

Statistical Methods in Counterterrorism

Game Theory, Modeling, Syndromic Surveillance,
and Biometric Authentication

Alyson G. Wilson
Gregory D. Wilson
David H. Olwell
Editors

Statistical Methods in Counterterrorism

Game Theory, Modeling,
Syndromic Surveillance, and
Biometric Authentication

 Springer

Alyson G. Wilson
Los Alamos National Laboratory
Los Alamos, NM 87545
USA
agw@lanl.gov

Gregory D. Wilson
Los Alamos National Laboratory
Los Alamos, NM 87545
USA
gdwilson@lanl.gov

David H. Olwell
Department of Systems Engineering
Naval Postgraduate School
Monterey, CA 93943
USA
dholwell@nps.edu

Library of Congress Control Number: 2006922769

ISBN-10: 0-387-32904-8 e-ISBN: 0-387-35209-0
ISBN-13: 978-0387-32904-8

Printed on acid-free paper.

© 2006 Springer Science+Business Media, LLC

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Springer acknowledges that there may be no basis for a claim to copyright with respect to a contribution prepared by an officer or employee of the United States Government as part of that person's official duties such as the contribution by William K. Sieber et al. (pp. 215–236).

Printed in the United States of America. (EB)

9 8 7 6 5 4 3 2 1

springer.com

For those who lost their lives on September 11, 2001, and the
men and women fighting the war on terror

Preface

In the months after September 11, 2001, in the aftermath of the attacks on the World Trade Center in New York, counterterrorism became a research interest for a broad range of Western scholars, statisticians among them. At the same time, the U.S. government, still in shock, repeated the same question during multiple hearings in Washington, D.C.: “All the data was out there to warn us of this impending attack, why didn’t we see it?” Data became a large part of the response to 9/11 as Americans tried to regain a rational grip on their world. Data from flight recorders was collected and analyzed, timelines were assembled to parse out explanations of what happened, sensitive data was removed from government websites, and the White House debated what data to release to investigators and the American public. “Data” was a frequently heard term in the popular media, one of the many things that we had to protect from the terrorists, and one of the most important things that we could use to defeat them.

In the statistical community, professionals wondered how they could help the government prevent terror attacks in the future by developing and applying advanced statistical methods. The federal government is a sizable consumer and producer of statistical data, as the 9/11 commission report noted.

The U.S. government has access to a vast amount of information. When databases not usually thought of as “intelligence,” such as customs or immigration information, are included, the storehouse is immense. But the U.S. government has a weak system for processing and using what it has. [KH04, pp. 416–417]

Additionally, government decision-makers are often skeptical about statistics. Understanding that the Washington audience wasn’t always receptive, the statistical community pondered how to put what they knew to work for the country. They felt specially qualified to help decision-makers see the important patterns in the oceans of data and detect the important anomalies in the seemingly homogeneous populations. At a round-table luncheon at the Joint Statistical Meetings in San Francisco in 2003, almost two years after

9/11, a dozen statisticians ate and pondered the same questions. “How do we get in the door?” “How do we get someone to let us help?”

It was hard to get in the door, because Washington was still trying to figure out what a response to terrorism in the homeland would begin to look like. The threat paradigm had shifted enough that no one quite knew what the appropriate questions were, let alone the appropriate responses. Potential bioterrorism is a case in point. Dread diseases like smallpox had been conceptualized and studied as diseases, as public health problems, and as potential battlefield weapons, but had not been extensively studied as agents terrorists might set loose in a major population center. When a set of anthrax mailings followed close on the heels of the World Trade Center bombings, it was as if our world-view had been fractured. Many old questions of interest faded away, many new ones appeared, others were yet to be discovered. Biologists, epidemiologists, biostatisticians, public health experts, and government decision-makers woke up the next day wondering where to begin. The same was true across many fronts and many lines of inquiry in those months. The U.S. government wound up organizing an entirely new Department of Homeland Security to address the raft of new problems that emerged after 9/11. In the decision-maker’s estimation, the new problems were different enough that existing structures like the Federal Bureau of Investigation, Centers for Disease Control and Prevention, and Immigration and Naturalization Services were not sufficient or appropriately specialized to address this new threat.

At the time of this writing, the science of counterterrorism is also still unfolding. The government has begun to engage the country’s research community through grants and collaborative opportunities, but across the sciences, and in statistics, the interesting problems and viable methodologies are still in a very speculative stage. Speculative is also exciting, though. Researchers feel lucky to be able to help define the landscape of a new research enterprise. This book encompasses a range of approaches to new problems and new problem spaces. The book is divided into four sections pertinent to counterterrorism: game theory, biometric authentication, syndromic surveillance, and modeling. Some of the chapters take a broad approach to defining issues in the specific research area, providing a more general overview. Other chapters provide detailed case studies and applications. Together they represent the current state of statistical sciences in the area of counterterrorism.

Game theory has long been seen as a valuable tool for understanding possible outcomes between adversaries. It played an important role in cold war decision and policymaking, but the opening section of this book rethinks game theory for the age of terrorism. In a world of asymmetric warfare, where your adversary is not a country with national assets and citizens at risk in the event of retaliation, the stakes are different. The section on game theory presented in this text provides an overview of statistical research issues in game theory and two articles that look specifically at game theory and risk analysis.

Biometric authentication has become a more prominent research area since 9/11 because of increased interest in security measures at border entry stations and other locations. Authentication of fingerprints, faces, retinal scans, etc., is usually an issue in the context of identity verification, i.e., does this passport match the person in front of me who is trying to use it? Beyond the logistics of collecting the information on everyone who applies for a passport or visa, storing it on the identity documents in a retrievable form, upgrading the computer equipment at all border crossings, and training border police to use the new technology, the issues of accurate identification are still to be worked out. Security agencies would also like to be able to use face recognition to pick known terrorists or criminals out of crowds using video cameras and real-time analysis software. The stakes for false positives are high — a man suspected as a potential terrorist bomber was held down by police and shot in the head in the London subway in 2005, and many individuals have wound up in long-term detention under the mere suspicion that they were members of terrorist organizations. Current technological shortcomings also have strong cultural implications: fingerprint authentication works less well with laborers who have worn skin and calluses on their hands; retinal scans work better with blue eyes than with brown. The section on biometric authentication in this book provides an overview of the history of its use with law enforcement and the courts and outlines some of the challenges faced by statisticians developing methods in this area. The two papers both address reducing error rates, specifically for authentication, although there are a myriad of other applications.

