Decision Analysis

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Improving the Value of Analysis for Biosurveillance

Henry H. Willis, Melinda Moore

DECISION ANALYSIS
Vol. 11, No. 1, March 2014, pp. 63–81
ISSN 1545-8490 (print) | ISSN 1545-8504 (online)
http://dx.doi.org/10.1287/deca.2013.0283

Improving the Value of Analysis for Biosurveillance

Henry H. Willis
RAND Corporation, Pittsburgh, Pennsylvania 15213, hwillis@rand.org
Melinda Moore
RAND Corporation, Arlington, Virginia 22202, mmoore@rand.org

Biosurveillance provides information that improves decisions about mitigating the effects of disease outbreaks and bioterrorism. The success of biosurveillance depends on the effectiveness of at least four key processes: data collection, data analysis and interpretation, data integration from across organizations, and action (including public responses) based upon results of the analysis. Questions typically arise about whether information from biosurveillance systems represents a threat that justifies a response. To begin answering these questions, the Institute of Medicine Standing Committee on Health Threats Resilience has been undertaking discussions of strategies that the Department of Homeland Security National Biosurveillance Integration Center could use to strengthen its decision support and decision analysis functions. As part of these discussions, this paper applies two standard decision analysis tools to biosurveillance–*decision trees* and *value-of-information* analysis—to assess the implications of strategies to enhance biosurveillance and to improve decisions about whether and how to act after detection of a biosurveillance signal. This application demonstrates how decision analysis tools can be used to improve public health preparedness decision making by developing a road map for how best to enhance biosurveillance through better analytic tools and methods.

Keywords: public health preparedness; applications: government; value of information; decision trees; biosurveillance
History: Received on December 31, 2012. Accepted by former Editor-in-Chief L. Robin Keller on May 12, 2013, after 1 revision. Published online in *Articles in Advance* December 23, 2013.

1. Biosurveillance: A Critical Component of Public Health Preparedness

"Biosurveillance" is a term that has evolved in national security and public health policy in the United States in the decade-plus since the September 2001 terrorist attacks (Figure 1). In this paper we use the term to refer to monitoring biological threats to human or animal health, whether arising from natural, accidental, or intentional causes.

"Biosurveillance" appears to emerge from both public health surveillance and national security surveillance and analysis at the still-evolving nexus of public health emergency preparedness, where these two sectors meet. Figure 2 depicts this nexus and also indicates the alignment of capabilities planning across federal programs, from the Department of Health and Human Services' (HHS) 10 essential public health services, to state and local preparedness capabilities supported by both the HHS and the Department of Homeland Security (DHS), to the DHS National Preparedness Guidelines and capabilities defined in the National Preparedness Goal (NPG).

The National Security Strategy of 2006 explicitly links naturally occurring public health challenges like HIV/AIDS and pandemic influenza to the national security agenda:

> The risks to social order [of such challenges] are so great that traditional public health approaches may be inadequate, necessitating new strategies and responses.... If left unaddressed, [nontraditional security challenges including infectious diseases] can threaten national security. (White House 2006, p. 47)

What are the traditional approaches that no longer suffice? From one side, public health surveillance is considered the cornerstone of public health—it has been defined as the ongoing systematic collection, analysis, interpretation, and dissemination of data regarding a health-related event for use in public health action to reduce morbidity and mortality and to improve health (Thacker 2000), or more simply, systematic information for public health action (Moore

Figure 1 Evolution of Biosurveillance in National Security and Public Health Policy

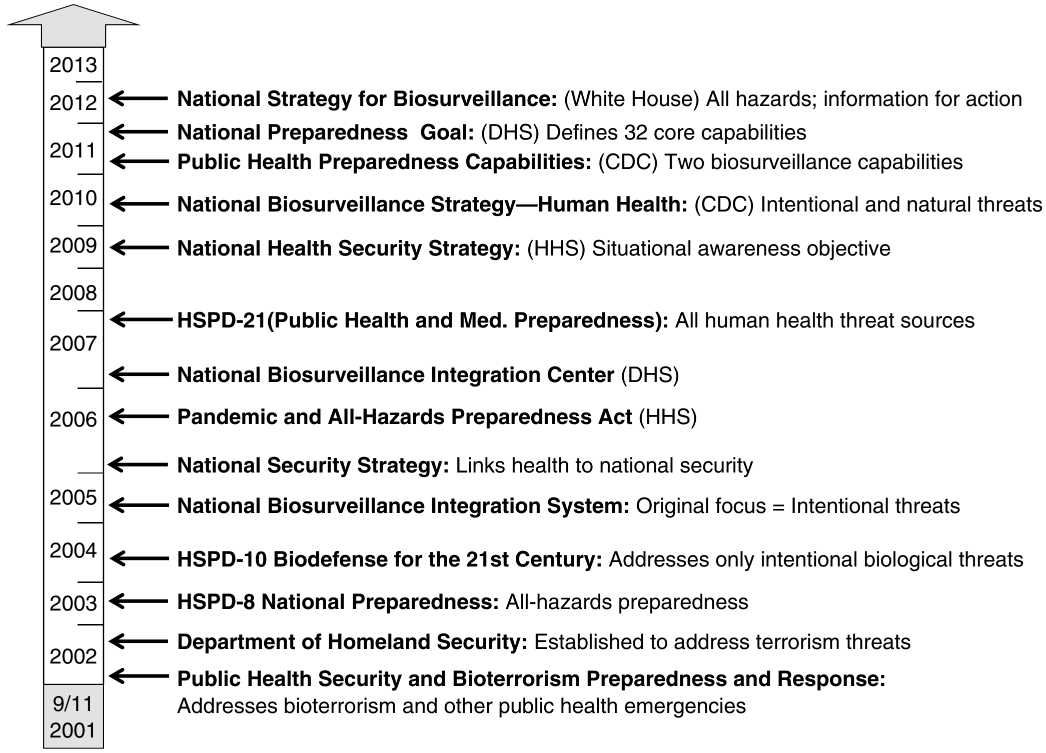

Note. CDC, Centers for Disease Control and Prevention.

Figure 2 "Surveillance" and "Biosurveillance" in Public Health and National Security

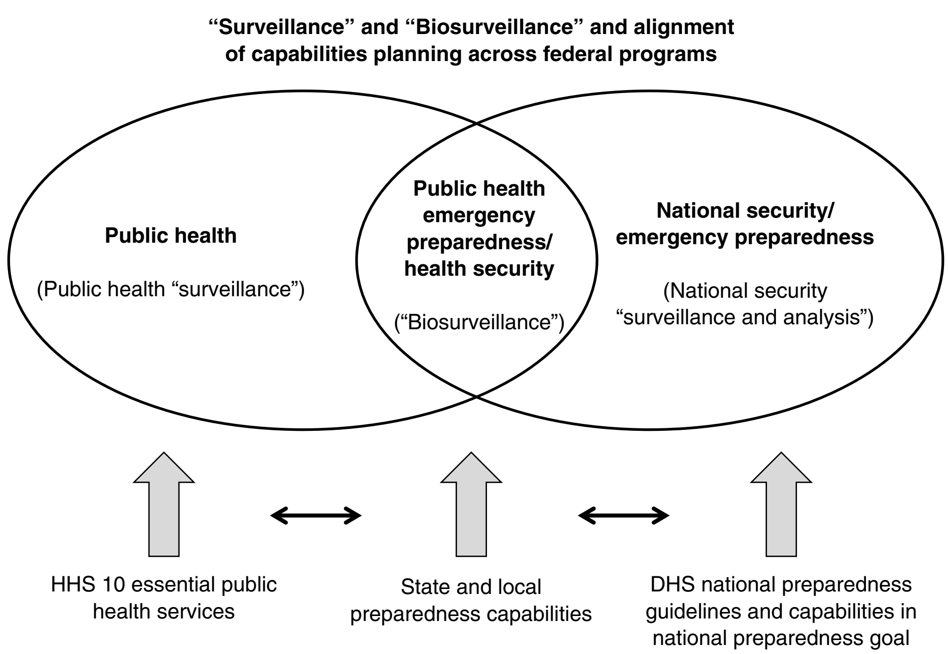

Source: Adapted from Centers for Disease Control and Prevention (2011, p. 3), "Public Health Preparedness Capabilities."

et al. 2008). From the other side, surveillance and analysis within the national security community (e.g., military, intelligence, homeland security, and diplomatic) has typically involved active data collection from different human and technology-oriented sources for purposes of assessing potential threats to security and, based on those threats, taking action to deter or mitigate them.

What new approaches are needed to address health-related threats to national security, and where does biosurveillance fit in? As can be seen in Figure 1, biosurveillance has evolved from a narrow focus on detecting intentional terrorism threats to health or security to a broader focus addressing all-hazard (intentional, accidental, or naturally occurring) threats to human or animal health. Biosurveillance has also become distinctly multisector and multidisciplinary in nature and explicitly oriented toward timely action. Therefore, new biosurveillance approaches must build upon traditional approaches, effectively encompass a broader range of sources and stakeholders, and harness technologies for data integration and near-real-time detection and tracking.

