helene Matheson, and Nancy Costigan (2010), *Pre-incident Indicators of Pbied: A Literature Review, a Draft,* Washington, D.C.: U.S. Department of Homeland Security.

Mullen, Paul E., et al. (2009), "The Fixated and the Pursuit of Public Figures," *Journal of Forensic Psychiatry and Practice*, Vol. 20, No. 1, pp. 33–47.

Nasiri, Omar (pseudonym) (2006), *Inside the Jihad: My Life with al Qaeda*, New York: Basic Books.

National Cholesterol Education Program, "Risk Assessment Tool for Estimating Your 10-Year Risk of Having a Heart Attack." As of April 8, 2013: http://hp2010.nhlbihin.net/atpiii/calculator.asp

National Commission on Terrorist Attacks (2004), *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*, New York: W. W. Norton & Company.

National Research Council (2003), *The Polygraph and Lie Detection*, Washington, D.C.: National Academies Press.

National Research Council (2008), *Pre-Milestone A and Early-Phase Systems Engineering*, Washington, D.C.: National Academies Press.

NBC News/Associated Press (2005), "Ticket Agent Recalls Anger in Atta's Eyes," March 7.

New York City (2012, August 8), "Mayor Bloomberg, Police Commissioner Kelly and Microsoft Unveil New, State-of-the-Art Law Enforcement Technology That Aggregates and Analyzes Existing Public Safety Data in Real Time to Provide a Comprehensive View of Potential Threats and Criminal Activity," press release on the Domain Awareness System, PR-291-12. As of April 22, 2013:
http://www.nyc.gov/portal/site/nycgov/menuitem.
c0935b9a57bb4ef3daf2f1c701c789a0/index.jsp?pageID=mayor_press_
release&catID=1194&doc_name=http%3A%2F%2Fwww.nyc.gov%2Fhtml%2Fo
m%2Fhtml%2F2012b%2Fpr291-12.html&cc=unused1978&rc=1194&ndi=1

Newman, Matthew L., Carla J. Groom, and Lori D. Handelman (2008), "Gender Differences in Language Use: An Analysis of 14,000 Text Samples," *Discourse Processes*, Vol. 45, No. 3.

Noricks, Darcy M. E. (2009), "The Root Causes of Terrorism," in Paul K. Davis, and Kim Cragin, eds., *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corporation, pp. 11–70. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG849.html

O'Hair, H. Dan, Daniel Rex Bernard, and Randy R. Roper (2011) "Communication-Based Research Related to Threats and Ensuing Behavior," in Cherie Chauvin (for Board on Behavioral, Cognitive, and Sensory Sciences), ed., *Threatening Communications and Behavior: Perspectives on the Pursuit of Public Figures*, Washington, D.C.: National Academy Press, pp. 33–74.

O'Sullivan, M. (2007), "Unicorns or Tiger Woods: Are Lie Detection Experts Myths or Rarities? A Response to 'On Lie Detection Wizards' by Bond and Uysal," *Law and Human Behavior*, Vol. 31, No. 1, pp. 117–123.

O'Sullivan, Maureen (2008), "Home Runs and Humbugs: Comment on Bond and DePaulo (2008)," *Psychological Bulletin*, Vol. 134, No. 4, pp. 493–497.

Office of the Surgeon General, National Center for Injury Prevention and Control, National Institute of Mental Health, and Center for Mental Health Services (2001), *Youth Violence: A Report of the Surgeon General*, Bethesda, Md.: National Center for Biotechnology Information. As of April 22, 2013:
http://www.ncbi.nlm.nih.gov/books/NBK44293/

Oppel, Richard, Mark Mazzetti, and Souad Mekhnnet (2010, January 4), "Attacker in Afghanistan Was a Double Agent," *New York Times*. As of April 22, 2013:
http://www.nytimes.com/2010/01/05/world/asia/05cia.html

Osborn, Andrew (2004, October 18), "Beslan Hostage-Takers 'Were on Drugs,'" *The Independent*.

Pantucci, Raffaello (2011), "A Typology of Lone Wolves: Preliminary Analysis of Lone Islamist Terrorists," *Developments in Radicalisation and Political Violence*.

Pape, Robert A. (2003), "The Strategic Logic of Suicide Terrorism," *American Political Science Review*, Vol. 97, No. 3, pp. 343–361.

