Skip all front matter: Jump to Page 1

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

RAND HOMELAND SECURITY AND DEFENSE CENTER

Modeling Terrorism Risk to the Air Transportation System

An Independent Assessment of TSA's Risk Management Analysis Tool and Associated Methods

Andrew R. Morral, Carter C. Price, David S. Ortiz, Bradley Wilson, Tom LaTourrette, Blake W. Mobley, Shawn McKay, Henry H. Willis

Prepared for the Transportation Security Administration
Approved for public release; distribution unlimited

The research described in this report was prepared for the Transportation Security Administration and conducted within the RAND National Defense Research Institute, a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Navy, the Marine Corps, the defense agencies, and the defense Intelligence Community under Contract W74V8H-06-C-0002.

Library of Congress Control Number: 2012953351

ISBN: 978-0-8330-7685-4

Published 2012 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665
RAND URL: http://www.rand.org
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Summary

To support policy and resource allocation decisions, the Department of Homeland Security (DHS) and the Transportation Security Administration (TSA) have developed a suite of tools and processes for conducting risk assessments. One such tool is the Risk Management Analysis Tool (RMAT) developed by the Boeing Company and TSA in consultation with private sector and governmental members of a risk management working group. In December 2010, TSA asked RAND to evaluate whether RMAT provides results that are valid for TSA's risk-assessment needs. This report describes RAND's approach to this assessment and our findings.

RMAT simulates terrorist behavior and success in attacking vulnerabilities in the domestic commercial air transportation system. In doing so, it draws on estimates of terrorist resources, capabilities, preferences, decision processes, intelligence collection, and operational planning. It describes how the layers of security protecting the air transportation system are likely to perform in the face of a range of more than 60 types of attack. It draws on detailed blast and other physical modeling to understand the damage produced by different weapons and attacks and calculates the direct and indirect economic consequences of that damage. As such, the tool is designed to provide vital information that can help TSA understand the risks to which the entire air transportation system is exposed and develop ways to improve it.

RAND's approach to validating RMAT required first establishing TSA's intended uses and requirements for risk assessment and then evaluating which of those requirements RMAT can satisfy. We have

not validated RMAT against a set of original requirements for the system, because RMAT has evolved over time without such a guiding set of requirements. Instead, therefore, we evaluate TSA's current and broad risk-assessment requirements that RMAT can validly support.

We divided the validation effort into four substantive research questions:

1. Are the adversary behavior and air transportation system conceptual models valid, and are the data used to support them adequate? (Chapters Two and Three)
2. Are the sources and methods for populating the RMAT model with data sufficient to ensure their validity? (Chapter Four)
3. Does the RMAT code, as implemented, perform in the way it was designed to? (Chapter Five)
4. Can risk estimates from RMAT be used in the ways TSA intends? (Chapter Six)

These validation efforts considered diverse sources of evidence, including published scientific literature, elicited judgments from subject matter experts, considerations of logic and reasonableness, historical evidence, and quantitative empirical analysis of RMAT and its outputs. Each of these chapters includes detailed observations about RMAT strengths and weaknesses and concludes with a set of recommendations for further developing RMAT. In this Summary, we highlight only the main findings from the study.

RMAT Suitability for TSA Risk Assessment

RMAT is built around two innovative conceptual models: an adversary model that simulates adversary efforts to select, plan for, and execute attacks; and a defender model that simulates how the air transportation system will react to each attack.

As one of the first general theories of terrorism designed to account for adversary resource constraints, intelligence collection, targeting, utility functions, and operations, the RMAT adversary model represents a potentially important contribution to the terrorism research community. Nevertheless, it also necessarily relies on assumptions that are speculative and requires data as inputs that are subject to great uncertainty. As such, we do not believe that the adversary model should be regarded as likely to accurately anticipate terrorist behavior. We do think that it is useful analytically, however, such as for exploring how plausible characteristics and choices of adversaries might affect risk, refining analysts' understanding of the complexity of terrorist behavior, or helping to focus intelligence collection activities on features of terrorist behavior that RMAT identifies as important. Such uses are valuable. Just as with other terrorism risk models we have examined, however, these uses do not include prediction of the most likely actual adversary strategies.