A key question in the public health community has been whether "public health emergency preparedness" is distinctly separate from or inherently part of routine public health practice. The term was initially defined in early 2007 as:

> the capability of the public health and healthcare systems, communities, and individuals, to prevent, protect against, quickly respond to, and recover from health emergencies, particularly those whose scale, timing, or unpredictability threatens to overwhelm routine capabilities. (Nelson et al. 2007, p. S9)

Nelson et al. (2007) also indicated that,

> as much as possible, [public health emergency preparedness] should be integrated with and expand upon day-to-day public health practices and build upon existing systems, not developed de novo. (p. S10)

The current U.S. Assistant Secretary for Preparedness and Response emphasizes that public health *system* preparedness is the foundation for public health *emergency* preparedness (Lurie 2011). Most agencies now refer to, and consider, "public health preparedness" as encompassing the routine public health systems and practices that are also essential in responding to an emergency affecting human or animal health. As such, both effective systems and preparedness planning are essential, and they are linked.

Central to preparedness planning are the information systems to detect and monitor threats and track resources over the course of disaster emergence, response, and recovery. Situational awareness, including biosurveillance, is a critical element of public health preparedness and one of 10 objectives in the HHS's National Health Security Strategy (NHSS; U.S. Department of Health and Human Services 2009). As noted in the NHSS, information contributing to situational awareness comes from all sectors and all levels of government, as well as from international and community-based sources. Stemming from the NHSS, the Centers for Disease Control and Prevention within the HHS subsequently issued the National Biosurveillance Strategy for Human Health (originally in 2008, updated in 2010; Centers for Disease Control and Prevention 2010) and guidance regarding public health preparedness capabilities (Centers for Disease Control and Prevention 2011). Later during 2011, the DHS issued the National Preparedness Goal, which principally defined a refined set of core capabilities that superseded their 2007 Target Capabilities List. One NPG capability explicitly includes biosurveillance (screening, search, and detection), and two others inherently do so (public information and warning, and situational assessment). Both HHS and DHS have supported state and local preparedness grants over the past several years, including surveillance activities, to enhance public health, health security, and national security. In July 2012, the White House issued the National Strategy for Biosurveillance (White House 2012); it focuses on "essential information for better decision making at all levels of government" (p. 1) and ties together several important threads: data collection, integration, dissemination/alert, and forecasting related to all biological threats—naturally occurring, accidental, or intentional disease events in humans, animals, or plants.

1.1. The Challenge of Biosurveillance in Preparedness

Biosurveillance involves a number of steps, including relevant inputs and outputs (Figure 3). The ultimate goal of biosurveillance is to provide information that improves decisions intended to mitigate the effects

Figure 3 The Biosurveillance Process

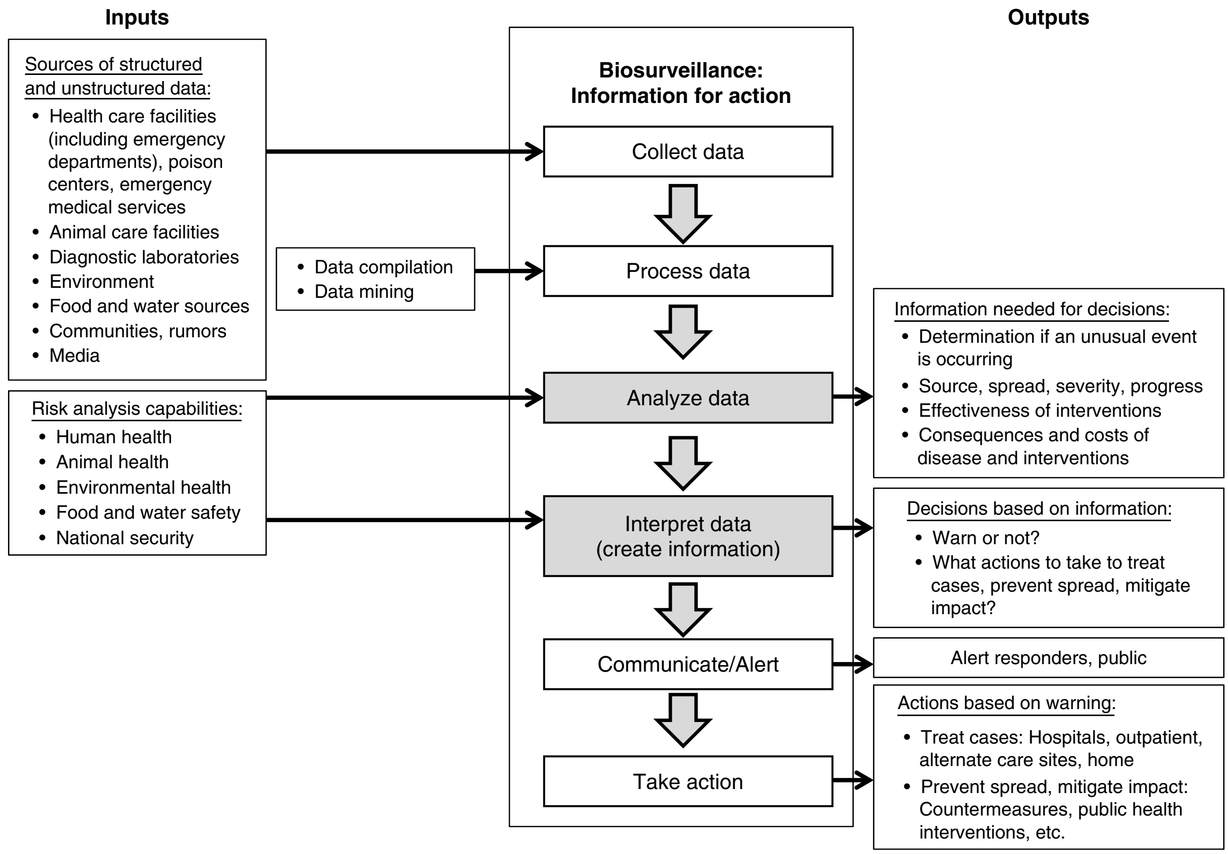

of disease outbreaks and bioterrorism. Thus, success of biosurveillance depends on the effectiveness of at least four key processes: data collection and processing, data analysis, data interpretation, and action based upon results of the analysis. Data integration from across organizations contributes to each of these processes and is arguably the most difficult challenge to biosurveillance, given the constraints posed by limitations and incentives distinct to the perspective of each organization within the response community (National Research Council 2011). This paper, however, focuses on how analysis can support biosurveillance when integration and collaboration are successfully addressed. Ultimately, a suitable analytic framework, such as that described in this paper, might be helpful in developing better integration and collaboration.

Collecting information is at the root of all biosurveillance activity. The quality of incoming data, and knowledge of that quality, is a sine qua non for high-quality biosurveillance. Biosurveillance information can be collected passively (i.e., reports are initiated by the data source, such as routine reporting of diagnosed cases of specified diseases seen at selected health-care facilities or of test results from public health laboratories) and/or actively (i.e., reports are initiated through outreach to data sources, with examples including environmental monitoring for potential bioterrorism pathogens and data mining from media sources). Surveillance data are processed either periodically (e.g., routine weekly disease surveillance) or continuously (e.g., environmental monitoring, data mining) and analyzed to determine whether an unusual health-related event has occurred. Translating a signal of an event into guidance for action requires interpreting the data analysis to understand the significance of the event and implications for action; ideally, this takes into account potential risks from the event, responses to the event, and consequences of the responses.

Analytic methods also play an essential role for collection and processing. Examples of approaches that are encompassed in this area include methods of encoding data in clinical settings (e.g., signs and symptoms versus coded diagnoses), algorithms for recognizing patterns of disease in syndromic surveillance or from unstructured data, and algorithms for identifying signals of pathogen release from distributed sensors. The effectiveness of these analytic methods, the disease prevalence, and the underlying signal-to-noise ratio of the disease event in the environment within which it occurs determine the likelihood that a signal actually reflects an event of interest, i.e., its *positive predictive value*. These challenges surrounding detecting a signal of a disease event remain important analytic issues and have been identified in previous studies by the NRC (National Research Council 2011).

Although data collection and processing are the foundation for biosurveillance and contribute directly to the quality and timeliness of data analysis and interpretation, the main focus of this paper is on how data analysis and interpretation improve responses to disease events. Questions typically arise about whether information from biosurveillance systems represents a threat that justifies a response. What biosurveillance signals rise to the level of requiring action? Are there circumstances in which action is not warranted? What actions should be taken? What is the value of improved biosurveillance in terms of the cost of new data and analysis, and the effect that information has on decisions? These are the critical questions addressed in this paper.

Answering these questions requires addressing several data and analytic challenges fundamental to public health security. First, communities must prepare for many types of naturally occurring biological threats, from food-borne outbreaks to pandemics, and terrorism-related events from any of a number of biological agents or toxins. Each of these events has a different—and for some, very low—likelihood of occurring, but together they represent a large number of scenarios to be considered. Second, these events affect communities in many ways (e.g., at a minimum through public health and financial consequences), and those consequences may affect some groups or places more than others. Third, for any given scenario, combinations of responses may be appropriate and will need to be compared. Fourth, decision making for disease events is dynamic. As a disease event unfolds, decision makers must seek and incorporate new information that unveils the characteristics of the disease (such as its virulence and infectiveness) and the effectiveness of the ongoing response. Finally and fortunately, history does not provide a body of evidence from which to evaluate the effectiveness of response to pandemics and bioterrorism because such events have been infrequent.