Patrick, Christopher J. (2008), "Psychophysiological Correlates of Aggression and Violence: An Integrative Review," *Philosophical Transactions of the Royal Society B-Biological Sciences*, Vol. 363, No. 1503, pp. 2543–2555.

Patrick, Christopher J., Margaret M. Bradley, and Peter J. Lang (1993), "Emotion in the Criminal Psychopath: Startle Reflex Modulation," *Journal of Abnormal Psychology*, Vol. 102, No. 1, pp. 82–92.

Paul, Christopher (2009), "How Do Terrorists Generate and Maintain Support," in Paul K. Davis and Kim Cragin, eds., *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corporation, pp. 113–209. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG849.html

Pedahzur, A., A. Perliger, and L. Weinberg (2003), "Altruism and Fatalism: The Characteristics of Palestinian Suicide Terrorists," *Deviant Behavior*, Vol. 24, pp. 405–423.

Peng, Cui, Wang Fei, Sun Li-Feng, Zhang Jian-Wei, and Yang Shi-Qiang (2012), "A Matrix-Based Approach to Unsupervised Human Action Categorization,"*IEEE Transactions on Multimedia*, Vol. 14, No. 1, pp. 102–110.

Pennebaker, James W., Roger J. Booth, and M. E. Francis (2007), *Linguistic Inquiry and Word Count: Operator's Manual*, Austin, Tex.: LIWC.net.

Pennebaker, James W., and Cindy K. Chung (2008), "Computerized Text Analysis of Al-qaeda Transcripts," April 20, 2013, in K. Krippendorff and M. A. Bock, eds., *A Content Analysis Reader*, Thousand Oaks, Calif.: Sage, pp. 453–464. As of April 20, 2013:
http://homepage.psy.utexas.edu/homepage/faculty/pennebaker/reprints/Pennebaker&Chung_Al-Qaeda.pdf

Pennebaker, James W., Cindy K. Chung, M. Ireland, A. Gonzales, and Roger J. Booth (2007), *The Development and Psychometric Properties of LIWC2007*, Austin, Tex.: LIWC.net.

Pennebaker, James W., and T. J. Mayne (1997), "Linguistic Predictors of Adaptive Bereavement," *Journal of Personality and Social Psychology*, Vol. 72, No. 4.

Pennebaker, James W., and L. D. Stone (2003), "Words of Wisdom: Language Use Over the Lifespan," *Journal of Personality and Social Psychology*.

Perry, Walter L., Claude Berrebi, Ryan Andrew Brown, John Hollywood, Amber Jaycocks, Parisa Roshan, Thomas Sullivan, and Lisa Miyashiro (2013), *Predicting Suicide Attacks Integrating Spatial, Temporal, and Social Features of Terrorist Attack Targets*, Santa Moncia, Calif.: RAND Corporation. As of April 25, 2013:
http://www.rand.org/pubs/monographs/MG1246.html

Perry, Walt, and James Moffat (2004), *Information Sharing Among Military Headquarters: The Effects on Decisionmaking*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/monographs/MG226.html

Perry, Walter L., David A. Signori, and John E. Boon (2004), *Exploring Information Superiority: A Methodology for Measuring the Quality of Information and Its Impact on Shared Awareness*, Santa Monica, Calif.: RAND Corporation. As of April 20, 2013:
http://www.rand.org/pubs/monograph_reports/MR1467.html

Perry, Walter L., and Harry E. Stephanou (1993), "A Quantitative Treatment of Multilevel Specificity and Uncertainty in Variable Precision Reasoning," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 23, No. 2, pp. 445–451.

Perry, William J., and Charles M. Vest (Chairmen) (2008), *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*, Washington, D.C.: National Academies Press.

Polikovsky, Senya, Yoshinari Kameda, and Yuichi Ohta (2009), "Facial Micro-Expressions Recognition Using High Speed Camera and 3d-Gradient Descriptor," in *Proceedings of 3d international Conference on Crime Detection and Prevention (ICDP 2009)*, pp. 1–6.

Pool, Robert (2011), *Field Evaluation in the Intelligence and Counterintelligence Context: A Workshop Summary*, Washington, D.C.: National Academy Press.

Poole, Robert W. (2009), "The Case for Risk-Based Aviation Security Policy," *World Customs Journal*, Vol. 3, pp. 3–16.

Poole, Robert W., and George Passatino (2003), *A Risk-Based Airport Security Policy*, Policy Study 308, Los Angeles, Calif.: Reason Foundation.