Syndromic surveillance has long been an issue of interest for biostatisticians, epidemiologists, and public health experts. After 9/11, however, more government funding became available to study issues related to sudden outbreaks of infectious diseases that might be the result of bioterrorism. Traditionally, research in this area would have looked at things like normal seasonal influenza cases, perhaps with an eye to preparing for possible flu pandemics caused by more virulent strains. But in the case of a bioterrorist incident, the concerns are a little different. For example, you want to be able to detect an outbreak of smallpox or cluster of anthrax infections as soon as possible so you can begin to respond. This may involve collecting and monitoring new data sources in near real-time: hospital admissions of patients with unusual symptoms, spikes in over-the-counter sales of cold medicines, etc. Collecting, integrating, and analyzing such new types of data involves the creation of new infrastructure and new methodologies. The section in this book on syndromic surveillance provides an overview of challenges and research issues in this growing area and includes articles on monitoring multiple data streams, evaluating statistical surveillance methods, and the spatiotemporal syndromic analysis.

Modeling is the bread and butter for many working statisticians and naturally is being applied to address issues in counterterrorism. Many of the speculative questions researchers and decision-makers have about terrorism

can be more practically and efficiently tested in computer models as opposed to actual physical experiments. As the section overview points out, “we cannot expose a population to a disease or chemical attack and see what happens.” This overview highlights the main issues addressed in the section and suggests future research directions. The section includes articles on developing large disease simulations, analyzing distributed databases, modeling of the concentration field in a building following release of a contaminant, and modeling the sensitivity of radiation detectors that might be deployed to screen cargo.

We would like to thank David Banks for suggesting this monograph, Sallie Keller-McNulty and Nancy Spruill for their ongoing support, and Hazel Kutac for her tireless editorial and production work.

Reference

- [KH04] Kean, T. H., and L. Hamilton. 2004. *The 9/11 commission report*. Washington, DC: National Commission on Terrorist Attacks upon the United States.

Los Alamos, NM
Los Alamos, NM
Monterey, CA

Alyson Wilson
Gregory Wilson
David Olwell

January 2006

Contents

Part I Game Theory

Game Theory in an Age of Terrorism: How Can Statisticians Contribute?

Ronald D. Fricker, Jr. 3

Combining Game Theory and Risk Analysis in Counterterrorism: A Smallpox Example

David L. Banks, Steven Anderson 9

Game-Theoretic and Reliability Methods in Counterterrorism and Security

Vicki Bier 23

Part II Biometric Authentication

Biometric Authentication

David L. Banks 43

Towards Statistically Rigorous Biometric Authentication Using Facial Images

Sinjini Mitra 47

Recognition Problem of Biometrics: Nonparametric Dependence Measures and Aggregated Algorithms

Andrew L. Rukhin 81

Part III Syndromic Surveillance

Data Analysis Research Issues and Emerging Public Health Biosurveillance Directions <i>Henry Rolka</i>	101
Current and Potential Statistical Methods for Monitoring Multiple Data Streams for Biosurveillance <i>Galit Shmueli, Stephen E. Fienberg</i>	109
Evaluating Statistical Methods for Syndromic Surveillance <i>Michael A. Stoto, Ronald D. Fricker, Jr., Arvind Jain, Alexis Diamond, John O. Davies-Cole, Chevelle Glymph, Gebreyesus Kidane, Garrett Lum, LaVerne Jones, Kerda Dehan, and Christine Yuan</i>	141
A Spatiotemporal Analysis of Syndromic Data for Biosurveillance <i>Laura Forsberg, Caroline Jeffery, Al Ozonoff, Marcello Pagano</i>	173
<hr/>	
Part IV Modeling	
<hr/>	
Modeling and Simulation for Defense and National Security <i>Wendy L. Martinez</i>	195
Modeling and Parameterization for a Smallpox Simulation Study <i>Sarah Michalak, Gregory Wilson</i>	199
Approaches to Modeling the Concentration Field for Adaptive Sampling of Contaminants during Site Decontamination <i>William K. Sieber, James S. Bennett, Abera Wouhib, Joe Fred Gonzalez, Jr., Myron J. Katzoff, Stanley A. Shulman</i>	215
Secure Statistical Analysis of Distributed Databases <i>Alan F. Karr, Xiaodong Lin, Ashish P. Sanil, Jerome P. Reiter</i>	237
Statistical Evaluation of the Impact of Background Suppression on the Sensitivity of Passive Radiation Detectors <i>Tom Burr, James Gattiker, Mark Mullen, George Tompkins</i>	263
Index	281

Game Theory

Game Theory in an Age of Terrorism: How Can Statisticians Contribute?

Ronald D. Fricker, Jr.

Department of Operations Research, Naval Postgraduate School,
rdfricker@nps.edu

In *The Law of Loopholes in Action* [Gel05], David Gelernter argues that “every loophole will eventually be exploited; every loophole will eventually be closed.” His thesis applied to terrorism means that terrorists will find security loopholes via continual exploration and that, once discovered, specific defensive measures have to be put in place to close each loophole.

The net effect of the Law of Loopholes, as anyone who flies regularly today knows, is an ever-expanding set of security rules and requirements. Such rules and requirements are useful for helping prevent the reoccurrence of a particular type of incident. But, when a determined adversary’s focus is on causing general destruction and mayhem, then as one loophole is plugged, the adversary simply shifts its attention and energies to looking for and trying to exploit a different loophole.

The problem, of course, is that it is impossible to defend all potential targets (and their associated loopholes) against all threats all of the time. While it *is* important to implement certain new and improved defensive tactics, precisely because it is impossible to protect everything at all times, it is equally as important (and arguably more important) to implement offensive strategies to deter and disrupt these adversaries.

The question is, how to identify effective offensive and defensive strategies and tactics?

One approach is through the use of *game theory*, the mathematically based study and analysis of adversarial conflicts. The classic text *The Compleat Strategyst* [Wil66] characterizes *games of strategy* as having the following characteristics:

- A conflict: the participants (e.g., individuals, organizations, countries; known as “players” in game theory parlance) are at cross-purposes or have opposing interests.
- Adversarial reaction and interaction: each player has some control over the course of the conflict or its outcome via one or more decisions.