As a result, planning for response to these events is an exercise in balancing priorities and resources for disease prevention, surveillance, and response. Choices must be made about the appropriate balance between prevention and response capabilities. Similarly, choices must be made about whether planning addresses events that are more commonly observed (such as seasonal flu outbreaks), events that are more catastrophic though less likely to occur (such as bioterrorism attacks), or both to an equal extent.

1.2. Analysis and the "Value-of-Information" in Biosurveillance

Improved biosurveillance will not eliminate the complexity decision makers must confront because of the uncertainty and low likelihood of catastrophic disease events or the multiplicity of response options and event consequences. However, the process of developing an analytic framework that structures choices inherent in disease surveillance and response and clarifies understanding of the factors that influence these choices could be useful for identifying ways to improve biosurveillance.

To begin this process, the IOM Standing Committee on Health Threats Resilience has been undertaking discussions of strategies that the DHS National Biosurveillance Integration Center could use to strengthen its decision support and decision analysis functions. This paper contributes to these discussions by applying two standard decision analysis tools to biosurveillance—*decision trees* and *value-of-information* analysis[1]—to assess the implications of strategies to

[1] For those not familiar with decision trees or value-of-information analysis, each tool is explained when it is introduced in subsequent sections of this paper.

enhance biosurveillance and to improve decisions about whether and how to act after detection of a biosurveillance signal.

Section 2 frames the interpretation of biosurveillance signals using a decision tree. This framing highlights three types of information that are needed to better inform response to disease events. Several analytic approaches are applicable to each type of information. Section 3 describes each category of analytic methods that could support biosurveillance decision making and also the role of decision support tools that could be used to interpret the results obtained from them.

Within decision analysis, the standard approach for evaluating whether to invest in acquiring new information is to consider how that information will improve the expected outcomes of decisions; i.e., the value of information. Section 4 describes how this framing can be applied to setting a research agenda for improving biosurveillance. This paper concludes with observations and recommendations on initial steps that could be taken to develop the analytic capabilities presented.

2. A Decision Analytic Perspective of Biosurveillance

Decision analysis is the disciplined study of how people should make important decisions, how they in fact do make decisions, and how tools can help people make better decisions. The theoretical basis underpinning the discipline of decision analysis for how people should make decisions is *expected utility theory*. This theory states that choosing the option with the greatest overall value is the best decision if a person has preferences among options, and that those preferences are logically consistent when choosing among options, depending on the probability of the option occurring, and not influenced by the addition of other options (von Neumann and Morgenstern 1947). Though expected utility theory is useful for identifying what decisions people should make in their best interest, Tversky and Kahneman (1974) showed that in practice our probability judgments are influenced by what we experience, how information is presented to us, and simple rules, i.e., heuristics, which we tend to fall back on when confronted with complex decisions. With the goal of helping people make smarter choices, decision analysts have developed tools and methods to describe choices and help people navigate these biases and decision-making heuristics (Keeney 1992, Hammond et al. 1999, Raiffa 1997, Howard and Matheson 1984).

2.1. What Decisions Are Inherent in Biosurveillance?

Biosurveillance aims to inform when an unusual disease event is happening by collecting, processing, and analyzing data, and, if it is, to help answer two questions:

- Is action required?
- What type of intervention should be initiated?

Answering these questions requires making decisions based on analysis and interpretation of biosurveillance data.

Figure 3 highlights these two decisions that are fundamental to the analysis and interpretation steps in biosurveillance. The first decision to be made is whether to act based on a signal. If the signal is interpreted as being too ambiguous or of too little concern, then no action may be warranted, and normal biosurveillance activities may carry on or additional information may be collected to verify the initial analyses and try to reduce the ambiguity. However, if the signal of an event warrants action, the second decision required is about what type of intervention should be implemented.

Logically, the first decision of whether to intervene depends on the choices in the second of how to intervene. If no feasible interventions exist, then the decision-related questions are moot. The decisions of whether and how to act therefore depend on how feasible interventions will change the expected outcomes of the suspected event. How to intervene also depends on several factors including what type of event is suspected, what options for intervention exist, how effective and costly (in both monetary and nonmonetary terms) the interventions are, and what the consequences will be if the event does not unfold as expected. Figure 4 captures these decisions using a common decision analysis tool: a decision tree.[2]

[2] A decision tree portrays the connections between decision points (shown as squares) and the probabilistic outcomes (shown as circles) that evaluations of decisions depend upon. To use a decision

Figure 4 A Decision Tree Perspective of Analysis and Interpretation Steps of Biosurveillance When Reacting to a Signal of an Unusual Event

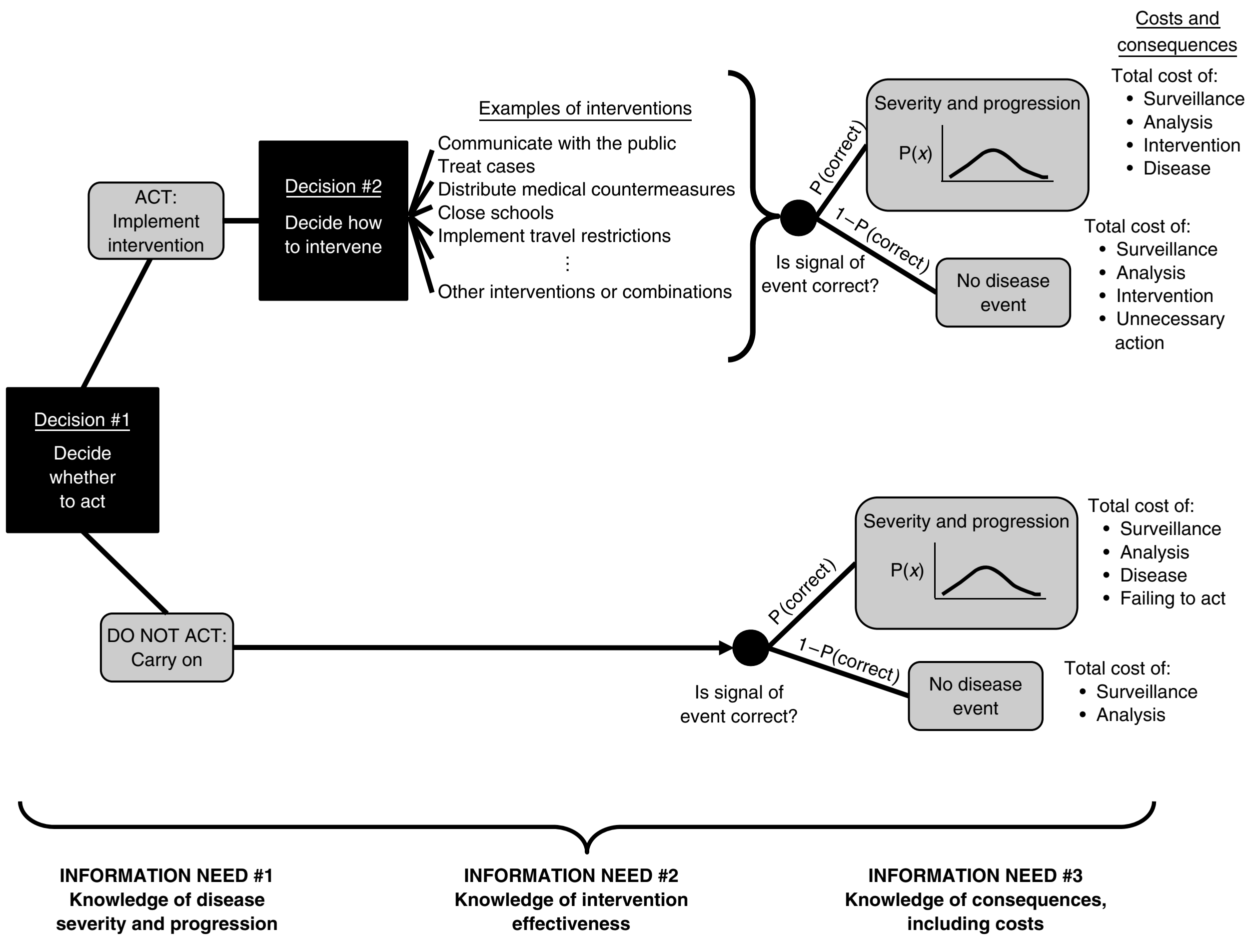

Figure 4 depicts a simplified view of biosurveillance decisions. As depicted, the decisions do not portray the full complexity resulting from uncertainty and low frequency of events, the temporal dynamics of diseases, or the full range of disparate consequences of concern. Despite these simplifications, the decision tree in Figure 4 highlights two unique aspects of biosurveillance decision making. First, the second decision is actually a complex choice among portfolios of interventions. Figure 4 lists several examples of interventions that are intended to reduce the spread of disease and/or treat those who have become ill.

Examples of such interventions include the following:

• Communicating with the public to provide facts and instructions, such as when a contaminated shipment of food has reached a region's grocery stores

• Treating cases at conventional health-care delivery points such as a hospital or provider's office

• Issuing restrictions to limit the extent to which infected people interact with susceptible populations such as closing schools or issuing travel restrictions

In practice, the public health response to a disease outbreak will typically include a combination of these types of interventions. Thus, analysis and interpretation of biosurveillance data must consider the expected effects of response strategies that constitute

tree to evaluate choices, one calculates the expected outcomes of choices working from right to left. The expected outcome is evaluated as the sum of the product of probability and consequences across all consequences at each probability node. At each decision point, the calculation selects the branch of the tree with the most favorable outcome. That planned decision is assumed when evaluating choices to the left in the tree.

a portfolio of interventions, not each intervention on its own.