Porges, Stephen W. (1995), "Orienting in a Defensive World: Mammalian Modifications of Our Evolutionary Heritage: A Polyvagal Theory," *Psychophysiology*, Vol. 32, pp. 301–318.

Porter, Stephen, and Leanne ten Brinke (2010), "The Truth About Lies: What Works in Detecting High-Stakes Deception?" *Legal and Criminological Psychology*, Vol. 15, No. 1, pp. 57–75.

Posamintier, Mette T., and Hervé Abdi (2003), "Processing Faces and Facial Expressions," *Neuropsychology Review*, Vol. 13, No. 3, pp. 113–143.

Post, Jerrold M. (2005), "When Hatred Is Bred in the Bone: Psycho-Cultural Foundations of Contemporary Terrorism," *Political Psychology*, Vol. 26, No. 4, pp. 615–636.

Pouilot, Louise, and Diego De Leo (2006), "Critical Issues in Psychological Autopsy Studies," *Suicide and Life-Threatening Behavior*, Vol. 36, No. 5.

Press, S. James (1989), *Bayesian Statistics: Principles, Models and Applications*, New York: Wiley.

Press, William H. (2009), "Strong Profiling is Not Mathematically Optimal for Discovering Rare Malfeasors," *Proceedings of the National Academy of Sciences*, Vol. 106, No. 6, pp. 1716–1719. As of April 23, 2013:
http://www.pnas.org/content/106/6/1716.abstract

Pugliese, Joseph (2008), "Biotypologies of Terrorism," *Cultural Studies Review*, Vol. 14, No. 2, pp. 49–66.

Putzel, Christof (2012), "Interview with Omar Hammami," Viewpoint. As of April 20, 2013:
http://current.com/shows/viewpoint/videos/exclusive-american-born-jihadist-tells-current-how-could-i-really-refrain-from-attacking-america/

Pyes, Craig (2000, January 21), "Canada Adds Details on Algerians' Suspected Bomb Plot," *New York Times*, pp. 1–2. As of April 22, 2013:
http://www.nytimes.com/2000/01/21/world/canada-adds-details-on-algerians-suspected-bomb-plot.html

Raiffa, Howard (1968), *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley.

Raine, Adrian (1996), "Autonomic Nervous System Activity and Violence," in David M. Stoff and Cairns Robert B., eds., *Aggression and Violence: Genetic, Neurobiological, and Biosocial Perspectives*, Mahwah, N.J.: Lawrence Erlbaum Associates, Inc., pp. 145–168.

Reeve, Simon (1999), *The New Jackals: Ramzi Yousef, Osama Bin Laden, and the Future of Terrorism*, Holliston, Mass.: Northeastern University Press. As of April 23, 2013:
http://books.google.co.uk/books?id=VQjpziNmoE4C&lpg=PP1&pg=PP1#v=onepage&q&f=false

Rieger, Thomas (2008), *Desperate Measures: Different Types of Violence, Motivations, and Impact on Stability*, Washington, D.C.: Gallup Consulting.

Robinson, William H., Jennifer E. Lake, and Lisa M. Seghetti (2005), *Border and Transportation Security: Possible New Directions and Policy Options*, Washington, D.C.: Congressional Research Service.

Rock, Margaret (2011, August 11), "NYPD to Scan Facebook, Twitter for Trouble," Forbes.com. As of April 23, 2013:
http://news.yahoo.com/nypd-scan-facebook-twitter-trouble-190941500.html

Roether, Claire L., Lars Omlor, Andra Christensen, and Martin A. Giese (2009), "Critical Features for the Perception of Emotion from Gait," *Journal of Vision: A Journal of Scientific Research and Biological Vision*, Vol. 9, No. 6. As of April 23, 2013:
http://www.journalofvision.org/content/9/6/15.full

Rosenfeld, Barry, and Steven D. Penrod (2011), *Resarch Methods in Forensic Psychology*, Hoboken, N.J.: John Wiley & Sons.

Ross, Brian, and ABC News Investigative Team (2011, September 6), "While America Slept: The True Story of 9/11," ABC News. As of April 20, 2013:
http://abcnews.go.com/Blotter/ten-years-ago-today-countdown-911/story?id=14191671#

Rotella, Sebastian (2012), "Militant Reaffirms Role of Pakistan in Mumbai Attacks," *Foreign Policy*, August 9.