- Outside forces: some aspects of the conflict are outside of the players' control and may be governed by chance or are unknown.

These characteristics clearly apply to the problem of thwarting terrorists and defeating terrorism.

The first extensive treatment of game theory was *Theory of Games and Economic Behavior* by John von Neumann and Oskar Morgenstern [VM44] in 1944. The seminal work on the subject, “Zur Theorie der Gesellschaftsspiele” by von Neumann [von28], was written in 1928. John von Neumann characterized the difference between games such as chess and games of strategy by saying “Chess is not a game. Chess is a well-defined form of computation. You may not be able to work out the answers, but in theory there must be a solution, a right procedure in any position. Now real games are not like that at all. Real life is not like that. Real life consists of bluffing, of little tactics of deception, of asking yourself what is the other man going to think I meant to do. And that is what games are about in my theory” [Pou92].

Game theoretic methods provide a structured way to examine how two adversaries will interact under various conflict scenarios. The results often provide insight into why real-world adversaries behave the way they do. In the middle and late 20th century, a great deal of game theoretic research focused on analyzing the arms race, nuclear brinkmanship, and Cold War strategies [Pou92]. While in the pre-9/11 era, game theory was also applied to terrorism, post-9/11 this work has expanded [SA03].

1 Game Theory Applied to Terrorism

In what is surely a gross oversimplification of the field (apologies to game theorists in advance), there are three broad categories of game theoretic methods applicable to the analysis of terrorism:

1. *Classic games* can generally be illustrated in a tabular form in which the players, their strategies, and their “payoffs” are completely specified. These types of games are often studied to determine whether there are a pair of strategies that result in an equilibrium between the two players (a “saddle point”) and how the players will behave given the existence or absence of a saddle point.
2. *Repetitive (or repeated) games*, which are games that occur over time and the opponents repeatedly interact in a series of conflicts. These games are studied to gain insight into how players behave and react to their opponent’s behavior and which behavioral strategies result in favorable or unfavorable final outcomes.
3. *Tabletop games* consisting of the simulation of an adversarial interaction with two or more actual (human) players using rules, data, and procedures designed to depict a conflict. “Tabletop” refers to the manner of older war games in which a battle was played out using miniature markers and

maps on a table, much like the board game Risk. These types of games are generally less structured than the previous types, meaning the players have a much larger set of strategies available than can be easily tabularized.

Recent applications of game theoretic methods to the study of terrorism include: assessing strategies for how nations allocate expenditures for terrorism deterrence and the resulting implications for being attacked [AST87, SL68]; measures evaluating how various military employment policies/strategies encourage or discourage states from sponsoring terrorism [Art04]; assessing insurance risks via models that explicitly account for malicious terrorist intent [Maj02]; determining whether or not a stated policy of nonnegotiation with terrorist hostage-takers deters such behavior and under what conditions [LS88]; and evaluating the effects of focusing national antiterrorism policy on deterrence or prevention [SA03].

2 Statistics and Game Theory

In the parlance of game theory, much of classical statistics is a “one-person game” because there is no adversary. Classic statistical problems, particularly inferential problems, concern the estimation of an unobserved parameter or parameters. In these problems, the “adversary” is nature, manifested as randomness in some form or another, not as a willful opponent.

A frequent assumption in statistical methods, analyses, and models is that the parameter or population under study is fixed and the most important uncertainty to quantify is that which comes from sampling variability. Even in those problems where the parameter may change over time, the usual assumption is that the underlying mechanism that generates an outcome is unaffected by that outcome. (For example, in a regression model we assume the dependent variable does not or cannot affect the independent variable.) Neither of these assumptions is likely to be true in a game theory problem, where the population of interest is an intelligent adversary capable of changing its form, tactics, and responses.

The upshot is that most statisticians are not used to thinking about problems such as those addressed by game theory. However, statisticians *are* used to addressing problems in which uncertainty is either a natural component or must be quantified, and there is a lot of uncertainty in game theoretic models about deterring, detecting, and thwarting terrorists.

3 How Can Statisticians Contribute?

Game theoretic models tend to be fairly abstract models of reality. This has not prevented the models from providing useful insights into strategies for addressing certain types of conflicts, but it does lead to two specific questions:

1. How well do the models fit observed data?
2. How can model uncertainty be quantified?

Both are questions that statisticians are well-suited to help address.

Possible ways statisticians could contribute to the further development of game theoretic methods, both in general and for terrorist problems in particular, include the following.

- Game theory models, including the strategies and their payoffs, are often defined in an ad hoc manner using expert judgment. A relevant statistical question is, how might data from past incidents and other knowledge be used to *infer* either the terrorist's "game" or the strategies they perceive or prefer? That is, how might a game be "fit" to observed data?
- The payoffs in game theory are utilities representing the desirability of the various outcomes to the players. In the absence of information, the utilities are often simply rankings of the various outcomes. A better methodology would be to elicit utilities from policymakers or subject-matter experts, much like one might elicit prior probabilities for a Bayesian analysis. Relevant questions include, what is (are) the best way(s) to elicit the utilities and how should utilities from multiple experts be combined?
- Once the payoffs are specified, the analysis of a game often treats them as fixed and known. How might the games be created, analyzed, and evaluated so that the uncertainty in payoffs is accounted for in the results, including the specification of the optimal strategy?
- Tabletop games are often useful for developing new insights and/or out-of-the-box potential strategies, but they also often can only explore a small portion of the "game space." Relevant questions include how to characterize and account for the uncertainty in game design (e.g., a terrorist opponent's capabilities) and how statistical methods might be used to help design a series of games to best explore the "capabilities/strategy space."
- Finally, for new types of games that incorporate uncertainty, as well as for a set or series of more traditional games, how can graphical methods be employed to best display important game results, including appropriate depictions of uncertainty and variability?

The two chapters that follow this one discuss and examine how risk analysis can be combined with game theory. In "Combining Game Theory and Risk Analysis in Counterterrorism: A Smallpox Example," Banks and Anderson describe how to use risk analysis to generate random payoff matrices, which are then used to estimate the probability that a given strategy is optimal. In "Game-Theoretic and Reliability Methods in Counterterrorism and Security," Bier discusses the literature on reliability and risk analytic methods for rare events, game theory, and approaches for combining the two methods for defending complex systems against terrorist attack.

These two efforts represent a promising start towards addressing some of the problems described above. Yet more remains to be done.