Second, the signal resulting from biosurveillance data collection and processing might be ambiguous with regard to the precise nature and magnitude of the potential threat. The potential courses of action will vary based on several factors—what the pathogen/disease is, how communicable and deadly it is, in whom and how far the outbreak has spread, and how it might evolve. The accuracy and timeliness of the signal are critical to the decisions regarding whether and how to act. Accuracy of the signal depends upon the sensitivity of the biosurveillance methods (i.e., the probability that if a disease is occurring, biosurveillance will detect it), the specificity of biosurveillance (i.e., the probability that if a disease isn't occurring, biosurveillance will not indicate one), and the prevalence of the disease. Together these factors determine the likelihood that a signal that a disease occurring is correct, i.e., the positive predictive value of biosurveillance. Figure 4 depicts such outcomes as dichotomous, probabilistic events in which either the signal is correct or not, where a correct signal means detection of a true event and an incorrect signal is a false positive—an indication from biosurveillance that there is a disease event when no underlying disease event is actually underway. In reality, this event is a distribution of possible evolutions of the disease event based on the pathogen itself, how interventions undertaken and behaviors of the population influence the spread of disease, and the inherent efficacy of treatment and other interventions.

At the same time, Figure 4 misrepresents one aspect of biosurveillance analysis and interpretation. This figure presents decision making as two discrete decisions. In reality, analysis and interpretation is a continuous process that begins with data collection and processing indicating that a disease event may be occurring, and continuing throughout the event until the situation is resolved and progressing to search for new emerging events. From this perspective, decision making is an iterative cycle of sensing, deciding, and responding in which the decision determines the next response (Figure 5).

In this type of iterative process, the cumulative cost of false alerts is an important consideration for biosurveillance and future research into approaches to biosurveillance. An overly sensitive screening tool

Figure 5 **Biosurveillance Analysis and Interpretation as an Iterative Decision Process**

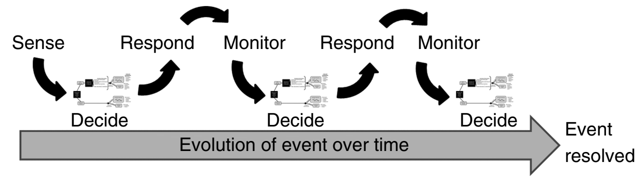

producing warnings that prove frequently unwarranted over time leads to accumulation of unnecessary response costs and could erode confidence in the value of surveillance. The simplified examples in the remainder of this paper, which analyze a one-time decision, do not capture these costs. Rather, they illustrate how the tools of decision analysis can be used to evaluate efforts to improve biosurveillance and can be extended to reflect decision making in a continuous process.

2.2. What Information Is Needed for Biosurveillance Analysis and Interpretation?

The decision tree depicted in Figure 4 points to three types of information required to analyze and interpret the signals that result from biosurveillance data collection and processing:

- Knowledge of disease severity and progression
- Knowledge of interventions available and their effectiveness
- Knowledge of consequences, including costs

Understanding each type of information is a starting point for considering what analysis can best provide it.

2.2.1. Knowledge of Disease Severity and Progression. As described in §2.1, the severity and progression of disease are influenced by the characteristics of the pathogen and the characteristics and behaviors of the population. The outcomes of these factors are captured in assessments of what populations and subpopulations are infected (both geographically and demographically) and how the disease progresses within and across these groups over time. To decide whether and how to intervene, it is necessary to know how a disease event is expected to unfold over time.

2.2.2. Knowledge of Interventions Available and Their Effectiveness. Interventions aim to quickly treat infected people and reduce the number of people who get exposed. The coverage and effectiveness of interventions determine the actual outcomes in terms of the progression of a disease outbreak. Just as disease progression can vary across populations, so can the effectiveness of interventions. Analyzing the decision tree for biosurveillance analysis and interpretation requires understanding how intervention strategies change the progression of a disease event.

2.2.3. Knowledge of Consequences, Including Costs. The consequences of a disease outbreak arise in many forms. Making smart decisions about whether and how to intervene requires an understanding of the magnitude of each type of consequence. These consequences include the following.

- *Consequences and costs of the disease*. The first sets of consequences are the deaths and injuries attributable to the disease event. Placing values on these consequences can be deeply controversial, though doing so increases the transparency of decisions. Drawing upon the economics literature related to estimating the value of life, these consequences can be monetized. Recent studies suggest that results from the health-effects valuation literature for environmental and safety hazards need to be augmented for bioterrorism and pandemics because people are more concerned about health effects from bioterrorism and pandemics than about effects of other health and safety incidents with equivalent expected morbidity and mortality effects (Viscusi 2009). Thus, estimates of disease costs should account for differences in how people perceive bioterrorism and pandemic risks compared to other risks. Furthermore, to be comprehensive, it is important to consider two components of the cost of disease that are frequently ignored. First, just as interventions can spur changes in consumption and behavior, so can the disease itself. If fear of disease keeps people from going to work, this could be attributed to be a cost of the disease. Second, fear of the disease could also have mental health consequences, leading to increased stress and other forms of psychological trauma. Though documented with respect to terrorist events and disasters, mental health trauma is often not counted among the costs of disease because it remains relatively poorly understood. Disagreement about and poor understanding of how some of these consequences are valued warrants further study of how to value the costs of disease.
- *Costs of interventions*. Another set of consequences of a decision to intervene involves the costs of implementing the intervention. These costs include resources required to acquire supplies and equipment, to conduct associated planning and training, and to support the personnel required to implement the intervention. Costs of intervention must account for all aspects of design, planning and preparation, implementation, and sustainment of the intervention. The costs of intervention also include economic disruptions associated with the intervention. For example, a decision to issue travel restrictions to the United States from a specific country could affect businesses in several ways. Cancellation of flights because of the travel restrictions could lead to business interruptions attributable to the intervention. If travelers decided to cancel travel plans to countries nearby the country targeted by travel restrictions out of fear of the disease outbreak, these changes in behavior and consumption could be attributable to the intervention. Estimating these derivative economic damages is tricky because of poor understanding of the public responses to disease incidents and the importance of distinguishing between true damages and transfers from one party to another. For example, if travelers cancel a trip to the affected country but instead travel to another part of the world, this is a loss for one country and a gain for another. From a societal cost perspective, this scenario does not reflect damages unless the traveler would have actually preferred to go to the targeted country and is left worse off being stuck with a second-best choice, or the transfer leads to a distribution of costs and benefits that is perceived to be inequitable.
- *Costs of surveillance and analysis*. This paper focuses mainly on how improving biosurveillance analysis can improve decisions. However, conducting such an analysis incurs costs. When considering whether to conduct additional activities to improve any aspect of biosurveillance—data collection, processing, analysis, or interpretation—it is necessary to consider the costs of developing and implementing them.
- *Consequences and costs of getting it wrong*. A consequence of uncertainty surrounding biosurveillance is that responses to signals of a disease event may

retrospectively be judged to be incorrect. If a disease event does not materialize with the severity projected, as was the case with the H1N1 influenza pandemic of 2009, the public can view interventions as unnecessary and wasteful. If the government fails to intervene quickly or intensely enough, as was the case with the SARS response in China in 2003, the public may view the response to be inadequate. In both cases, the result can be a loss of public confidence that can jeopardize other initiatives or the effectiveness of responses to future incidents. Thus, although these consequences can be difficult to evaluate, they could be a pivotal factor in some decisions, and further study of the public responses would be valuable.

2.3. Summary

Framing analysis and interpretation tasks of biosurveillance using a decision tree focuses the information needs of biosurveillance on answering two key questions: whether to act and how to act.

In addition, this framing identifies three types of information needed to optimally address these two questions. When considering which analytic approaches can most benefit biosurveillance analysis and interpretation, we concentrate on those analyses that provide the three types of information described in §2.2.

3. Analytic Approaches to Supporting Biosurveillance Decision Making

Modeling, simulation, and quantitative analysis are being applied in many ways to the study of infectious diseases and public health responses to disease. Researchers across disciplines are using many different approaches that are capable of providing information required for biosurveillance analysis and interpretation. This paper does not attempt to provide a comprehensive review of all modeling approaches or all applications of a specific modeling approach. Table 1 lists a few examples of selected approaches to modeling, simulation, and analysis that can be used to provide information needed for biosurveillance analysis and information. A brief review of these examples demonstrates the breadth of opportunities to use analysis in the field of biosurveillance.