Rowe, Neil C., Ahren A. Reed, Riqui Schwamm, Jeehee Cho, Jose J. Flores, and Arijit Das (2012) "Networks of Simple Sensors for Detecting Emplacement of Improvised Explosive Devices," *Critical Infrastructure Protection*, WIT Press, pp. 241–254. As of April 20, 2013:
http://faculty.nps.edu/ncrowe/critinfsense_rowe.htm

Rubin, Alissa J. (2011), "Assassination Deals Blow to Peace Process in Afghanistan," *New York Times*, p. A1.

Rubin, Alissa J., Ray Rivera, and Jay Healy (2011, September 14), "U.S. Embassy and NATO Headquarters Attacked in Kabul," *New York Times*, p. A1. As of April 20, 2013:
http://www.nytimes.com/2011/09/14/world/asia/14afghanistan.html

Rubin, Philip E. (2011), "Behavioral Science and Security Statement to Subcommittee on Investigations and Oversight Committee on Science, Space, and Technology, U.S. House of Representatives." As of April 8, 2013:
http://science.house.gov/sites/republicans.science.house.gov/files/documents/hearings/2011%2004%2001%20RubinTestimony.pdf

Rude, Stephanie S., Eva-Maria Gortner, and James W. Pennebaker (2004), "Language Use of Depressed and Depression-vulnerable College Students," *Cognition & Emotion*, Vol. 18, No. 8.

Ruderman, Wendy (2012, August 7), "Court Prompts Twitter to Give Data to Police in Threat Case," *New York Times*. As of April 23, 2013:
http://www.nytimes.com/2012/08/08/nyregion/after-court-order-twitter-sends-data-on-user-issuing-threats.html?_r=0

Runeson, Sverker, and Gunilla Frykholm (1983), "Kinematic Specification of Dynamics as an Informational Basis for Person-and-Action Perception: Expectation, Gender Recognition, and Deceptive Intention," *Journal of Experimental Psychology*, Vol. 112, No. 4, pp. 585–615. As of April 22, 2013:
http://www.skidmore.edu/~flip/Site/Lab/Entries/2007/9/10_Biological_Motion_files/Runeson1983.pdf

Russell, James A. (1995), "Facial Expressions of Emotion: What Lies Beyond Minimal Universality?" *Psychological Bulletin*, Vol. 118, No. 3, pp. 379–391.

Russell, Tamara A., Elvina Chu, and Mary L. Phillips (2006), "A Pilot Study to Investigate the Effectiveness of Emotion Recognition Remediation in Schizophrenia Using the Micro-expression Training Tool," *British Journal of Clinical Psychology*, Vol. 45, No. 4, pp. 579–583.

Russell, Tamara A., Melissa J. Green, Ian Simpson, and Max Coltheart (2008), "Remediation of Facial Emotion Perception in Schizophrenia: Concomitant Changes in Visual Attention," *Schizophrenia Research*, Vol. 103, Nos. 1–3, pp. 248–256.

Saaty, Thomas L. (1999), *Decision Making for Leaders: The Analytic Hierarchy Process for Decisions in a Complex World, New Edition 2001 (Analytic Hierarchy Process Series, Vol. 2)*, RWS Publications.

Sageman, Marc (2004), *Understanding Terror Networks*, Philadelphia: University of Pennsylvania Press.

——— (2008), *Leaderless Jihad: Terror Networks in the Twenty-First Century*, Philadelphia: University of Pennsylvania Press.

Savage, Charlie (2012), "Plane Bomb Plot Detailed in Support of Life Term," *New York Times*, February 11, p. A15.

Savva, Nikolaos, and Nadia Bianchi-Berthouze (2011), "Automatic Recognition of Affective Body Movement in a Video Game Scenario," web4.cs.ucl.ac.uk. As of April 8, 2013:
http://web4.cs.ucl.ac.uk/uclic/people/n.berthouze/paper/SavvaBerthouze11.pdf

Sayette, Michael A., Kasey G. Creswell, John D. Dimoff, Catharine E. Fairbairn, Jeffrey F. Cohn, Bryan W. Heckman, Thomas R. Kirchner, John M. Levine, and Richard L. Moreland (2012), "Alcohol and Group Formation," *Psychological Science*, Vol. 23, No. 8, pp. 869–878.

Scalora, Mario J., J. Baumgartner, W. Zimmerman, D. Callaway, M. Hatch Maillette, C. Covell, R. Palarea, J. Krebs, and D. Washington (2002), "An Epidemiological Assessment of Problematic Contacts to Members of Congress," *Journal of Forensic Science*, Vol. 47, No. 6.