References

- [Art04] Arthur, K. 2004. "Understanding the military's role in ending state-sponsored terrorism." Master's thesis, Naval Postgraduate School.
- [AST87] Atkinson, S. E., T. Sandler, and J. T. Tschirhart. 1987. "Terrorism in a bargaining framework." *Journal of Law and Economics* 30:1–21.
- [Gel05] Gelernter, D. 2005. "The law of loopholes in action." *Los Angeles Times*, B13, May 6.
- [LS88] Lapan, H. E., and T. Sandler. 1988. "To bargain or not to bargain: That is the question." *American Economic Review* 78:16–20.
- [Maj02] Major, J. A. 2002. Advanced techniques for modeling terrorism risk. National Bureau of Economic Research Insurance Group Conference, February 1, 2002. http://www.guycarp.com/portal/extranet/pdf/major_terrorism.pdf, downloaded on June 30, 2005.
- [Pou92] Poundstone, W. 1992. *Prisoner's dilemma*. New York: Doubleday.
- [SA03] Sandler, T., and D. G. Arce M. 2003. "Terrorism and game theory." *Simulation and Gaming* 34:319–337.
- [SL68] Sandler, T., and H. E. Lapan. 1968. "The calculus of dissent: An analysis of terrorists' choice of targets." *Synthese* 76:245–261.
- [von28] von Neumann, J. 1928. "Zur Theorie der Gesellschaftsspiele." *Mathematische Annalen* 100:295–300.
- [VM44] von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- [Wil66] Williams, J. D. 1966. *The compleat strategist*, rev. ed. New York: McGraw-Hill Book Company.

Combining Game Theory and Risk Analysis in Counterterrorism: A Smallpox Example

David L. Banks¹ and Steven Anderson²

¹ Institute of Statistics and Decision Sciences, Duke University,
banks@stat.duke.edu

² Center for Biologics Evaluation and Research, U. S. Food and Drug
Administration, AndersonSt@cber.fda.gov

1 Introduction

The U.S. government wishes to invest its resources as wisely as possible in defense. Each wasted dollar diverts money that could be used to harden crucial vulnerabilities, prevents investment in future economic growth, and increases taxpayer burden. This is a classic conflict situation; a good strategy for the player with fewer resources is to leverage disproportionate resource investment by its wealthy opponent. That strategy rarely wins, but it makes the conflict sufficiently debilitating that the wealthy opponent may be forced to consider significant compromises.

Game theory is a traditional method for choosing resource investments in conflict situations. The standard approach requires strong assumptions about the availability of mutual information and the rationality of both opponents. Empirical research by many people [KT72] shows that these assumptions fail in practice, leading to the development of modified theories with weaker assumptions or the use of prior probabilities in the spirit of Bayesian decision theory. This paper considers both traditional game theory (minimax solution for a two-person, zero-sum game in normal form) and also a minimum expected loss criterion appropriate for extensive-form games with prior probabilities. However, we emphasize that for terrorism, the zero-sum model is at best an approximation; the valuation of the wins and the losses is likely to differ between the opponents.

Game theory requires numerical measures of payoffs (or losses) that correspond to particular sets of decisions. In practice, those payoffs are rarely known. Statistical risk analysis allows experts to determine reasonable probability distributions for the random payoffs. This paper shows how risk analysis can support game theory solutions and how Monte Carlo methods provide insight into the optimal game theory solutions in the presence of uncertainty about payoffs.

Our methodology is demonstrated in the context of risk management for a potential terrorist attack using the smallpox virus. The analysis we present here is a simplified version that aims at methodological explanation rather than analysis or justification of specific healthcare policies. As a tabletop exercise, the primary aim is only to provide a blueprint for a more rigorous statistical risk analysis. The underlying assumptions, modeling methods used here, and any results or discussion of the modeling are based on preliminary and unvalidated data and do not represent the opinion of the Food and Drug Administration (FDA), the Department of Health and Human Services, or any branch of the U.S. government.

2 Game Theory for Smallpox

The smallpox debate in the United States has focused upon three kinds of attack and four kinds of defense. The three attack scenarios suppose that there might be:

- No smallpox attack,
- A lone terrorist attack on a small area (similar to the likely scenario for the anthrax letters), or
- A coordinated terrorist attack upon multiple population centers.

The four defense scenarios that have been publicly considered by U.S. agency officials are:

- Stockpile smallpox vaccine,
- Stockpile vaccine and develop biosurveillance capabilities,
- Stockpile vaccine, develop biosurveillance, and inoculate key personnel, and
- Provide mass vaccination to nonimmunocompromised citizens in advance.

Although there are many refinements that can be considered for both the attack and the defense scenarios, these represent the possibilities discussed in the public meetings held in May and June 2002 [McK02].

Suppose that analysts used game theory as one tool to evaluate potential defense strategies. Then the three kinds of attack and four kinds of defense determine a classic normal-form payoff matrix for the game (Table 1).

Table 1. Attack—defense cost matrix

	No Attack	Single Attack	Multiple Attack
Stockpile Vaccine	C_{11}	C_{12}	C_{13}
Biosurveillance	C_{21}	C_{22}	C_{23}
Key Personnel	C_{31}	C_{32}	C_{33}
Everyone	C_{41}	C_{42}	C_{43}

The C_{ij} entries are the costs (or payoffs) associated with each combination of attack and defense, and we have used abbreviated row and column labels to identify the defenses and attacks, respectively, as described before.

For each of the 12 attack–defense combinations, there is an associated cost. These costs may include dollars, human lives, time, and other resources. For our calculation, all of these costs are monetized, according to principles detailed in Sect. 3. The monetized value of a human life is set to \$2.86 million, following the Department of Transportation’s figures for cost–benefit analyses of safety equipment.

Note that there is very large uncertainty in the C_{ij} values. Portions of the cost (e.g., those associated with expenses already entailed) may be known, but the total cost in each cell is a random variable. These random variables are not independent, since components of the total cost are common to multiple cells. Thus it is appropriate to regard the entire game theory table as a multivariate random variable whose joint distribution is required for a satisfactory analysis that propagates uncertainty in the costs through to uncertainty about best play.