Table 1 Examples of Analytic Approaches That Can Be Used to Provide Information Needed for Biosurveillance Analysis and Interpretation

Information need	Selected analytic approaches	Examples of selected analytic approaches
Disease severity and progression	• Dynamic epidemic progression models • Disease transmission models • Exposure and behavior modeling	Anderson and May (1991) Epstein et al. (2007, 2008) Epstein et al. (2007) Keeling and Danon (2009)
Intervention effectiveness	• Agent-based modeling • Probabilistic risk analysis • System dynamics models • Operations research and analysis	Lee et al. (2010a, b) Carley et al. (2003) Jackson and Faith (2013) Moore et al. (2008)
Costs and consequences	• Economic input–output analysis • Computable general equilibrium • Stated and revealed preference studies	Gyrd-Hansen et al. (2007) Santos et al. (2009) Dixon et al. (2010)
Decision support	• Decision trees and cost benefit analysis • Policy-level models and portfolio analysis tools	Lee et al. (2009) Davis et al. (2008)

3.1. Approaches for Analyzing Disease Severity and Progression

System dynamics models for analyzing disease severity and progression describe the state of populations that are susceptible to a disease, infected with a disease, and recovered from a disease. So-called susceptible-infected-recovered models describe the dynamic progression of an epidemic within a population (Anderson and May 1991). They can be used to assess the total morbidity and mortality of a disease outbreak as well as how that progression evolves in populations of different ages, of different demographic characteristics, with different behavioral risk factors, and in different places (Keeling and Danon 2009).

The results of these models depend on many parameters such as the rate at which susceptible people become infected upon exposure, the frequency with which susceptible people are exposed, the rate at which infected people recover, and the rate at which infected people die. More recent advances to this form of modeling have focused on better understanding some of these parameters. For example, Epstein et al. (2007) modeled disease transmission by describing how changes in travel can affect transmission rates and thus disease progression, and Epstein et al. (2008) modeled how behavioral reactions to disease, including fear, can affect exposure and thus disease progression.

Models and analyses like these can help decision makers anticipate the progression of disease scenarios based upon characteristics of the disease, the exposure scenario, and the population within which the disease outbreak occurs.

3.2. Approaches to Analyzing the Effectiveness of Interventions

To evaluate intervention strategies to curb the consequences of disease outbreaks, decision makers must know how interventions reduce the number of people who become exposed, become sick, or die during a disease outbreak. Analysts have adopted many approaches to estimating this information using models and simulations. A review of selected examples provides a sense of the diversity of approaches that are available and is a starting point for considering the value of each approach in terms of how it improves decision making when responding to disease incidents.

3.2.1. System Dynamics Modeling. Just as the susceptible-infected-recovered models of this class can be used to assess the expected progression of a disease outbreak, they can also be used to estimate the efficacy of interventions. When used in this way, analysts estimate how interventions would influence key parameters in the model, such as how treatment affects the rate of transition from an infected population to a recovered population, how treatment affects the rate of death of those in the infected population, or how travel restrictions affect the rate at which the susceptible population encounters the infected population. Examples of this approach include studies of the efficacy of school closures and employee vaccination programs following pandemic disease outbreaks (Lee et al. 2010a, b)

3.2.2. Agent-Based Modeling. Building off of similar methods used to estimate disease severity and progression, agent-based models describe how the perceptions, incentives, and decisions at an individual level culminate in group behavior (Epstein et al. 2007, 2008). These methods have been used to understand how different assumptions about travel behavior, decisions to seek treatment, or adherence to public warnings and directions affect the outcomes associated with different interventions. For example, one such simulation, BioWar (Carley et al. 2003), simulates the implementation of a large number of public health interventions for infectious diseases so their efficacy can be estimated in several U.S. cities and under different assumptions about background levels of other disease symptoms and behaviors that affect disease transmission and rates of treatment.

3.2.3. Operations Research and Probabilistic Risk Analysis. These methods describe an outbreak and an intervention as a chain of discrete events and reactions to those events, and have been applied to both technical and social problems. These analyses deliver insight by describing the causal relationships between events and how successful completion of each event depends upon the outcome of preceding events. The relationships described include the likelihood of each event, the dependence of each event on outcomes of preceding events, and the time required to complete each event. When these relationships are characterized, operations research and probabilistic risk analysis can be used to estimate the overall expected consequences of a series of events. For example, Jackson and Faith (2012) used a probabilistic risk analysis method, Failure Mode and Critical Effects Analysis, to describe the overall performance of a system designed to dispense medical countermeasures following a mass anthrax exposure event and estimate the likelihood that the preparedness plan will deliver the intended capability. Similarly, Moore et al. (2008) used operations research methods such as analysis of queues to assess how the coverage, timeliness, and accuracy of components of public health surveillance and reporting lead to the overall likelihood that a novel strain of disease is confirmed, and the time required to achieve this result.

3.3. Approaches to Analyzing the Costs and Consequences of Diseases and Interventions

Conventional approaches for estimating the costs and consequences of diseases and interventions account for easily identified and readily tangible factors. These include the costs to establish plans, develop and acquire equipment, and construct facilities, and the costs of personnel time and consumable materials needed to implement plans. Equally important, but sometimes omitted, are costs associated with required training and ongoing maintenance or support required to ensure plan success. Many of the costs for preparedness planning are incurred well in advance of an event; others are incurred "just in time" in response to an event. The costs and feasibility of pre-event and just-in-time preparedness must be factored into decisions regarding whether and how to act in response to a biosurveillance signal, and thus present a decision about how to best balance resources for response planning.

The direct consequences of disease, of course, include the effects of illnesses and deaths. In the health economics literature, the consequences of illness are frequently monetized so that avoidance of these costs can be directly compared to the costs of avoiding life-threatening hazards. The methods in this literature either ask people to state how much value they place on avoiding health and safety consequences or infer this value from compensation people demand (e.g., hazard-adjusted wages). A frequent finding of this literature has been that not all sources of death and illness are valued equally. Though not yet a robust finding in the literature, recent studies suggest that individuals value preventing deaths and injuries from pandemic diseases more than those from more common life-threatening events (Viscusi 2009, Gyrd-Hansen et al. 2007).

While necessary, methods of accounting for the direct consequences of disease and interventions often miss several categories of costs (described in §2) associated with the cost stemming from reactions to the event and disruptions caused by the intervention. Though these costs are more difficult to estimate, several approaches have been used to better understand them.

The most traditional approach to estimating the broader indirect effects of diseases and interventions is input–output economic analysis. The core of this class of models are tables developed by the U.S. Bureau of Economic Analysis describing dependence among sectors of the U.S. economy. These data, for example, describe how economic activity in the transportation equipment sector stimulates activity in primary metals, energy, and agriculture, and vice versa. Similarly, these models can be used to describe how a decrease in economic activity ripples throughout the economy, thus providing an estimate of the overall economic impacts of an event beyond those directly attributed to morbidity, mortality, and response. Santos et al. (2009) used a version of this analysis to analyze the economic consequences of recovery from a pandemic.

The principal criticism of input–output economic models is that the method relies on a steady-state view of the economy. In the short term, this perspective fails to capture how decisions people make to adapt to the disruptions affect the consequences of the event. Some of these decisions could decrease the consequences; for example, people choosing to eat lettuce instead of spinach following a food contamination incident involving only spinach. Others could increase the consequences; for example, people choosing to cancel vacation plans because of a disease outbreak. In the long term, input–output analysis does not capture the economic consequences associated with decisions leading to new ways of doing business; for example, a structural shift in the use of business travel versus telecommuting following a terrorist event.

Another form of the economic impact model was developed to address these issues: computable general equilibrium (CGE) models. These models attempt to capture the dynamic shifts in markets through equations that represent the elasticity of sectors of the economy to changes in demand and supply of things such as labor, natural resources, services, and commodities. When data are available to estimate the parameters of these equations, CGE models provide a more refined description of the economic impacts of a disruption. Dixon et al. (2010) applied these models to evaluate the economic consequences of several types of terrorism-related and naturally occurring incidents, including pandemic flu.

3.4. Decision Support Tools to Support Biosurveillance Efforts

The examples of information needs and analytic methods described in the preceding sections reveal the complexity and importance of biosurveillance.

To interpret and act upon these complex sources of information, decision makers need tools to help them structure choices and understand the implications of their decisions on expected outcomes. Like the decision tree presented in Figure 4, decision support tools provide a type of metainformation that can help improve public health preparedness decisions.

Decision support tools are generally found in two varieties. The most common use of these tools is to evaluate a decision about a single intervention. The decision tree depicted in §2 is one example of this type of tool. Another example of this type of analysis is cost–benefit analysis. Cost–benefit analysis uses a variety of the approaches described in the preceding sections to estimate both the costs of implementing an intervention and the benefits that are expected from the intervention in terms of reduced adverse impact on health and safety. Both decision trees and cost–benefit analysis are commonly used tools in the analysis of public health policies and interventions. For example, a study by Lee et al. (2009) on pandemic influenza prevention in Singapore demonstrates the utility of these tools in this context.

The practical limitation of decision trees and cost–benefit analysis arise when decision makers must evaluate decisions under conditions of extreme uncertainty about the likelihood of events occurring, must plan for a very large set of scenarios, have many different types of interventions that are feasible, and must make decisions within the context of time and fiscal constraints. In these situations, cost–benefit analysis may not help decision makers determine the best course of action (Greenfield et al. 2012).

Two alternative approaches have been developed for problems like this: *policy models* and *portfolio analysis tools*. To address many sources of uncertainty and plan for a large number of scenarios, the use of low-resolution policy models facilitates understanding of the expected outcomes of different intervention strategies comprised of multiple intervention approaches. To address the breadth of uncertainties and scenarios, policy models must allow for rapid evaluation of a large number of relevant scenarios. To be valid, the parameters of the models and results must be consistent with more detailed studies. However, to manage the complexity in communicating the results of this analysis, policy-level models must be used in conjunction with tools that enable evaluation of a portfolio of interventions across a large number of scenarios. Though most applications of policy-level models and portfolio analysis methods have been focused on national security and defense issues, they are in principle also applicable to the problems surrounding naturally occurring and terrorist-related disease outbreaks (Davis et al. 2008).