Scarpa, Angela, S. C. Haden, and Akiho Tanaka (2010), "Being Hot-tempered: Autonomic, Emotional, and Behavioral Distinctions Between Childhood Reactive and Proactive Aggression," *Biological Psychology*, Vol. 84, No. 3, pp. 488–496.

Scarpa, Angela, and Adrian Raine (1997), "Psychophysiology of Anger and Violent Behavior," *Psychiatric Clinics of North America*, Vol. 20, No. 2, pp. 375–394. As of April 22, 2013:
http://home.mdconsult.com/das/journal/view/N/9441384?ja=91610&PAGE=1.html&ANCHOR=top&source=HS,MI

Schauer, Maggie, and Thomas Elbert (2010), "Dissociation Following Traumatic Stress," *Zeitschrift für Psychologie/Journal of Psychology*, Vol. 218, No. 2, pp. 109–127. As of April 22, 2013:
http://mandaladesign.com.au/startts/winter2011/schauer-elbert-dissociation.pdf

Scherer, Klaus R. (1970), "Analyse Der Aussersprachlichen. Aspekte Von Interaktionsverhalten," *Buske.*

Schmitt, Eric, and Eric Lipton (2010, January 1), "Focus on Internet Imams as al Qaeda Recruiters," *New York Times*, p. A14.

Sciolino, Elaine, and Don van Natta (2005, July 10), "For a Decade, London Thrived as a Busy Crossroads of Terror," *New York Times.*

Sebanz, Natalie, and Maggie Shiffrar (2009), "Detecting Deception in a Bluffing Body: The Role of Expertise," *Psychonomic Bulletin & Review*, Vol. 16, No. 1, pp. 170–175. As of April 23, 2013:
http://web.mac.com/gknoblich/page3/assets/2009_SebanzShiffrar.pdf

Sengupta, Somini (2009, January 7), "Dossier Gives Details of Mumbai Attacks," *New York Times*, p. A5. As of April 20, 2013:
http://www.nytimes.com/2009/01/07/world/asia/07india.html?_r=2&

Sentz, Kari, and Scott Ferson (2002), *Combination of Evidence in Dempster-Shafer Theory*, Albuquerque, N.M.: Sandia National Laboratories.

Shafer, Glen A. (1976), *A Mathematical Theory of Evidence*, Princeton University Press.

Shafer, Glenn, and Judea Pearl, eds. (1990a), *Bayesian and Belief-function Formalisms for Evidential Reasoning: A Conceptual Anaysis*, San Mateo, Calif.: Morgan Kaufmann Publishers Inc.

Shafer, Glenn, and Judea Pearl, eds. (1990b), *Readings in Uncertain Reasoning*, San Mateo, Calif.: Morgan Kaufman

Shannon, Claude (1948), "A Mathematical Theory of Communications," *Bell Systems Technical Journal*, Vol. 27, pp. 370–423; 623–656.

Shaughnessy, Larry (2012, July 20), "Hasan's E-Mail Exchange with al-Awlaki: Islam, Money and Matchmaking," CNN.com. As of April 20, 2013:
http://security.blogs.cnn.com/2012/07/20/
hasans-e-mail-exchange-with-al-awlaki-islam-money-and-matchmaking/

Shaver, Russell D., and Michael Kennedy (2004), *The Benefits of Positive Passenger Profiling on Baggage Screening Requirements*, Santa Monica, Calif.: RAND Corporation. As of April 8, 2013:
http://www.rand.org/pubs/documented_briefings/DB411.html

Simonite, Tom (2012, July/August), "What Facebook Knows," *Technology Review.*

Skeem, Jennifer L., D. Polaschek, C. Patrick, and S. Lilienfeld (2011), "Psychopathic Personality: Bridging the Gap Between Scientific Evidence and Public Policy," *Psychological Science in the Public Interest*, Vol. 12, No. 3, pp. 95–162.

Skillicorn, David B., and Ayron Little (2010), "Patterns of Word Use for Deception in Testimony," *Annals of Information Systems*, Vol. 9, pp. 25–39.

Smarandache, Florentin, and Jean Dezert, eds. (2009a), *Advances and Applications of DSmT for Information Fusion*, Rehoboth: American Research Press.