Classical game theory [Mye91, Chap. 3] determines the optimal strategies for the antagonists via the minimax theorem. This theorem asserts that for any two-person cost matrix in a strictly competitive game (which is the situation for our example), there is an equilibrium strategy such that neither player can improve their expected payoff by adopting a different attack or defense. This equilibrium strategy may be a pure strategy, in which case optimal play is a specific attack–defense pair. This happens when the attack that maximizes the minimum damage and the defense that minimizes the maximum damage coincide in the same cell. Otherwise, the solution is a mixed strategy, in which case the antagonists pick attacks and defenses according to a probability distribution that must be calculated from the cost matrix. There may be multiple equilibria that achieve the same expected payoff, and for large matrices it can be difficult to solve the game.

Alternatively, one can use Bayesian decision theory to solve the game. Here a player puts a probability distribution over the actions of the opponent, and then chooses their own action so as to minimize the expected cost [Mye91, Chap. 2]. Essentially, one just multiplies the cost in each row by the corresponding probability, sums these by row, and picks the defense with the smallest sum. This formulation is easier to solve, but it requires one to know or approximate the opponent’s probability distribution, and it does not take full account of the mutual strategic aspects of adversarial games (i.e., the assigned probabilities need not correspond to any kind of “if I do this, then he’ll do that” reasoning). Bayesian methods are often used in extensive-form games, where players make their choices over time, conditional on the actions of their opponent.

In developing our analysis of the smallpox example we make two assumptions about time. First, we use only the information available by June 1, 2002; subsequent information on the emerging program costs is not included. This

keeps the analysis faithful in spirit to the decision problem actually faced by U.S. government policymakers in the spring of 2002 (their initial plan was universal vaccination, but ultimately they chose the third scenario with stockpiling, biosurveillance, and very limited vaccination of some first responders). Second, all of the estimated cost forecasts run to October 1, 2007. The likelihood of changing geopolitical circumstances makes it unrealistic to attempt cost estimates beyond that fiscal year.

3 Risk Analysis for Smallpox

Statistical risk analysis is used to estimate the probability of undesirable situations and their associated costs. In the same way that it is used in engineering (e.g., for assessing nuclear reactor safety [Spe85]) or the insurance industry (e.g., for estimating the financial costs associated with earthquakes in a specific area [Bri93]), this paper uses risk analysis to estimate the costs associated with different kinds of smallpox attack/defense combinations.

Risk analysis involves careful discussions with domain experts and structured elicitation of their judgments about probabilities and costs. For smallpox planning, this requires input from physicians, public health experts, mathematical epidemiologists, economists, emergency response administrators, government accountants, and other kinds of experts. We have not conducted the in-depth elicitation from multiple experts in each area that is needed for a fully rigorous risk analysis; however, we have discussed the cost issues with representatives from each area, and we believe that the estimates in this section are sufficiently reasonable to illustrate, qualitatively, the case for combining statistical risk analysis with game theory for threat management in the context of terrorism.

Expert opinion was typically elicited in the following way. Each expert was given a written document with background on smallpox epidemiology and a short description of the attacks and defenses considered in this paper. The expert often had questions; these were discussed orally with one of the authors and, to the extent possible, resolved on the basis of the best available information. Then the expert was asked to provide a point estimate of the relevant cost or outcome and the range in which that value would be expected to fall in 95% of similar realizations of the future. If these values disagreed with those from other experts, then the expert was told of the discrepancy and invited to alter their opinion. Based on point estimate and the range, the authors and the expert chose a distribution function with those parameters, which also respected real-world requirements for positivity, integer values, known skew, or other properties. As the last step in the interview, the expert was given access to all the other expert opinions obtained to that point and asked if there were any that seemed questionable; this led, in one case, to an expert being recontacted and a subsequent revision of the elicitation. But it should be emphasized that these interviews were intended to be short and did

not use the full range of probes, challenges, and checks that are part of serious elicitation work.

The next three subsections describe the risk analysis assumptions used to develop the random costs for the first three cells (C_{11} , C_{21} , C_{31}) in the game theory payoff matrix. Details for developing the costs in the other cells are available from the authors. These assumptions are intended to be representative, realistic, and plausible, but additional input by experts could surely improve upon them. Many of the same costs arise in multiple cells, introducing statistical dependency among the entries. (That is, if a given random payoff matrix assumes an unusually large cost for stockpiling in one cell of the random table, then the same high value should appear in all other cells in which stockpiling occurs.)

3.1 Cell (1,1): Stockpile Vaccine/No Attack Scenario

Consider the problem of trying to estimate the costs associated with the (1,1) cell of the payoff matrix, which corresponds to no smallpox attack and the stockpiling of vaccine. This estimate involves combining costs with very different levels of uncertainty.

At the conceptual level, the cost C_{11} is the sum of four terms:

$$C_{11} = ET_{\text{dry}} + ET_{\text{Avent}} + ET_{\text{Acamb}} + \text{VIG} + \text{PHIS},$$

where ET_{dry} and ET_{Avent} are the costs of efficacy and safety testing for the Dryvax and Aventis vaccines, respectively; ET_{Acamb} is the cost of new vaccine production and testing from Acambis; VIG is the cost of producing sufficient doses of vaccinia immune globulin to treat adverse reactions and possible exposures; and PHIS is the cost of establishing the public healthcare infrastructure needed to manage this stockpiling effort.

There is no uncertainty about ET_{Acamb} ; the contract fixes this cost at \$512 million. But there is substantial uncertainty about ET_{dry} and ET_{Avent} since these entail clinical trials and may require follow-on studies; based on discussions with experts, we believe these costs may be realistically modeled as independent uniform random variables, each ranging between \$2 and \$5 million. There is also large uncertainty about the cost for producing and testing sufficient doses of VIG to be prepared for a smallpox attack; our discussions suggest this is qualitatively described by a normal random variable with mean \$100 million and a standard deviation of \$20 million. There is great uncertainty about PHIS (which includes production of bifurcated inoculation needles, training, storage costs, shipment readiness costs, etc.). Based on the five-year operating budget of other government offices with analogous missions, we assume this cost is normally distributed with mean \$940 million and standard deviation \$100 million.

3.2 Cell (2,1): Biosurveillance/No Attack Scenario

Biosurveillance programs are being piloted in several major metropolitan areas. These programs track data, on a daily basis, from emergency room admission records to quickly discover clusters of disease symptoms that suggest bioterrorist attack. Our cost estimates are based upon discussions with the scientists working in the Boston area [RKD02] and with the Pittsburgh team that developed monitoring procedures for the Salt Lake City Olympic games.