3.5. Summary

The examples described in this section demonstrate the existence of multiple opportunities for modeling, simulation, and analysis to improve biosurveillance and decision making in the event of disease outbreaks or bioterrorism. An effective strategy for improving analysis used within biosurveillance will draw upon a portfolio of these methods to provide information about expected disease progression, effectiveness of interventions, and the consequences of the diseases and interventions—including their monetary and nonmonetary costs—to help decision makers interpret this information. A starting point for building this portfolio is consideration of the added value from incorporating these analytic methods into biosurveillance information collection, processing, analysis, and interpretation.

4. Assessing the Value of Information in Biosurveillance

The analytic methods described in §3 provide many options for improving the analysis and interpretation of biosurveillance information. One way to choose among these approaches and prioritize efforts to develop analytic capabilities is to consider the value of the information provided by analysis to decision making. *Expected value of perfect information* is a formal concept in the field of decision analysis that describes the extent to which new information about probabilistic uncertainties would cause changes in decisions and as a result leads to better outcomes (Howard et al. 1972).

To understand the applicability of the expected value of perfect information framework to biosurveillance, consider the following simplified example of a public health emergency decision under uncertainty presented in Figures 6–10. This example is designed to illustrate how the value-of-information framework

Figure 6 Decision Tree Depicting the Choice of Whether and How to Intervene Without Any Biosurveillance Information

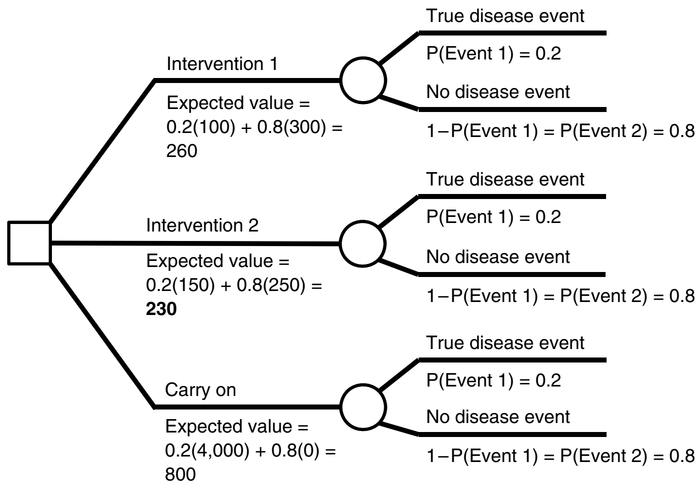

Total		Cost of intervention		Cost of disease		Cost of inappropriate reaction
110	=	100	+	**10**	+	0
300	=	100	+	**0**	+	200
150	=	50	+	**100**	+	0
250	=	50	+	**0**	+	200
4,000	=	0	+	**2,000**	+	2,000
0	=	0	+	**0**	+	0

provides insight into which types of new analysis are most useful. The numbers used in the example are purely notional and would also vary based on the specific disease threat. They can be interpreted as monetized consequences. However, units have not been assigned in an effort to minimize interpretation of the actual values presented. In a real application, the numbers would come from a combination of sources including field data, results of analysis using the methods described in §3, and expert judgment. It will be impossible to monetize all consequences. However, this need not limit the application of analysis. Experience with similar analysis for public policy analysis, such as regulatory cost–benefit analysis, suggests that analysis that monetizes consequences when possible and clearly identifies and describes those consequences that cannot be monetized facilitates informed decision making (Obama 2011).

For this simple example, only two interventions exist (see Figure 6). Intervention 1 is twice as costly as Intervention 2 (see the second column of cost data in Figure 6), but it also leads to one-tenth the consequences of the disease if it occurs (see the third column of cost data in Figure 6). If the health department decides to intervene, but there is no disease event, the community faces serious consequences of an unnecessary intervention, perhaps due to business disruptions and loss of confidence in the county health department. On the other hand, if the health department chooses not to intervene and there is a disease event, the community faces disastrous consequences from deaths and illness (see the third column). The consequences of failing to act also result in severe loss of confidence in the health department (Figure 6, column 4). With this information, the decision tree allows calculation of the expected consequences of each set of decisions. For example, the expected outcome of choosing to act with no information and implementing Intervention 2 is calculated as being $[0.2 \times 150] + [0.8 \times 250]$, or 230.

Now consider a case in which routine collection and processing of biosurveillance data costs the same for all pathways and that cost is assumed to be 10. Suppose the biosurveillance system detects a signal of a possible disease outbreak. However, the signal is not perfectly accurate, and the likelihood that it reflects a true positive event is relatively low. Figure 7 indicates that the probability of a true positive given that the system gave a positive signal of a disease event is $P(E_1/I_1) = 18/34$. The county public health department must decide how to act based on this information.

With this system in place, officials have several choices (see Figure 8). They could choose to use the biosurveillance system, pay the costs to operate it, and make the best choice available when signals of an unusual event are indicated. Instead, they could

Figure 7 **Estimated Performance of the Notional Biosurveillance System Depicted in the Example Analysis**

	Actual situation		Total number of events
	True disease event (E_1)	True no disease event (E_2)	
Biosurveillance system indication			
Positive signal says there is a disease event (I_1)	True positive 18	False positive 16	34
Negative signal says there is no disease event (I_2)	False negative 2	True negative 64	66
Total number of events	20	80	100

forgo the costs of operating the biosurveillance system and make the best choice among interventions, presumably later and after the occurrence of the event is obvious. For simplicity, we assume that this delay in action does not lead to worse disease outcomes. Under these conditions, the optimal decision is to use the biosurveillance system. The expected value of this choice is 157 (i.e., $0.34 \times 197 + 0.66 \times 121 + 10$). This is better (lower) than the expected value of any of the other three interventions without use of the biosurveillance systems (see Figure 8).

Suppose now that the county health department has the opportunity to develop better sources and/or methods for data collection and processing, for example by integrating existing sources of data, adding new data sources, or adding new approaches for identifying a signal in the data collected. These methods would improve the surveillance efforts dramatically, so much so that any signal detected could be perfectly accurate (see Figure 9). Note that the probability $\mathrm{P}(E_1)$ of a disease event occurring remains at 0.2. Should the county health department invest in this new analytic capability and what would it be worth?

To answer this question, consider the revised decision tree presented in Figure 10, which assesses the value with perfect prediction. For the initial assumption that the probability that a disease event will occur is 0.2, it is logical to assume that with perfect prediction the probability is 0.2 that a perfect biosurveillance system will indicate that an event is occurring. Under these conditions, using the biosurveillance system is, not surprisingly, still preferred, and the expected value is lower (i.e., 22). In turn, the expected value of perfect information for biosurveillance is the difference of the expected value with perfect prediction (see Figure 10) and the expected value under uncertainty with no information (see Figure 6), in this case, $230 - 22$, or 208. This calculated value of the expected value of perfect information provides an upper bound on how much should be invested in improving the quality of information produced by biosurveillance efforts.

Value-of-information analysis of analytic capabilities cannot in practice achieve the comprehensive, quantitative precision presented in these examples. This does not, however, invalidate the value-of-information framework as an approach for prioritizing opportunities to improve analysis. Careful accounting of how decisions to develop new analytic capability improve both monetized and nonmonetized consequences can form the basis of a reasoned justification of an analytic agenda, even if estimated improvements are coarse, ordinal assessments.

5. Recommendations

The preceding sections of this paper describe several ways that analytic methods can be used to improve the analysis and interpretation capabilities of biosurveillance and thereby the ability of biosurveillance to inform appropriate decisions regarding whether and how to intervene. Opportunities exist to improve knowledge related to disease severity and progression, effectiveness of interventions, and consequences and costs of disease and interventions. Opportunities also exist to use decision analysis tools to interpret the information resulting from biosurveillance and act upon it. The goal of efforts to improve analytic capability is the same as the goal of biosurveillance—to improve decisions of whether and how to respond to disease incidents.