——— (2009b), "An Introduction to DSmT," in Florentin Smarandache, and Jean Dezert, eds., *Advances and Applications of DSmT for Information Fusion*, Rehoboth, N.M.: American Research Press. As of April 22, 2013: http://fs.gallup.unm.edu//IntroductionToDSmT.pdf

Smith, Brent L., Kelly R. Damphousse, and Paxton Roberts (2006), *Pre-incident Indicators of Terrorist Incidents: The Identification of Behavioral, Geographic, and Temporal Patterns of Preparatory Conduct*, NCJ 214217, Rockville, Md.: National Institute of Justice.

Speckhard, Anne (2008), "The Emergence of Female Suicide Terrorists," *Studies in Conflict and Terrorism*, Vol. 31, No. 11, pp. 995–1023.

Speckhard, Anne, Beatrice Jacuch, and Valentijn Vanrompay (2012), "Taking on the Persona of a Suicide Bomber: A Thought Experiment," *Perspectives on Terrorism*, Vol. 6, No. 2, pp. 51–73. As of April 22, 2013: http://www.terrorismanalysts.com/pt/index.php/pot/article/view/speckhard-taking-on-the-persona/html

Stewart, Mark G., and John Mueller (2012), "Terrorism Risks and Cost-Benefit Analysis of Aviation Security," *Risk Analysis*, pp. 1–15.

Stiglitz, Joseph (2010), "A Trade War with China Isn't Worth It," *Guardian*. As of April 22, 2013: http://www.guardian.co.uk/commentisfree/cifamerica/2010/apr/07/united-states-china-currency-manipulation

Stirman, Shannon Wiltsey, and James W. Pennebaker (2001), "Word Use in the Poetry of Suicidal and Nonsuicidal Poets," *Psychosomatic Medicine*, Vol. 63, No. 4, pp. 517–522.

Stone, Lawrence D., Carl A. Barlow, and Thomas L. Corwin (1999), *Bayesian Multiple Target Tracking*, Norwood, Mass.: Artech House.

Stone, Lori D., and James W. Pennebaker (2002), "Trauma in Real Time: Talking and Avoiding Online Conversations About the Death of Princess Diana," *Basic and Applied Social Psychology*, Vol. 24, pp. 172–182.

Tausczik, Yia R., and James W. Pennebaker (2010), "The Psychological Meaning of Words: Liwc and Computerized Text Analysis Methods," *Journal of Language and Social Psychology*, Vol. 29, No. 1, pp. 24–54.

Tavernise, Sabrina, and Eric Schmitt (2010), "Virginia Man Is Charged in Plot on Capital Subway," *New York Times*, October 28, p. A23.

Timberg, Craig, and Ellen Nakashima (2012, July 25), "Skype Makes Chats and User Data More Available to Police," *Washington Post*, p. A1. As of April 20, 2013: http://www.washingtonpost.com/business/economy/skype-makes-chats-and-user-data-more-available-to-police/2012/07/25/gJQAobI39W_story.html

Tombak, Mati, Ain Isotamm, and Tõnu Tamme (2001), "On Logical Method for Counting Dedekind Numbers," *Lecture Notes in Computer Science*, Vol. 2138, pp. 424–427.

Trahan, Jason, Todd J. Gillman, and Scott Godstein (2009), "Dallas Bomb Plot Suspect Told Landlord He Was Moving Out," *Dallas Morning News*, September 26.

Turow, Joseph (2011), *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth*, New Haven, Conn.: Yale University Press.

U.S. Army Training and Doctrine Command (2006), *Suicide Bombing in the COE*, DCSINT Handbook No. 1.03. As of April 22, 2013: http://www.fas.org/irp/threat/terrorism/sup3.pdf

U.S. Court of Appeals for the Ninth Circuit (2010, February 2), *U.S. v. Ressam*, Volume 1 of 2 (No. 09-30000), Seattle, Wash.: Ninth Circuit. As of April 20, 2013: http://cdn.ca9.uscourts.gov/datastore/opinions/2010/02/02/09-30000.pdf

U.S. Department of Justice, ed. (2007), *Policing 2020: Exploring the Future of Crime, Communities, and Policing*, Quantico, Va.: Futures Working Group, Behavioral Science Unit, FBI Academy.

——— (2008), *Findings of the Suspicious Activity Report*, Washington, D.C.