The cost C_{21} includes the cost C_{11} since this defense strategy uses both stockpiling of vaccine and increased biosurveillance. Thus

$$C_{21} = C_{11} + \text{PHIB} + \text{PHM} + \text{NFA} \times \text{FA},$$

where PHIB is the cost of the public health infrastructure needed for biosurveillance, including the data input requirements and software; PHM is the cost of a public health monitoring center, presumably at the Centers for Disease Control and Prevention, that reviews the biosurveillance information on a daily basis; NFA is the number of false alarms from the biosurveillance system over five years of operation; and FA is the cost of a false alarm.

For this exercise, we assume that PHIB is normally distributed with mean \$900 million and standard deviation \$100 million (for a five-year funding horizon); this is exclusive of the storage, training, and other infrastructure costs in PHIS, and it includes the cost of hospital nursing-staff time to enter daily reports on emergency room patients with a range of disease symptoms (not just those related to smallpox). PHM is modeled as a normal random variable with mean \$20 million and standard deviation \$4 million (this standard deviation was proposed by a federal administrator and may understate the real uncertainty). False alarms are a major problem for monitoring systems; it is difficult to distinguish natural contagious processes from terrorist attacks. We expect about one false alarm per month over five years in a national system of adequate sensitivity, and thus FA is taken to be a Poisson random variable with mean 60. The cost for a single false alarm is modeled as a normal random variable with mean \$500,000 and standard deviation \$100,000.

3.3 Cell (3,1): Key Personnel/No Attack Scenario

One option, among several possible policies that have been discussed, is for the United States to inoculate about 500,000 key personnel, most of whom would be first-responders in major cities (i.e., emergency room staff, police, and public health investigators who would be used to trace people who have come in contact with carriers). If chosen, this number is sufficiently large that severe adverse reactions become a statistical certainty.

The cost of this scenario subsumes the cost C_{21} of the previous scenario, and thus

$$C_{31} = C_{21} + \frac{\text{NKP} \times \text{IM}}{25,000} + \text{PAE} \times \text{NKP} \times \text{AEC},$$

where NKP is the number of key personnel; IM is the cost of the time and resources needed to inoculate 25,000 key personnel and monitor them for adverse events; PAE is the probability of an adverse event; and AEC is the average cost of one adverse event.

We assume that NKP is uniformly distributed between 400,000 and 600,000 (this reflects uncertainty about how many personnel would be designated as “key”). The IM is tied to units of 25,000 people, since this is a one-time cost and represents the number of people that a single nurse might reasonably inoculate and maintain records upon in a year. Using salary tables, we approximate this cost as a normal random variable with mean \$60,000 and standard deviation \$10,000.

The probability of an adverse event is taken from Anderson [And02], which is based upon Lane et al. [LRN70]; the point estimate for all adverse events is 0.293, but since there is considerable variation and new vaccines are coming into production, we have been conservative about our uncertainty and assumed that the probability of an adverse event is uniformly distributed between 0.15 and 0.45. Of course, most of these events will be quite minor (such as local soreness) and would not entail any real economic costs.

The AEC is extremely difficult to estimate. For purposes of calculation, we have taken the value of a human life to be \$2.86 million (the amount used by the National Highway Transportation Safety Administration in cost-benefit analyses of safety equipment). But most of the events involve no cost, or perhaps a missed day of work that has little measurable impact on productivity. After several calculations and consultations, this analysis assumes that AEC can be approximated as a gamma random variable with mean \$40 and standard deviation \$100 (this distribution has a long right tail).

4 Analysis

The statistical risk analysis used in Sect. 3, albeit crude, shows how expert judgment can generate the random payoff matrices. The values in the cells of such tables are not independent, since many of the cost components are shared between cells. In fact, it is appropriate to view the table as a matrix-valued random variable with a complex joint distribution.

Random tables from this joint distribution can be generated by simulation. For each table, one can apply either the minimax criterion to determine an optimal strategy in the sense of von Neumann and Morgenstern [VM44], or a minimum expected loss criterion to determine an optimal solution in the sense of Bayesian decision theory [Mye91, Chap. 2]. By doing this repeatedly, for many different random tables, one can estimate the proportion of time that each defense strategy is superior.

Additionally, it seems appropriate to track not just the number of times a defense strategy is optimal, but also weight this count by some measure of the difference between the costs of the game under competing defenses. For

example, if two defenses yield game payoffs that differ only by an insignificant amount, it seems unrealistic to give no credit to the second-best strategy. For this reason we also use a scoring algorithm in which the score a strategy receives depends upon how well separated it is from the optimal strategy. Specifically, suppose that defense strategy i has value V_i on a given table. Then the score S_i that strategy i receives is

$$S_i = 1 - \frac{V_i}{\max V_j},$$

and this ensures that strategies are weighted to reflect the magnitude of the monetized savings that accrue from using them. The final rating of the strategies is obtained by averaging their scores from many random tables.

4.1 Minimax Criterion

We performed the simulation experiment described above 100 times and compared the four defense strategies in terms of the minimax criterion. Although one could certainly do more runs, we believe that the approximations in the cost modeling are so uncertain that additional simulation would only generate spurious accuracy.

Among the 100 runs, we found that the Stockpile strategy won 9 times, the Biosurveillance strategy won 24 times, the Key Personnel strategy won 26 times, and the Vaccinate Everyone strategy won 41 times. This lack of a clear winner may be, at some intuitive level, the cause of the widely different views that have been expressed in the public debate on preparing for a smallpox attack.

If one uses scores, the results are even more ambiguous. The average score for the four defense strategies ranged between 0.191 and 0.326, indicating that the expected performances were, on average, quite similar.

From public policy standpoint, this may be a fortunate result. It indicates that in terms of the minimax criterion, any decision is about equally defensible. This gives managers flexibility to incorporate their own judgment and to respond to extra scientific considerations.

4.2 Minimum Expected Loss Criterion

The minimax criterion may not be realistic for the game theory situation presented by the threat of smallpox. In particular, the normal-form game assumes that both players are ignorant of the decision made by their opponent until committed to a course of action. For the smallpox threat, there has been a vigorous public discussion on what preparations the United States should make. Terrorists know what the United States has decided to do, and presumably this will affect their choice of attack. Therefore the extensive-form version of game theory seems preferable. This form can be thought of as a

decision tree, in which players alternate their moves. At each stage, the player can use probabilistic assessments about the likely future play of the opponent.