TSA and Boeing recognized the limitations of the adversary model and have adopted strategies for working with RMAT that reduce reliance on it. RMAT can be run in such a way that adversaries are forced to attempt specific attacks, thereby circumventing some more speculative parts of the model that simulate adversary preferences and choice behavior. Even in this mode, however, the adversary must construct a plan, gather intelligence and resources, conduct dry runs, and perform other activities that affect the likelihood of success. Thus, results continue to depend on adversary modeling that is subject to major assumptions and important sources of uncertainty. The assumptions may be reasonable for some adversaries, in which case the results might be quite good. For others, the results may be quite poor. As such, the model results cannot be assumed to reliably anticipate system terrorism risks.

The defender model, which characterizes the U.S. domestic air transportation security system, is a particular strength of RMAT. The current abstract air transportation system modeled in RMAT appears to capture the key features relevant to security at most airports. If we have good information about an adversary's capabilities and intentions, the RMAT defender model can provide credible and useful estimates of the likelihood of detecting and interdicting the adversary. Moreover, modification of the generalized airport configuration is straightforward, so to the extent they fall within the scope of the RMAT "world," new places, processes, and vulnerabilities can be incorporated.

There are some gaps in the defender model that TSA and Boeing should consider remedying, as these limit its scope and validity and could introduce unwanted biases in RMAT results. Chief among these recommended improvements is broadening the scope of the domestic air transportation system to include its interfaces with non-U.S. airports and inbound foreign flights. Additional improvements to this portion of the model may include expanding the range of security threats considered under RMAT, expanding the range of attack pathways available to attackers, and inclusion of off-airport freight processing, catering, general aviation, mass transit, air traffic control, and booking information systems. Some of these recommended changes have already been planned by the RMAT development team.

Both the adversary and the defender models place heavy demands on the identification, validation, and maintenance of the roughly 4,300 input values quantifying aspects of airports, security operations, terrorists, attack outcomes, and their valuations. To fulfill model data requirements, Boeing and TSA have undertaken repeated data collection efforts that have relied on elicitations from subject matter experts; assessments of technical, red-teaming, and scientific data; review of TSA policies and procedures; and other data sources.

In reviewing a sample of RMAT data inputs, the RAND team was able to validate the reasonableness, if not necessarily the real-world correctness, for more than half of the parameter values either on logical grounds (e.g., nonmetallic knives are not explosive and do not contain any metal) or by confirming values through literature searches or consultation with subject matter experts. Nevertheless, some values appeared wrong to us, and others required estimates that call for information that either does not exist or is subject to such profound uncertainty that we judged they should not be estimated as point values but, rather, explicitly treated as ranges and sources of deep uncertainty affecting RMAT estimates. Boeing considers the specific parameters used in the model to be proprietary, so we are prohibited from illustrating this point with examples. Suffice it to say, however, that portions of the model attempt to parameterize quite specific features of terrorists' decision criteria, including specification of their risk tolerance, preferences, knowledge, and learning. Such RMAT variables are important in the model but require precision beyond what intelligence or academic research can credibly provide. Estimating ranges of values is often more plausible, but understanding the implications of these uncertainties on model results should be a priority in future work with the model.

RMAT uses subject matter experts as sources for roughly two-thirds of its input data. After reviewing TSA and Boeing methods for eliciting subject matter expert judgments, we offer several suggestions for improving the elicitation process and results. In addition, we recommend relying less on subject matter expert judgments when scientific or empirical literature is available as relevant input data and expanding the pool of experts used to provide judgments on RMAT conceptual models and parameter values.

To produce valid and useful results, RMAT requires more than just good conceptual models and valid input data; it needs code that faithfully characterizes the conceptual models and change management processes that help ensure that the code remains faithful through periodic modifications necessary for new case studies of risk, when improvements are made to the conceptual model, or when coding errors are corrected. To evaluate the RMAT software, we ran third-party software quality diagnostic tools on its code, we conducted sensitivity analysis experiments to establish whether input and output variables are associated in predictable ways, and we evaluated the change management processes used to maintain the software.

RMAT was originally developed as a prototype, and it has evolved continuously to fit new uses and requirements. The result is a complex program with less organization and efficiency than would be expected of a production model. Given the rapid pace of progress and changing requirements, it is easy to understand the software's current state, but it implies significant challenges for expanding, revising, debugging, testing, and managing the code—all of which threaten its ongoing reliability and validity. Boeing is aware of these code issues, of course, and reports that it is in the process of improving the RMAT source code. Our point here is not to criticize, because such complex and ambitious undertakings often exhibit these types of problems along the way, but to point out that moving forward to something stable, solid, and adaptable will be important but challenging.