U.S. District Court (2009), "Warrant for Arrest of Hosam Maher Husein Smadi." As of April 23, 2013: http://graphics8.nytimes.com/packages/pdf/us/Smadi_complaint.pdf

U.S. House of Representatives (2011, April 6), *Behavioral Science and Security: Evaluating TSA's SPOT Program, Hearing Before the Subcommittee on Investigations and Oversight Committee on Science, Space, and Technology House of Representatives, One Hundred Twelfth Congress, First Session*, Serial No. 112-11, Washington, D.C.: Government Printing Office. As of April 24, 2013: http://www.gpo.gov/fdsys/pkg/CHRG-112hhrg65053/pdf/CHRG-112hhrg65053.pdf

U.S. Senate Committee on Homeland Security and Governmental Affairs (2012), "Zachary Chesser: A Case Study in Online Islamist Radicalization and Its Meaning for the Threat of Homegrown Terrorism."

U.S. Senate (2012), *Federal Support for and Involvement in State and Local Fusion Centers, chaired by Carl Levin and Tom Coburn*. As of April 22, 2013:
http://www.hsgac.senate.gov/download/
report_federal-support-for-and-involvement-in-state-and-local-fusions-centers

Vance, Ashlee, and Brad Stone (2011, November 22), "Palantir, the War on Terror's Secret Weapon," *Bloomberg Businessweek*. As of April 20, 2013:
http://www.businessweek.com/magazine/palantir-the-vanguard-of-cyberterror-security-11222011.html

Vigluicci, Vincent V. (2009), "Calculating Credibility: State V. Sharma and the Future of Polygraph Admissibility in Ohio and Beyond," *Akron Law Review*, Vol. 42, pp. 319–354.

Villar, Gina, Joanne Arciuli, and Helen Paterson (2012), "Vocal Pitch Production During Lying: Beliefs About Deception Matter," *Psychiatry, Psychology, and Law*, Vol. 20, No. 1, pp. 1–10. As of April 22, 2013:
http://www.tandfonline.com/doi/abs/10.1080/13218719.2011.633320

Vossekuil, Bryan, Robert A. Fein, Marisa Reddy, Randy Borum, and William Modzeleski (2002), *The Final Report and Findings of the Safe School Initiative: Implications for the Prevention of School Attacks in the United States*, Washington, D.C.: U.S. Secret Service and Department of Education.

Vrij, Aldert (2006), "Nonverbal Communication and Deception," in Valerie Manusov, and Miles L. Patterson, eds., *Sage Handbook of Nonverbal Communication*, Thousand Oaks, Calif.: Sage Publications, Chapter 18.

——— (2010), *Detecting Lies and Deceit: Pitfalls and Opportunities (2nd Ed.)*, New York: John Wiley & Sons Ltd.

Vrij, Aldert, and Pär Anders Granhag (2012, June), "Eliciting Cues to Deception and Truth: What Matters Are the Questions Asked," *Journal of Applied Research in Memory and Cognition*, pp. 110–117.

Vrij, Aldert, Par Anders Granhag, Samantha Ann Mann, and Sharon Leal (2011a), "Lying About Flying: The First Experiment to Detect False Intent," *Psychology, Crime & Law*, Vol. 17, No. 7, pp. 611–620.

Vrij, Aldert, Sharon Leal, and Samantha Ann Mann (2011b), "A Comparison Between Lying About Intentions and Past Activities: Verbal Cues and Detection Accuracy," *Applied Cognitive Psychology*, Vol. 25, pp. 212–218.

Wasserman, Todd (2012, August 2), "83 Million Facebook Accounts Are Fake," Mashable.com. As of April 8, 2013:
http://mashable.com/2012/08/02/fake-facebook-accounts/

Weinberger, Sharon (2010), "Airport Security: Intent to Deceive?" *Nature International Weekly Journal of Science*, Vol. 465, pp. 412–415. As of April 20, 2013:
http://www.nature.com/news/2010/100526/full/465412a.html

——— (2011), "Terrorist 'Pre-Crime' Detector Field Tested in United States," *Nature International Weekly Journal of Science.* As of April 8, 2013: http://www.nature.com/news/2011/110527/full/news.2011.323.html

Wells, G. L., and E. Olsen (2002), "Eyewitness Identification: Information Gain from Incriminating and Exonerating Behaviors," *Journal of Experimental Psychology, Applied*, Vol. 8, pp. 155–167.