The minimum expected loss criterion requires more information than does the minimax criterion. The analyst needs to know the probabilities of a successful smallpox attack conditional on the United States selecting each of the four possible defenses. This is difficult to determine, but we illustrate how one can do a small sensitivity analysis that explores a range of probabilities for smallpox attack.

Table 2 shows a set of probabilities that we treat as the baseline case. We believe it accords with a prudently cautious estimate of the threat of a smallpox attack.

Table 2. Baseline probabilities of attack for different defenses

	No Attack	Single Attack	Multiple Attack
Stockpile Vaccine	0.95	0.040	0.010
Biosurveillance	0.96	0.035	0.005
Key Personnel	0.96	0.039	0.001
Everyone	0.99	0.005	0.005

To interpret Table 2, it says that if the United States were to only stockpile vaccine, then the probability of no smallpox attack is 0.95, the probability of a single attack is 0.04, and the probability of multiple attacks is 0.01. Similarly, one reads the attack probabilities for other defenses across the row. All rows must sum to one.

The minimum expected loss criterion multiplies the probabilities in each row of Table 2 by the corresponding costs in the same row of Table 1, and then sums across the columns. The criterion selects the defense that has the smallest sum.

As with the minimax criterion, one can simulate many payoff tables and then apply the minimum expected loss criterion to each. In 100 repetitions, Stockpile won 96 times, Biosurveillance won 2 times, and Vaccinate Everyone won twice. The scores showed roughly the same pattern, strongly favoring the Stockpile defense.

We now consider two alternative sets of probabilities shown in Tables 3 and 4. Table 3 is more pessimistic and has larger attack probabilities. Table 4 is more optimistic and has smaller attack probabilities. A serious sensitivity analysis would investigate many more tables, but our purpose is illustration and we doubt that the quality of the assessments that underlie the cost matrix can warrant further detail.

For Table 3, 100 simulation runs found that Stockpile won 15 times, Biosurveillance won 29 times, Key Personnel won 40 times, and Vaccinate Everyone won 16 times. In contrast, for Table 4, the Stockpile strategy won 100 times in 100 runs. The scores for Table 3 ranged from 18.2 to 38.8, which are

Table 3. Pessimistic probabilities of attack for different defenses

	No Attack	Single Attack	Multiple Attack
Stockpile Vaccine	0.70	0.20	0.10
Biosurveillance	0.80	0.15	0.05
Key Personnel	0.85	0.10	0.05
Everyone	0.90	0.05	0.05

Table 4. Optimistic probabilities of attack for different defenses

	No Attack	Single Attack	Multiple Attack
Stockpile Vaccine	0.980	0.0100	0.0100
Biosurveillance	0.990	0.0050	0.0050
Key Personnel	0.990	0.0050	0.0050
Everyone	0.999	0.0005	0.0005

quite similar. In contrast, for Table 4 nearly all the weight of the score was on the Stockpile defense.

These results show that the optimal strategy is sensitive to the choice of probabilities used in the analysis. Determining those probabilities requires input from the intelligence community and the judgment of senior policymakers.

5 Conclusions

This paper has outlined an approach combining statistical risk analysis with game theory to evaluate defense strategies that have been considered for the threat of smallpox. We believe that this approach may offer a useful way of structuring generic problems in resource investment for counterterrorism.

The analysis in this paper is incomplete.

1. We have focused upon smallpox, because the problem has been framed rather narrowly and quite definitively by public discussion. But a proper game theory analysis would not artificially restrict the options of the terrorists, and should consider other attacks, such as truck bombs, chemical weapons, other diseases, and so forth (which would get difficult, but there may be ways to approximate). It can be completely misleading to seek a local solution, as we have done.
2. Similarly, we have not fully treated the options of the defenders. For example, heavy investment in intelligence sources is a strategy that protects against many different kinds of attacks and might well be the superior solution in a less local formulation of the problem.
3. We have not considered constraints on the resources of the terrorists. The terrorists have limited resources and can invest in a portfolio of different kinds of attacks. Symmetrically, the United States can invest in a portfolio

of defenses. This aspect of the problem is not addressed — we assume that both parties can fund any of the choices without sacrificing other goals.

4. The risk analysis presented here, as discussed previously, is not adequate to support public policy formulation.

Nonetheless, despite these limitations, the methodology has attractive features. First, it is easy to improve the quality of the result through better risk analysis. Second, it automatically raises issues that have regularly emerged in policy discussions. Third, it captures facets of the problem that are not amenable to either game theory or risk analysis on their own, because classical risk analysis is not used in adversarial situations and because classical game theory does not use random costs.

Appendix: Background on Smallpox

Although the probability that the smallpox virus (*Variola major*) might be used against the United States is thought to be small, the public health and economic impact of even a limited release would be tremendous. Any serious attack would probably force mass vaccination programs, causing additional loss of life due to adverse reactions. Other economic consequences could easily be comparable to those of the attacks of September 11, 2001.

A smallpox attack could potentially be initiated through infected humans or through an aerosol [HIB99]. In 12 to 14 days after natural exposure patients experience fever, malaise, body aches, and a body rash [FHA88]. During the symptomatic stages of the disease the patient can have vesicles in the mouth, throat, and nose that rupture to spread the virus during a cough or sneeze. Person-to-person spread usually occurs through inhalation of virus-containing droplets or from close contact with an infected person. As the disease progresses, the rash spreads to the head and extremities and evolves into painful, scarring vesicles and pustules. Smallpox has a mortality rate of approximately 30%, based on data from the 1960s and 1970s [Hen99].

Various mathematical models of smallpox spread exist and have been used to forecast the number of people infected under different exposure conditions and different public health responses [KCW02, MDL01]. There is considerable variation in the predictions from these models, partly because of differing assumptions about the success of the “ring vaccination” strategy that has been planned by the Centers for Disease Control and Prevention (CDC) [CDC02], and this is reflected in the public debate on the value of preemptive inoculation versus wait-and-see preparation. However, the models are in essential agreement that a major determinant of the size of the epidemic is the number of people who are exposed in the first attack or attacks.

The current vaccine consists of live vaccinia or cowpox virus and is effective at preventing the disease. Also, vaccination can be performed within the first 2 to 4 days postexposure to reduce the severity or prevent the occurrence of the disease [Hen99].