Our sensitivity tests found that most, but not all, relationships between inputs and outputs are in the expected direction. Moreover, several relationships we expected to find were not present, such as associations between the probability of successfully entering the flight deck or the probability of air marshals being onboard and either the attractiveness or success of hijack attempts. We allowed the probability of federal air marshals being present on the hijacked aircraft to vary across a wide range, yet the model suggests that their presence or absence has no significant influence over hijack success rates—a result that is hard to understand.

Of the variables that appear most often to have a significant influence on outcomes, many are those we consider to be difficult or impossible to estimate with precision. These include judgments about how much perceived risk might color the decisions of current and future terrorists or how large the maximum possible size of terrorist cells might be when considering known and unknown groups.

These are all parameters that are subject to deep uncertainty and, no doubt, to wide variation across terrorist groups. Subject matter experts and intelligence analysts cannot credibly supply meaningful point estimates of these values. That these parameters also happen to explain a large portion of the variance in RMAT outcomes suggests the need for caution when interpreting model results based on rough estimates for these uncertain parameters.

Finally, our review finds that RMAT is capable of supporting several of TSA's risk-assessment intended uses but that its design is not always conducive to these purposes. TSA must make high-stakes resource allocation decisions designed to counter threats that are not well known, that are continuously evolving, and that may intelligently adapt to circumvent our security measures. This is a complex problem and one for which there may not be one best answer.

A recurring theme in our review is that it is a serious error to imagine that the correct values of all the parameters of a good model such as RMAT can be established. The analytic endeavor should be conceived differently so as to acknowledge the need for exploratory analysis under uncertainty. Because future terrorism risks are subject to sources of deep uncertainty, TSA should not seek security solutions that are optimized for a set of plausible and carefully collected input values. Instead, it should search for solutions that perform well across a range of input values selected to span the space of plausible future conditions. This type of exploratory analysis is invaluable for identifying robust solutions and understanding the conditions under which different solutions might be expected to perform well.

TSA has some procedures in place to explore the implications of key uncertainties. For instance, it uses sensitivity analysis methods to consider how risk reduction might vary with different RMAT assumptions about attacker capabilities, the expected number of attacks per year, or expectable improvements in technology. This approach provides insights into the robustness of the RMAT results to differing assumptions and represents an important advance in TSA approaches to understanding the risks it is charged with managing. But this approach still assumes that many other uncertain variables are estimated accurately.

A better approach, we believe, would use a simplified low-resolution terrorism risk model abstracted from RMAT to highlight the key sources of uncertainty affecting outcomes and then use exploratory analysis to evaluate the space of possible future outcomes using a spanning set of test cases or parameter values for all important and uncertain parameter values and assumptions. A similar approach is now reflected in Department of Defense planning (as in the 2009 *Quadrennial Defense Review*). In this report, we discuss how insights and parameter estimates from the RMAT might be used to support low-resolution models that could provide TSA with transparent analyses of the effects of deep uncertainties on risk and the decisions TSA must make.

Conclusions

RMAT has proven to be of great value to TSA in driving a more sophisticated understanding of terrorism risks to the air transportation system. Indeed, at the time RMAT was begun, TSA's approach to risk analysis and risk management was rudimentary. The process of developing RMAT led TSA to increasingly sophisticated understandings of the nature of terrorism threats, vulnerabilities, and consequences, as demonstrated in its current risk doctrine.

This is an example of one principal value that high-resolution models such as RMAT offer. Specifically, they can be invaluable for facilitating understanding of important phenomena and for recording, structuring, and conveying information that is complex and not well understood. Such models can become essential textbooks for training analysts and leaders to think more clearly and productively about complex phenomena and for driving developments in our theory and conceptual models for these phenomena. RMAT is clearly well suited for such purposes.

In addition, we find that RMAT fully or partially satisfies 16 of TSA's 19 high- and medium-priority risk-assessment requirements, making it a valuable addition to TSA's collection of analytic tools and methods.

As with all other terrorism risk models, however, it is not well suited for revealing how the future is likely to unfold. Even if the conceptual models on which RMAT is built were sound and comprehensive, the input data requirements exceed what subject matter experts or science can estimate with precision, and the imprecision of those estimates is subject to unknown sources and ranges of error. That said, we recommend that TSA make RMAT a component of a new exploratory and multiresolution modeling approach for supporting resource allocation and high-level policy questions.