Accordingly, the value of improved biosurveillance capabilities rests in the value of information that improved analysis brings. More specifically, improving disease surveillance and response involves balancing resources to address a number of fundamental analytic challenges inherent in disease surveillance and response, as described in §1.1. The demonstrations of decision trees and value-of-information analysis presented in this paper illustrate how a decision analytic framework can increase transparency and clarity of

Figure 8 **Decision Tree Depicting the Choice of Whether and How to Intervene Based on an Imperfect Biosurveillance System**

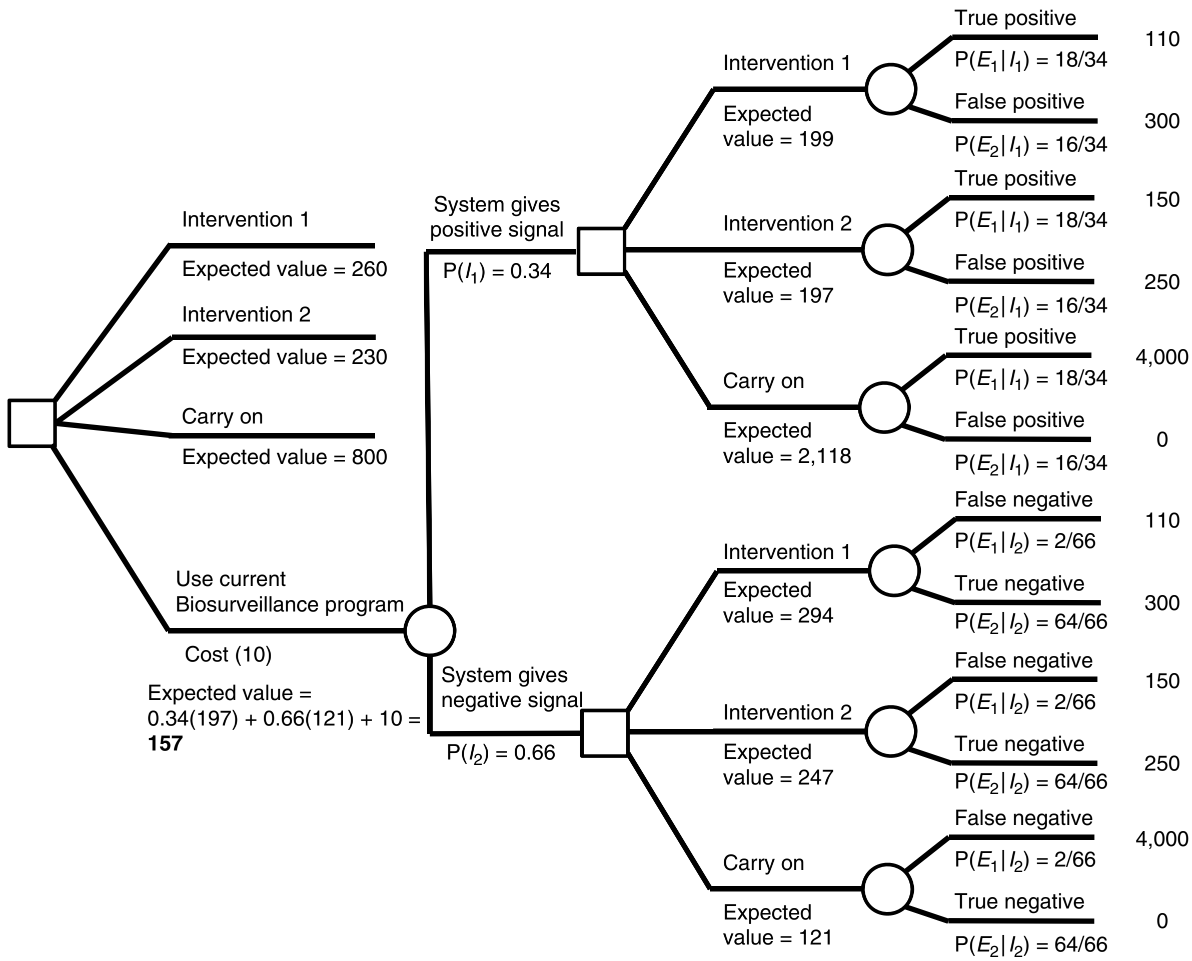

efforts to improve the National Biosurveillance Information Center's analytic capabilities as a contribution to the whole-of-government approach to disease surveillance.

Figure 9 **Estimated Performance of a Perfect Biosurveillance System**

	Actual situation		Total number of events
	True disease event (E_1)	True no disease event (E_2)	
Perfect biosurveillance system			
Positive signal says there is a disease event (I_1)	True positive 20	False positive 0	20
Negative signal says there is no disease event (I_2)	False negative 0	True negative 80	80
Total number of events	20	80	100

In many cases, application of these new analytic capabilities may be possible without a drawn-out development effort. For example, existing data sources and tools from across agencies and programs potentially can be integrated into a more robust national biosurveillance system. Also, given uncertainties inherent in the biosurveillance area, simple approaches may provide adequate insight into choices and may be preferred to more complex, less transparent models that are more difficult to develop.

Developing a detailed road map for improving biosurveillance requires analysis beyond that included in this paper. It requires clearly understanding the objectives and goals of the biosurveillance program; understanding what opportunities exist to leverage existing data and methods and what the demands are from decision makers for biosurveillance information;

Figure 10 **Decision Tree Depicting the Choice of Whether and How to Intervene Based on Perfect Biosurveillance (Assuming No Cost for the Surveillance System)**

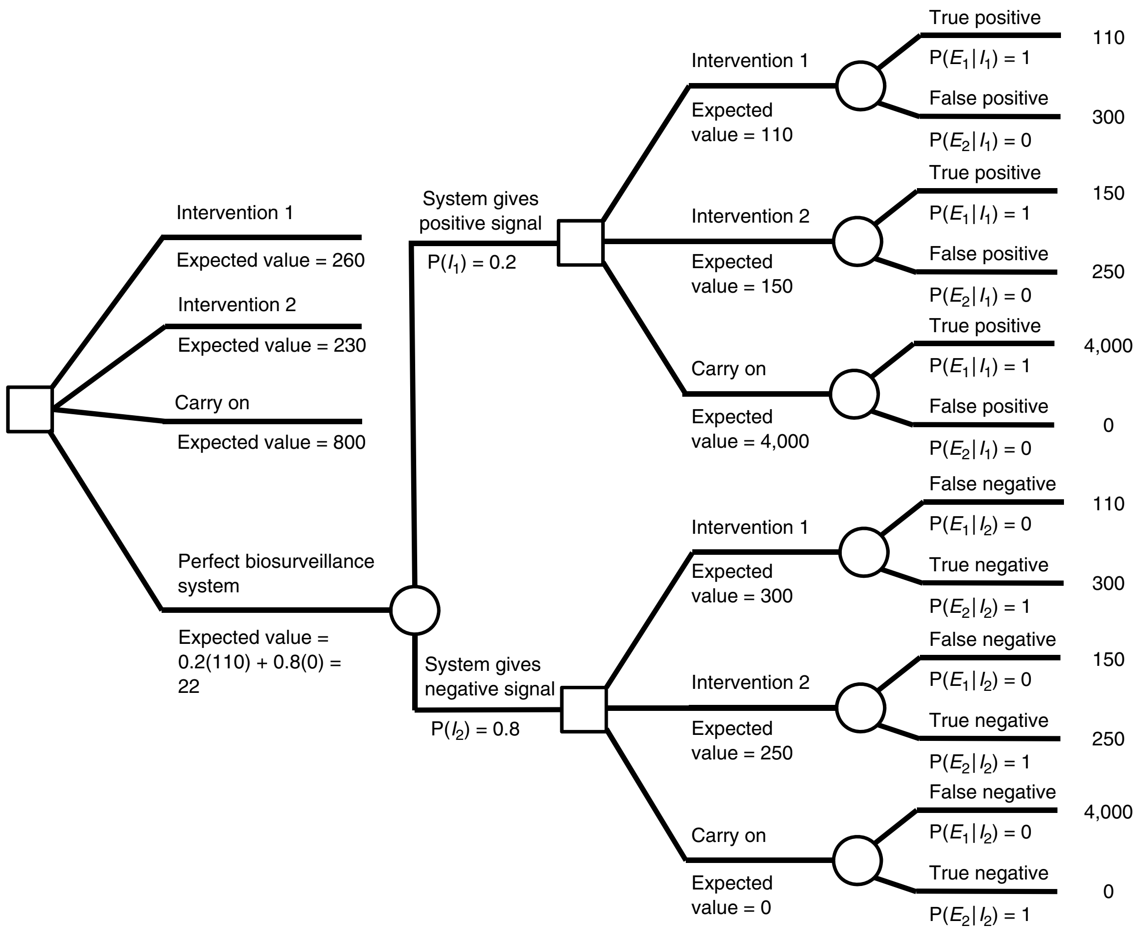

and discussing with decision makers how biosurveillance information can improve how they manage disease outbreaks. To develop a five-year plan to improve specific analytic capabilities for biosurveillance, we recommend completing the following three foundational elements of a strategic plan.

5.1. Clarify Objectives for the Biosurveillance Efforts

The first step in establishing a strategy to improve the analytic capabilities for biosurveillance analysis and interpretation is to clarify the guiding objectives for biosurveillance. For example, the information required to detect and respond to bioterrorism is different than that required for naturally occurring disease outbreaks like seasonal influenza. Because federal, state, and local organizations have distinct capabilities, roles, and responsibilities during response to a disease incident, the information they each require is different. Approaches to surveillance intended to monitor trends in existing diseases are different from those aimed at detecting rare events (Moore et al. 2008). A clear set of objectives will connect a biosurveillance effort to the national strategies it is meant to support, position it with respect to other complementary biosurveillance efforts, and define the capabilities it must provide.

5.2. Identify Available Analytic Capabilities

Table 1 lists several analytic approaches and examples of specific methods associated with each approach. However, this review is not comprehensive. Analysis of these types occurs across the Department of Homeland Security, Department of Defense, and

Department of Health and Human Services. It occurs in universities and nongovernmental organizations. Before any new efforts are made to develop analytic methods to support biosurveillance analysis and interpretation, the body of existing analytic methods should be reviewed as well as the body of existing surveillance data sources and systems that could be "stitched together" to enhance biosurveillance and situational awareness across sectors and levels of government from local to national. A survey of these existing biosurveillance capabilities can be compared to defined biosurveillance objectives and capabilities to identify potential gaps in supporting biosurveillance analysis and interpretation.