Vaccination is not without risk; the major complications are serious infections and skin disease such as progressive vaccinia, eczema vaccinatum, generalized vaccinia, and encephalitis. Approximately 12 people per million have severe adverse reactions that require extensive hospitalization, and about one-third of these die — vaccinia immune globulin (VIG) is the recommended therapy for all of these reactions except encephalitis. Using data from Lane et al. [LRN70], we estimate that 1 in 71,429 people suffer postvaccinial encephalitis, 1 in 588,235 suffer progressive vaccinia, 1 in 22,727 suffer eczema vaccinatum, and 1 in 3,623 suffer generalized vaccinia. Additionally, 1 in 1,656 people suffer accidental infection (usually to the eye) and 1 in 3,289 suffer some other kind of mild adverse event, typically requiring a person to miss a few days of work. Other studies give somewhat different numbers [NLP67a, NLL67b]. People who have previously been successfully vaccinated for smallpox are less likely to have adverse reactions, and people who are immunocompromised (e.g., transplant patients, those with AIDS) are at greater risk for adverse reactions [CDC02, Guide B, parts 3, 5, and 6].

Because the risk of smallpox waned in the 1960s, vaccination of the U.S. population was discontinued in 1972. It is believed that the effectiveness of a smallpox vaccination diminishes after about 7 years, but residual resistance persists even decades later. It has been suggested that people who were vaccinated before 1972 may be substantially protected against death, if not strongly protected against contracting the disease [Coh01].

The United States currently has about 15 million doses of the Wyeth Dryvax smallpox vaccine available. The vaccine was made by scarification of calves with the New York City Board of Health strain and fluid containing the vaccinia virus was harvested by scraping [RMK01]. Recent clinical trials on the efficacy of diluted vaccine indicate that both the five- and ten-fold dilutions of Dryvax achieve a take rate (i.e., a blister forms at the inoculation site, which is believed to be a reliable indicator of immunization) of at least 95%, so the available vaccine could be administered to as many as 150 million people should the need arise [FCT02, NIA02].

The disclosure by the pharmaceutical company Aventis [Ens02] of the existence in storage of 80 to 90 million doses of smallpox vaccine that were produced more than 30 years ago has added to the current stockpile. Testing is being done on the efficacy of the Aventis vaccine stock, including whether it, too, could be diluted if needed.

Contracts to make new batches of smallpox vaccine using cell culture techniques have been awarded to Acambis. The CDC amended a previous contract with Acambis in September 2001 to ensure production of 54 million doses by late 2002. Another contract for the production of an additional 155 million doses was awarded to Acambis in late November 2001, and the total cost of these contracts is \$512 million. After production, additional time may be needed to further test the safety and efficacy of the new vaccine [RMK01].

References

- [And02] Anderson, S. 2002. "A risk-benefit assessment of smallpox and smallpox vaccination." Technical Report, Office of Biostatistics and Epidemiology, Center for Biologics Evaluation and Research, U.S. Food and Drug Administration, Rockville, MD.
- [Bri93] Brillinger, D. R. 1993. "Earthquake risk and insurance." *EnviroMetrics* 4:1–21.
- [CDC02] Centers for Disease Control and Prevention. 2002. Smallpox response plan and guidelines version 3.0. <http://www.bt.cdc.gov/agent/smallpox/response-plan/index.asp>.
- [Coh01] Cohen, J. 2001. "Smallpox vaccinations: How much protection remains?" *Science* 294:985.
- [Ens02] Enserink, M. 2002. "New cache eases shortage worries." *Science* 296:25–26.
- [FHA88] Fenner, F., D. Henderson, I. Arita, Z. Jezek, and I. Ladnyi. 1988. *Smallpox and its eradication*. Geneva: World Health Organization.
- [FCT02] Frey, S. E., R. B. Couch, C. O. Tacket, J. J. Treanor, M. Wolff, F. K. Newman, R. L. Atmar, R. Edelman, C. M. Nolan, and R. B. Belshe. 2002. "Clinical responses to undiluted and diluted smallpox vaccine." *New England Journal of Medicine* 346 (17): 1265–1274.
- [Hen99] Henderson, D. A. 1999. "Smallpox: Clinical and epidemiological features." *Emerging Infectious Diseases* 5:537–539.
- [HIB99] Henderson, D. A., T. V. Inglesby, J. G. Bartlett, M. S. Ascher, E. Eitzen, P. B. Jahrling, J. Hauer, M. Layton, J. McDade, M. T. Osterholm, T. O'Toole, G. Parker, T. Perl, P. K. Russell, and K. Tonat. 1999. "Smallpox as a biological weapon — Medical and public health management." *Journal of the American Medical Association* 281 (22): 2127–2137.
- [KT72] Kahneman, D., and A. Tversky. 1972. "Subjective probability: A judgment of representativeness." *Cognitive Psychology* 3:430–454.
- [KCW02] Kaplan, E., D. Craft, and L. Wein. 2002. "Emergency response to a smallpox attack: the case for mass vaccination." *Proceedings of the National Academy of Sciences* 99:10935–10940.
- [LRN70] Lane, J. M., F. L. Ruben, J. M. Neff, and J. D. Millar. 1970. "Complications of smallpox vaccination, 1968: Results of ten statewide surveys." *Journal of Infectious Diseases* 122:303–309.
- [McK02] McKenna, M. A. 2002. "No mass smallpox vaccinations, panel recommends." *Atlanta Journal-Constitution*, June 21, p. 1.
- [MDL01] Meltzer, M. I., I. Damon, J. W. LeDuc, and J. D. Millar. 2001. "Modeling potential responses to smallpox as a bioterrorist weapon." *Emerging Infectious Diseases* 7:959–969.
- [Mye91] Myerson, R. B. 1991. *Game theory: Analysis of conflict*. Cambridge, MA: Harvard University Press.
- [NIA02] National Institute of Allergy and Infectious Diseases. March 28, 2002. "NIAID study results support diluting smallpox vaccine stockpile to stretch supply." *NIAID News*. <http://www.niaid.nih.gov/newsroom/releases/smallpox.htm>.
- [NLP67a] Neff, J. M., J. M. Lane, J. P. Pert, H. Moore, and J. Millar. 1967. "Complications of smallpox vaccination, I: National survey in the United States, 1963." *New England Journal of Medicine* 276:1–8.