Wells, Linton, and Barry Horowitz (2007), "Implementation of a Methodology for the Prioritizing of Suicide Attacker Recruitment Preferences," *Journal of Homeland Security and Emergency Management*, Vol. 4, No. 2.

White House Commission on Aviation Security (1997), *Final Report to President Clinton of the White House Commission on Aviation Safety and Security.*

Wilkening, Dean (1999), "A Simple Model for Calculating Ballistic Missile Defense Effectiveness," Vol. 8, No. 2, pp. 183–215.

Williams, Carol J. (2001, November 18), "Love Letter Written by Suspected Hijacker Reportedly Surfaces," *Los Angeles Times*. As of April 20, 2013: http://articles.latimes.com/2001/nov/18/news/mn-5616

Willis, Henry H., James J. Bonomo, Paul K. Davis, and Richard Hillestad (2006), *Capabilities Analysis Model for Missile Defense (CAM-MD): User's Guide*, Santa Monica, Calif.: RAND Corporation, not available to general public.

Willis, Larry (2011), *Behavioral Science and Security: Evaluating TSA's Spot Program*, Washington, D.C.: U.S. House of Representatives Subcommittee on Investigations and Oversight. As of April 22, 2013: http://www.gpo.gov/fdsys/pkg/CHRG-112hhrg65053/pdf/CHRG-112hhrg65053.pdf

Willis, Larry E. (2008), *Privacy Impact Assessment for the Experimental Testing of Project Hostile Intent Technology*, Department of Homeland Security. As of April 21, 2013: http://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_st_phi.pdf

Witten, I. H., F. Eiber, and M. A. Hall (2011), *Data Mining: Practical Machine Learning Tools and Techniques*, Elsevier.

Yoon, K. Paul, and Ching-Lai Hwang (1995), *Multi Attribute Decision Making: An Introduction*, New York: Sage.

Zadeh, Lofti A. (1978), "Fuzzy Sets as a Basis for a Theory of Possibility," *Fuzzy Sets and Systems*, Vol. 1, No. 1, pp. 3–28. As of April 22, 2013: http://ece.ut.ac.ir/classpages/F83/Fuzzysets/Papers/PossibilityTheory/Zadeh_Fuzzy_Sets_as_a Basis_for_a_Theory_of_Possibility_BW.pdf

——— (1984), "Review of Books: a Mathematical Theory of Evidence," *The A.I. Magazine*, Vol. 5, No. 3, pp. 81–83.

Zhongfei, Zhang (2002), "Mining Surveillance Video for Independent Motion Detection," in *Proceedings of 2002 IEEE International Conference on Data Mining (ICDM)*, pp. 741–744.

Zhou, Lina, Judee K. Burgoon, Doug Twitchell, and Tiantian Qin (2004), "A Comparison of Classification Methods for Predicting Deception in Computer-Mediated Communication," *Journal of Management Information Systems*, Vol. 20, No. 4, pp. 139–165.

Zubaydah, Abu (no date), "Arch-Terrorst's Recruiting Video." As of April 8, 2013: http://videos.mcclatchydc.com/vmix_hosted_apps/p/ media?id=64088231&item_index=1&all=1&sort=NULL

Zwerdling, Daniel (2009), "Walter Reed Officials Asked: Was Hasan Psychotic," National Public Radio, November 11.

Government organizations have put substantial effort into detecting and thwarting terrorist and insurgent attacks by observing suspicious behaviors of individuals at transportation checkpoints and elsewhere. This report reviews the scientific literature relating to observable, individual-level behavioral indicators that might—along with other information—help detect potential violent attacks. The report focuses on new or nontraditional technologies and methods, most of which exploit (1) data on communication patterns, (2) "pattern-of-life" data, and/or (3) data relating to body movement and physiological state. To help officials set priorities for special attention and investment, the report proposes an analytic framework for discussion and evaluation; it also urges investment in cost-effectiveness analysis and more vigorous, routine, and sustained efforts to measure real-world effectiveness of methods. One cross-cutting conclusion is that methods for behavioral observation are typically not reliable enough to stand alone; success in detection will depend on information fusion across types of behaviors and time. How to accomplish such fusion is understudied. Finally, because many aspects of using behavioral observations are highly controversial, both scientifically and because of privacy and civil-liberties concerns, the report sharpens the underlying perspectives and suggests ways to resolve some of the controversy while significantly mitigating problems that definitely exist.

RAND

NATIONAL DEFENSE RESEARCH INSTITUTE

www.rand.org

$32.95