5.3. Use Value of Information to Inform Improvements in Biosurveillance Analysis

Clear objectives and awareness of ongoing biosurveillance activities provide a foundation for developing strategies for improving biosurveillance analysis and interpretation. Each of these strategies can provide an approach for integrating new analytic tools with existing tools to achieve clearly stated objectives. The strategies should reflect portfolios of methods to provide insight into the four types of information described in §3 and at the beginning of this section. The value-of-information approach described in §4 provides insight into which strategy will most improve biosurveillance efforts across a range of biothreat scenarios. The selected strategy will provide a road map and priorities for an agenda to improve biosurveillance analysis and interpretation.

5.4. Summary

The steps outlined in this section provide the basis for an analytic agenda and recommend a value-of-information approach to improve biosurveillance analysis and interpretation in response to signals of potential disease events. This agenda will describe strategies for improving biosurveillance analysis and interpretation and the rationale for implementing one or more of these strategies. Ideally, each strategy will include recommendations for developing near-term (one- to two-year) and midterm (three- to five-year) analytic capabilities. Completing these elements will require time and funding. However, a modest investment over approximately six months can clarify what information would be most useful and what gaps in analytic methods currently exist. With this information, efforts to enhance biosurveillance can avoid wasting effort on the development of low-value or redundant analytic methods.

An analytic strategy grounded in the value-of-information approach provides an overarching framework to guide improvements in biosurveillance data collection, processing, and analysis. Although adopting this approach is desirable, it is not sufficient to ensure effective biosurveillance capabilities. As mentioned in §1, effective biosurveillance requires coordination among agencies and organizations responsible for collecting and processing biosurveillance data, detecting signals, making decisions about whether and how to act based on such signals, and implementing appropriate interventions. Thus, as the biosurveillance community considers ways to improve analytic methods at its disposal, continued attention is warranted to coordination and integration of components of the biosurveillance community.

Acknowledgments

The authors thank the anonymous reviewers and the journal's (former) editor L. Robin Keller for constructive comments that improved this paper. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the Institute of Medicine, its standing committee, the United States Department of Homeland Security, the University of Southern California, or CREATE. This paper was supported by funding from the Institute of Medicine's Standing Committee on Health Threats Resilience and by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) [Award 2010-ST-061-RE0001].

References

Anderson RM, May RM (1991) *Infectious Diseases of Humans* (Oxford University Press, Oxford, UK).

Carley K, Fridsma D, Casman E, Altman N, Chang J, Kaminsky B, Nave D, Yahja A (2004) BioWar: Scalable multi-agent social and epidemiological simulation of bioterrorism events. *NAACSOS Conf. 2003* (IEEE SMCA03-11-0274, Pittsburgh).

Centers for Disease Control and Prevention (2010) National biosurveillance strategy for human health. (Originally published in 2008.) Accessed July 29, 2013, http://www.cdc.gov/osels/pdf/NBSHH_v2.pdf.

Centers for Disease Control and Prevention (2011) Public health preparedness capabilities: National standards for state and local planning. Accessed December 5, 2013, http://www.cdc.gov/phpr/capabilities/DSLR_capabilities_July.pdf.

Davis PK, Shaver RD, Beck J (2008) *Portfolio-Analysis Methods for Assessing Capability Options*, MG-662 (RAND Corporation, Santa Monica, CA).

Dixon PB, Rimmer MT, Rose AZ, Verikios G (2010) Effects on the U.S. economy of a serious H1N1 epidemic: Analysis with the USAGE model. General Paper G-202, Centre of Policy Studies, Monash University, Clayton, Australia.

Epstein J, Hammond R, Parker J, Cummings D (2008) Coupled contagion dynamics of fear and disease: Mathematical and computational explorations. *PLoS ONE* 3(12):e3955.

Epstein J, Bobashev G, Goedecke D, Morris R, Wagner D, Yu F (2007) Controlling pandemic influenza: The role of international travel restrictions. *PLoS ONE* 2(5):e401.

Greenfield VA, Willis HH, LaTourrette T (2012) *Assessing the Benefits of U.S. Customs and Border Protection Regulatory Actions to Reduce Terrorism Risk*, CF-301 (RAND Corporation, Santa Monica, CA).

Gyrd-Hansen G, Halvorsen PA, Kristiansen IS (2007) Willingness-to-pay for a statistical life in the times of a pandemic. *Health Econom.* 17(1):55–66.

Hammond JS, Keeney RL, Raiffa H (1999) *Smart Choices: A Practical Guide to Making Better Decisions* (Harvard Business School Press, Watertown, MA).

Howard RA, Matheson JE, eds. (1984) *Readings on the Principles and Applications of Decision Analysis* (Strategic Decisions Group, Menlo Park, CA).

Howard RA, Matheson JE, North DW (1972) The decision to seed hurricanes. *Science* 176(4040):1191–1201.

Jackson BA, Faith KS (2013) The continuing challenge of measuring emergency preparedness: Integrating component metrics to build system-level measures for strategic national stockpile deployment and dispensing operations. *Disaster Medicine Public Health Preparedness* 7(1):96–104.

Keeling MJ, Danon L (2009) Mathematical modelling of infectious diseases. *British Medical Bull.* 92(1):33–42.

Keeney R (1992) *Value-Focused Thinking—A Path to Creative Decisionmaking* (Harvard University Press, Cambridge, MA).

Lee BY, Brown ST, Cooley P, Potter MA, Wheaton WD, Voorhees RE, Stebbins S, et al. (2010a) Simulating school closure strategies to mitigate an influenza epidemic. *J. Public Health Management Practice* 16(3):252–261.

Lee BY, Brown ST, Cooley PC, Zimmerman RK, Wheaton WD, Zimmer SM, Grefenstette JJ, et al. (2010b) A computer simulation of employee vaccination to mitigate an influenza epidemic. *Amer. J. Preventive Medicine* 38(3):247–257.

Lee VJ, Tok MY, Chow VT, Phua KH, Ooi EE, Tambyah PA, Chen MI (2009) Economic analysis of pandemic influenza vaccination strategies in Singapore. *PLoS ONE* 4(9):e7108.

Lurie N (2011) Personal communication with RAND researchers, including Melinda Moore, August 23.

Moore M, Chan E, Lurie N, Schaefer AG, Varda DM, Zambrano JA (2008) Strategies to improve global influenza surveillance: A decision tool for policymakers. *BMC Public Health* 8:186.

National Research Council (2011) *BioWatch and Public Health Surveillance: Evaluating Systems for the Early Detection of Biological Threats: Abbreviated Version* (The National Academies Press, Washington, DC).

Nelson C, Lurie N, Wasserman J, Zakowski S (2007) Conceptualizing and defining public health emergency preparedness. *Amer. J. Public Health* 97(Suppl. 1):S9–S11.

Obama B (2011) Executive Order 13,563: Improving regulation and regulatory review. *Federal Register* 76(14):3821–3823.

Raiffa H (1997) *Decision Analysis: Introductory Readings on Choices Under Uncertainty* (McGraw Hill, Columbus, OH).

Santos JR, Orsi MJ, Bond EJ (2009) Pandemic recovery analysis using the dynamic inoperability input–output model. *Risk Anal.* 29(12):1743–1758.

Thacker S (2000) Historical development. *Principles and Practice of Public Health Surveillance*. Teutsch S, Churchill RE, eds. (Oxford University Press, New York), 1–16.

Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.

U.S. Department of Health and Human Services (2009) National Health Security Strategy. Accessed June 17, 2013, http://www.phe.gov/preparedness/planning/authority/nhss/strategy/documents/nhss-final.pdf.

U.S. Department of Homeland Security (2011) National preparedness goal. http://www.fema.gov/pdf/prepared/npg.pdf.

Viscusi WK (2009) Valuing risks of death from terrorism and natural disasters. *J. Risk Uncertainty* 38(3):191–193.

von Neumann J, Morgenstern O (1947) *Theory of Games and Economic Behavior*, 2nd ed. (Princeton University Press, Princeton, NJ).

White House (2006) The National Security Strategy of the United States of America. Accessed December 5, 2013, http://www.comw.org/qdr/fulltext/nss2006.pdf.

White House (2012) National strategy for biosurveillance. Accessed March 12, 2013, http://www.whitehouse.gov/sites/default/files/National_Strategy_for_Biosurveillance_July_2012.pdf.

Henry H. Willis is director of the RAND Homeland Security and Defense Center and a professor of policy analysis at the Pardee RAND Graduate School. His research applies risk analysis tools to resource allocation and risk management decisions in the areas of national security policy, public health, and emergency preparedness. Dr. Willis serves on the editorial board of the journal *Risk Analysis*. He earned his PhD from the Department of Engineering and Public Policy at Carnegie Mellon University.

Melinda Moore is a senior health researcher at RAND. She earned her MD and Master of Public Health degrees from Harvard University. Since 2005, she has led several projects at RAND related to health security and emergency preparedness, biosurveillance, and public health. In one, she developed a quantitative, interactive decision support tool related to the probability and timeliness of influenza surveillance. Before 2005, Dr. Moore served 25 years as a medical officer at the Centers for Disease Control and Prevention and the HHS Secretary's Office of Global Health Affairs.